

VI

HOW TO DO THINGS WITH PLANS

1. A FRESH START

From childhood on, we are taught that there is a crucial difference between what others think is right or wrong and what *is* right or wrong. Just because everyone does it does not mean that we should do it. We are repeatedly told that the rules of ethical behavior apply to us regardless of whether other people accept them as well.

But in the realm of law, the legal positivist claims, this admonition is out of place. What is *legally* right or wrong *does* depend on other people and certain other people in particular. According to H.L.A. Hart, if judges accept a rule requiring you to jump off the proverbial bridge, then it is legally wrong for you *not* to plunge into the icy waters below.

This claim follows from the positivist's picture of morality and law as distinct domains with correspondingly distinct ground rules. According to this picture, the proper way to establish the existence of moral rules is to engage in substantive moral argument. It is never enough simply to say: "That's what we do round here." While a convention may of course be morally relevant, it is because some moral fact ultimately deems it to be so. In the case of law, on the other hand, rules must satisfy the specific criteria for legal validity, and these criteria can only be discovered through empirical observations of the relevant legal communities. To divine the set of legally valid rules, in other words, one must know what legal officials think, intend, claim and do. For the legal positivist, it is simply irrelevant to point out that these criteria of validity are morally illegitimate, or that they sanction undesirable rules. Regardless of the merits, the law is just what certain people think, intend, claim and do around here.

In the past three chapters we rehearsed the various arguments legal positivists have used to buttress this story and found them lacking. In the next two chapters, I want to present an alternative to these arguments, which I believe captures the power of the positivistic picture of law while also addressing the important limitations we noted earlier. My strategy is to show that there is another realm whose norms can only be discovered through social, not moral, observation, namely, the realm of *planning*. The proper way to establish the existence of plans, as I argue below, is simply to point to the fact of their adoption and acceptance. Whether I have a plan to go to the store today, or we have a plan to cook dinner together tonight, depends not on the desirability of these plans, but simply on whether we have in fact adopted (and not yet rejected) them. In other words, positivism is trivially and uncontroversially true in the case of

plans: the existence of a plan is one thing, its merits or demerits quite another.

As I hope will become clear in what follows, my purpose here is not to draw an analogy between laws and plans, but to flesh out an implication. The existence conditions for law are the same as those for plans because the fundamental rules of legal systems *are* plans. Their function is to structure legal activity so that participants can work together and thereby achieve goods and realize values that would otherwise be unattainable. For that reason, the existence of legal authority can only be determined *sociologically*: the question of whether a body has legal power is never one of its moral legitimacy; it is a question of whether the relevant officials of that system accept a plan that authorizes and requires deference to that body.

I am going to argue here that understanding fundamental laws as plans not only vindicates the positivist conception of law, but provides a compelling solution to our earlier question about how legal authority is possible. For the picture that emerges is one in which the creation and persistence of the fundamental rules of law is grounded in the authority that all individuals possess to adopt plans. As I attempt to show, this power is not conferred on us by morality. On the contrary, it follows from the fact that we are planning creatures. As the philosopher Michael Bratman has shown in his groundbreaking work on intention and action, human beings have a special kind of psychology: we not only have desires to achieve complex ends, but we also have the capacity to settle on such goals and to organize our behavior over time and between persons to attain them.

Building on Bratman's insights, I want to show that understanding the law entails understanding our special psychology and the norms of rationality that regulate its proper functioning. For that reason, I am going to spend a significant amount of time describing the activity of planning, the structure of plans, the motivation for creating plans and the rationality constraints that attend this activity. I will begin by constructing simple hypotheticals involving one person planning his own actions and then move on to more complicated examples, such as group planning in hierarchical and non-hierarchical contexts among both small and large numbers of people.

One of my main goals in this chapter is to show that planning is a surprisingly diverse activity. Not only can it be carried out in very different ways, but it comprises many distinct stages. In fact, multiple individuals can engage in the planning process at the same time: one person can formulate a plan, another can adopt it, a third can apply the plan and the fourth can enforce it. Plans are also complex entities: they have a rich structure and assume diverse forms. As our hypotheticals will illustrate, planners are able to combine different kinds of plans to construct new and sophisticated technologies of planning, which enable participants

in shared activities to navigate complex, contentious and arbitrary environments.

In the next chapter, I want to develop my central argument that legal activity is best understood as social planning and that legal rules themselves constitute *plans*, or *plan-like* norms. I realize that this claim is not self-evidently true and the relationship between legality and planning is not yet apparent. But as the nature of planning becomes more explicit, and our examples become more complex, the connection between the two phenomena will become clearer. Or at least that is the plan.

2. INDIVIDUAL PLANNING

The Partiality of Plans

I am sitting at my desk in my office and thinking about what to do for dinner tonight. Should I eat out or cook dinner at home? Since I feel a bit guilty about having frequented restaurants so often lately, I decide on the latter option. I now have a plan, namely, to cook dinner at home tonight. Admittedly, it is not much of a plan, because I have no food at home to cook. So the question with which I started – where to eat? – has been replaced with a new query – where should I get food to cook?

I respond to this new query by forming an intention to buy the food from a supermarket. And so I now have two plans: one to cook dinner tonight and one to buy food at the supermarket. These plans are clearly related to one another. Buying food at the supermarket is a means to cooking dinner at home tonight. When one plan specifies a means for accomplishing, or a way of realizing, the end fixed by another plan, we will say that it is a “sub-plan” of the second. Thus, the plan to buy food at the supermarket is a sub-plan of the initial plan to cook dinner tonight.

Of course, by adopting these two plans, I have also created a third plan, namely, the plan to cook dinner *by* buying food at the supermarket. This larger plan, we might say, has two parts to it: the first is the plan to cook dinner tonight and the second is the plan to buy food for dinner at the supermarket. These parts are related as means to end: the second part is a sub-plan of the first.

As Michael Bratman has shown, planning typically involves the creation of these larger plans.¹ When I initially form my intention to cook dinner tonight, my plan simply identifies my end goal. But if my plan is going to work – that is, if it is to organize my behavior so that I may attain the goal I set for myself – I have to specify the means as well. I must decide which meal to make, what kind of food to buy, where and when to buy it, whether to make enough for leftovers, which knife to use when preparing the food and so on.

Bratman notes that these larger plans are typically partial. They begin as empty shells and, as more details are added, they become more comprehensive and useful. Plans are almost never exhaustive because

there is rarely a need for a full specification of every step necessary to achieve a goal. My plan to cook dinner tonight will not specify the correct way to hold the knife when I cut the food because I have the skills necessary to accomplish the task without deliberation or reflection.

As plans are filled in, they thus naturally assume a nested structure. My plan to make dinner tonight specifies the overall goal I wish to achieve. My plan to buy food from a supermarket, as we mentioned, is a sub-plan of the overall plan of making dinner. My intention to buy chicken at the supermarket after work is, in turn, a sub-plan of the plan to buy food at the supermarket and, thus, a sub-sub-plan of the overall plan to cook dinner tonight.

The nested structure of plans explains how past deliberation shapes present planning. When constructing my plan, I take my prior decisions about means and ends as given. These plans and sub-plans are settled and not up for reconsideration. Rather, my present deliberation is confined solely to those options that are not ruled out by past decisions. If I have decided to go to Stop & Shop to buy food, I figure out how to get there, not whether it might be better to go to Pathmark instead.

As Bratman points out, plans not only organize our behavior, they also organize our thinking about how to organize our behavior.² The planner sets ends to be achieved and determines which means are best suited to achieve those ends. Once selected, these means are treated as new ends and lead the planner to determine which new means ought to be adopted. By fleshing out plans in this manner, the planner ensures that, according to his beliefs, he will perform all the necessary actions in the right sequence and thus realize the overall ends of the plan.

Planning Ahead

Clearly, if we did not seek to achieve complex ends, there would be no need to engage in planning about the future. Non-human animals have little or no use for plans, whether because they do not have goals that require significant degrees of behavioral organization or because they do but are hard-wired with the proper instincts to fulfill them. It is plausible to suppose that dogs, cats and mice act purposely insofar as they have desires and that they act on those desires in light of their beliefs. But they probably do not *plan* since they lack both the need and capacity to do so.

We are planning creatures in part because we have desires for many ends that demand substantial coordination. But there is another aspect of our psychology that compels us to planning, namely, that our rationality is limited. If we were like chess computers, able to look ahead millions of moves on each turn and choose the best play among the myriad alternatives, we might have little use for planning. Since we are not, however, mentally omnipotent and rational deliberation is costly, we must conserve our energies. I cannot spend every second of the day thinking about what to do and reviewing every one of my past judgments or I

would never get anything done. It is generally far more efficient to decide on a course of action before hand and follow it when the time for action arrives.

Planning ahead is not, however, a solely economizing measure; we often plan out of paternalistic concerns for our future selves. Deliberation is a risky endeavor. If I were to engage in perpetual reasoning about what I ought to do, I would almost certainly find myself in a poor state to make decisions. I would lack the composure, energy and will either to think through all of the possibilities or to resist temptation. Making up my mind well in advance allows me to pick a good occasion for reflection, provides ample time to puzzle things out and saves me from choosing a course of action deemed less than ideal from my current perspective.

We have good reason, therefore, to be planners: planning guides and organizes our behavior over time, enabling us to achieve ends that we might not be able to achieve otherwise. As Bratman has argued, this pragmatic rationale for planning suggests that the activity is subject to several different norms of rationality.³ Suppose that having decided to make dinner at home tonight, I do not give that decision any more thought. I do not contemplate how I am going to pull off this feat, e.g., where to get the food, what to eat, when to cook and so on. These omissions would be irrational because I would not be able to achieve the end that I set for myself. I cannot just *cook dinner*. Cooking dinner is not a simple action like raising my arm – it is a multi-step process, requiring that I make preparations, string numerous actions together and perform them in the proper order.

When we set ends for ourselves, rationality thus demands that we flesh out our plans. Of course we need not settle all outstanding issues at once. While I should soon decide when to buy the food for dinner, I can probably wait until I arrive at the supermarket to decide what to make and how much to cook. And I certainly can wait until I get to the kitchen before settling on which knife to use. Indeed, there is a pragmatic argument for leaving certain aspects of plans open until the time for action nears. Settling on a course of conduct far ahead of time in the absence of complete information is a risky thing to do. By leaving our plans for future actions somewhat sketchy, we provide the measure of flexibility necessary to enable us to fill in the details as our visibility substantially improves.

Strictly speaking, rationality does not demand that the planners formulate courses of action *themselves*. Others may tailor the means and communicate the plan to the person committed to the end in question. My foodie friend may tell me what food to buy and how to cook it. Rationality does not forbid taking instruction from others; indeed, it requires it when they are more reliable or when doing so is economical. When we say that planners are rationally obligated to “fill in” their plans, we mean they are required to *adopt* the means to their ends, not that they are required to figure out what those means are themselves.

Rationality not only demands that we fill in our plans over time; it also counsels us to settle on plans of actions which are internally consistent and consistent with each other. In this respect, plans are different from desires. Desires typically conflict, but plans must not. There is nothing irrational about wanting to lose weight and wanting to have dessert, but it is incoherent to go on a no-dessert diet and, at the same time, order dessert. In the same way, one's plans must be consistent with one's beliefs about the world. One should not adopt a plan that one believes cannot successfully be carried out. Again, these consistency demands are supported by the pragmatic rationale for planning: consistency within plans is necessary if we are to achieve the ends of the plan; consistency between plans is necessary if we are to achieve the ends of all our plans; and consistency with one's beliefs ensures that the plans we have adopted can be achieved in the world in which we find ourselves.

Finally, if planning is to compensate for our limited cognitive capacities and reduce deliberation costs, our plans must be fairly stable, which is to say that they must be reasonably resistant to reconsideration. Suppose on my way home from the office I ask myself: "Should I eat out or at home?" After thinking about the issue, and weighing the ease of dining out against the economy of eating in, I settle on the same option I chose earlier, namely, cooking dinner at home. My reconsideration of the issue of where to eat, therefore, rendered my prior decision moot. I did not derive any benefit from my earlier planning, for I ended up engaging in the same thought processes that I followed earlier.

To be sure, choosing a plan does not set it in stone. Reconsideration is rational when, but only when, there is good enough reason to do so. If I find out, for example, that the power is off at home, then I *should* of course reconsider my earlier decision. If nothing much has changed, however, it would be irrational to upend my earlier judgment. It would defeat the purpose of having plans if I were to review their wisdom without an otherwise compelling reason to do so.

Top-Down vs. Bottom-Up Planning

As we saw in the last section, planning never occurs in a vacuum. Past decisions form a framework that constrains and guides present deliberation. When a rational planner contemplates whether to pursue a certain end, she attempts to determine whether the goal can be achieved in a manner compatible with this framework of prior decision-making.

In fleshing out her plans, the planner may pursue one of two options: "top-down" or "bottom-up" planning. In instances of top-down planning, the planner starts with the overall action to achieve (cook dinner) and breaks it up into a few major tasks (buy food, cook food, clean up). She then refines each major task into its component parts (buy food → drive to store, pick up food, buy food, load car and drive home). The planner continues this process of refinement at each step until she reaches a point

at which the relevant actions can be accomplished without further planning (get in car, start car, make right at State Street, etc. ...)

In cases of bottom-up planning, the planner starts with a vague sense of the goals to be achieved (I want to eat some soup for dinner) and proceeds to think through the lower-level tasks in great detail (make the stock → fill pot with water, throw in carrots, celery, onions and chicken, skim when boils, simmer for one hour). Any decision to carry out a simple task in a certain way constrains how other simple tasks will be carried out (making chicken stock requires going to a market that sells chickens). Once other basic tasks are planned, she attempts to combine them to see whether they fit together. They might fail to connect up for two reasons: either they are inconsistent with each other (it may not be possible to get to Safeway and cook the soup in the available time) or they are consistent, but insufficient, to accomplish any higher-level task (something more must be added to the stock to make soup). In the first case, consistency must be restored through fiddling with one or both of the conflicting tasks (go to Pathmark instead). In the second case, new lower-level tasks must be added to achieve the necessary effect (add rice to stock). Once the sub-plans are adjusted, the new higher-level tasks are then combined to see whether they fit together (is there enough time to make the rice? is soup enough for dinner?). The process of planning ends when all the tasks settled on are sufficient to achieve the ultimate goal.

Bottom-up planning is especially useful when the planner is unsure which tasks she must undertake or how they all will hang together. In such cases, she cannot start from her main aim and methodically work her way down the planning tree because she lacks an abstract appreciation of how the various tasks connect up. By starting with lower-level tasks whose contours she understands, her detailed planning of one part of the project constrains how the closely-related tasks must be performed. She can proceed to fill in adjacent slots, moving slowly across and up the planning hierarchy and eventually establishing a coherent and complete plan of action.

The downside of bottom-up planning is that the ordering of tasks is not informed by a full sense of the overall structure of the activity. Too much attention to low-level detail may unwittingly cause the planner to lose the forest for the trees and result in plans that are riddled with inconsistencies, gaps and redundancies. By contrast, if the functional shape of the project is well-understood, a top-down approach is usually more appropriate. To be sure, planning in real-life usually combines both elements of top-down and bottom-up planning, with the best mix determined by how well the planner understands the nature of the activity she intends to perform.

Applying Plans

There would be little point in making plans if we did not use them to guide our conduct. If my cooking plan is to be useful to me, it is not enough to formulate and adopt it: it must be applied it as well.

As I employ the term, to “apply” a plan means to use it to guide or evaluate conduct. A plan is applied prospectively when it is used to determine which actions are required, permitted or authorized in the circumstances; a plan is applied retrospectively when it is employed to assess whether an action conformed, or failed to conform, to the plan in question. (A note of caution: sometimes, when we say that we are “applying” a plan, we mean that we are carrying it out. Thus, I apply my chicken-on-sale plan when I actually buy the chicken because it is on sale. When I speak of “applying” plans, however, I will be referring to the use of plans to guide or evaluate action, reserving “carrying out” for the process of following through on them.)

Applying a plan is a three-step process. The plan-applier must determine: 1) the content of the plan; 2) the context of its application; and 3) how to conform to the plan in that context. Thus, if I apply my plan during the afternoon, I must decide what that plan is (i.e., “to cook dinner tonight”), what the world is like (e.g., do I have enough food to cook?) and what I should do to execute the plan at that point (i.e., must head out soon to buy food). The planner might find out that the plan is not applicable to a particular situation, in which case there is nothing that the plan requires, permits or authorizes the subject to do or not to do.

Just as someone need not formulate a plan she adopts, she need not apply it herself either. If I ask the butcher for a pot roast at the supermarket, my friend might say to me: “Wait, I thought you told me that you were going to buy the chicken if it’s on sale, and look, it is on sale.” My friend, thus, applied my plan for me.

Regardless of who applies the plan, rationality requires that the plan-adopter make sure that *someone* does. To adopt a plan and not use it, or use it incorrectly, is irrational. In other words, a planner is subject to criticism when she forgets that she adopted a certain plan, cannot figure out the content of the plan, does not bother to find out what the world is like, fails to use her beliefs to determine the application of the plan, uses these beliefs incorrectly or simply does not carry out the plan that she believes applies.

In order to determine the content of the plan, the planner must be careful not to engage in deliberation about its merits. As we have seen, the value of a plan is that it does the thinking for us. If in order to determine the content of my cooking plan I had to deliberate about whether I should cook dinner tonight, then adopting my plan would have been useless. Plans cannot do the thinking for us if, in order to discover their counsel, we are required to repeat the same sort of reasoning.⁴

Let me end this section on individual planning by saying a few words about what I mean by the term “plan.” By a “plan,” I am not referring to the mental state of “having a plan.” Intentions are not plans, but rather take plans as their objects. For my purposes, plans are abstract propositional entities that require, permit or authorize agents to act, or not act, in certain ways under certain conditions.

In Chapter Two, I characterized a norm as an abstract object that functions as a guide for conduct and a standard for evaluation. In keeping with this characterization, plans too are norms. They are guides for conduct, insofar as their function is to pick out courses of action that are required, permitted or authorized under certain circumstances. They are also standards for evaluation, insofar as they are supposed to be used as measures of correct conduct, if not by others then at least by the subjects of the plans themselves.

When a person adopts a personal plan, she thus places herself under the governance of a norm. This power of self-governance is conferred on her by the principles of instrumental rationality. Planning creatures, in other words, have the rational authority to subject themselves to norms. Indeed, this authority explains the efficacy of planning. Planning psychology is unique not only because it enables planners to form mental states that control future conduct but insofar as it enables them to recognize that the formation of these states generates rational pressure to act accordingly. Thus, when an individual adopts a self-governing plan, the disposition to follow through is not akin to a brute reflex; it is instead mediated by the recognition that the plan is a justified standard of conduct and imposes a rational requirement to carry it out.

While all plans are norms, not all norms are plans. The laws of logic and the principles of morality, for example, are norms but they are not usually considered plans. Plans are “positive” entities – they are created via adoption and sustained through acceptance. By contrast, logical and moral norms exist simply by virtue of their ultimate validity. They are not created by anyone. Plans are also typically partial norms which are supposed to be fleshed out over time, whereas it makes no sense to talk about incrementally developing the laws of logic or morality.

Plans are also purposive entities. They are norms that are not only created, but are created *to be norms*. I adopted a plan to cook dinner tonight precisely so that it would guide my conduct in the direction of cooking dinner. Customary norms, on the other hand, may exist even though they were not created in order to be used in decision-making. The custom to eat turkey on Thanksgiving, for example, may have arisen spontaneously and not for the purpose of getting people to choose to eat turkey on Thanksgiving.

In general, we can say that a norm is a plan as long as it was created by a process that is supposed to create norms. In the case of individual planning, the process is the psychological activity of intending. In institutional contexts, however, as we will see in the next chapter, a plan

may be created even though the one who adopted it did not intend to create a norm. As long as the institutionally prescribed procedure is followed, he will be acting in accordance with a process that is supposed to create norms and will therefore be capable of adopting a plan.

While all plans are positive purposive norms, not all positive purposive norms are plans. Threats, requests and advice are created by human action, and are created to guide action, but they are not typically structured norms: unlike plans, they are not characteristically partial, composite or nested. More importantly, these norms do not aim to guide conduct by settling questions about how to act. Threats, requests and advice merely purport to be one factor among many to be considered. It shows no irrationality or disrespect to deliberate about whether to capitulate to a threat, grant a request or accept advice. By contrast, when one has adopted a plan, for oneself or for another person, the plan is supposed to preempt further deliberations about its merits.

Finally, a norm is a plan only if it is created by a process which disposes the subjects of a norm to follow it. If I plan to cook dinner tonight, I will be disposed to cook dinner tonight. This does not entail that I *will* cook dinner, only that under normal conditions I will. It follows then that decisions that do not instill dispositions in their subjects to comply do not generate plans. If a madman “plans” to withdraw the United States Army from Iraq, no withdrawal plan exists because the madman’s decisions have absolutely no effect on troop movements.

To conclude, a plan is a special kind of norm. First, it has a typical structure, namely, it is partial, composite and nested. Second, it is created by a certain kind of process, namely, one that is incremental, purposive and disposes subjects to comply with the norms created.

3. PLANNING FOR SMALL-SCALE SHARED ACTIVITIES

My Part and Your Part

Having decided that I will cook dinner at home, it occurs to me that it would be fun to cook with someone else. I therefore call up my friend, Henry, invite him over to cook together, and he agrees. We now have a plan: that is, to cook dinner together tonight.

Of course, this plan won’t not be of much use to us unless we fill it in. But here matters become complicated. Whereas I was previously able to resolve all issues regarding cooking by myself, I must now consult Henry, at least with respect to the major tasks. It would be unfair, let alone rude, to decide unilaterally what we are going to eat, when we should start cooking and so on. In addition, we have a new set of questions that must be answered, such as who should get the food, who should cook which part of the meal, who should clean up and so on. Planning for two involves organizing behavior not only across time but between persons as well.

Let's say that Henry and I decide to cook fish and make a salad. I opt to get the fish and he opts to get the ingredients for the salad. How many plans do we now have (or, as I will sometimes say, how many plans do we now *share*)? Again, the answer depends. In one sense, we have adopted five plans: we cook dinner together tonight; we cook fish together tonight; we make a salad together tonight; I get the fish before dinner; you get the salad ingredients before dinner.

In another sense, we share only one plan, namely, the plan to cook dinner together tonight. Cooking fish and making salad together are sub-plans of the overall plan of cooking dinner together. We cook dinner together *by* cooking fish and making salad together. Likewise, my purchasing the fish before dinner is a sub-plan of our cooking fish together tonight, Henry's procuring the salad ingredients before dinner is a sub-plan of our making salad together and each is a sub-sub-plan of the overall plan of cooking together tonight.

As the foregoing suggests, the structure of shared plans is similar to that of individual plans. Shared plans too are typically partial: they are developed over time, beginning with a settling of ends and a progressive divvying up of steps each member is to take. Shared plans are also normally composite: they have parts which are themselves plans. Our plan to cook dinner, for example, includes plans to buy and cook the food. Finally, shared plans are usually nested: they identify the overall end to be achieved by the group and specify in their sub-plans the parts that everyone is to take. When fleshing out how we are to cook dinner together, we take our cooking together as settled and deliberate only about which courses of action each of us should take so that our combined activity adds up to a tonight's dinner.

Planning for the Group

Although planning for a group can be a complicated affair, especially when it is also performed *by* a group, the benefits of planning normally outweigh the costs. As with individual planning, participants in a group activity will not always be able to ponder the optimality of their next move. Since Henry and I are not hi-tech deliberation machines, programmed for precisely this purpose, we need to map out our some of actions beforehand so that when the time for execution arrives we can each consult our respective parts of the shared plan and proceed accordingly.

There are, nonetheless, reasons to plan for the group which are quite independent of these benefits of planning ahead. To see what they are, let us begin by considering the advantages of acting together. According to David Hume: "When every individual person labours a-part, and only for himself, his force is too small to execute any considerable work; his labour being employ'd in supplying all his different necessities, he never attains a perfection in any particular art; and as his force and success are not at all

times equal, the least failure in either of these particulars must be attended with inevitable ruin and misery.”⁵ As Hume points out, individual effort is often too feeble, amateurish and/or risky to accomplish many of the ends we wish to accomplish (think of building a house all by yourself). By pooling efforts in an orderly fashion, we are able to supplement our energies, engage in specialization and minimize the risk of failure.

In a shared activity, then, the actions of the participants must be coordinated with one another in order to benefit from the pooling of talent. The utility of any course of action cannot be evaluated in isolation but only as part of a total vector of concerted effort. Rational deliberation in a shared activity is, therefore, inherently strategic: what one person ought to do depends on what others will do.

We can imagine two basic ways in which participants to a shared activity might attempt to order their affairs. The first way is completely improvised: at each moment, each person assesses the various options open to them based on their predictions about how the others will act and chooses the option that they judge to be best. When an activity is completely improvised, no guidance is provided to any participant; each is left to their own deliberative devices.

While this kind of improvisation is effective in many contexts, such as leisurely walks, doubles tennis and jazz riffs, there are a number of reasons why it cannot be a universal method for coordinating shared activity. First, participants might not always be able to trust one another to make the right decisions: some participants might be less informed and mistakenly judge certain choices to be the best; some might have all the necessary information, but become overwhelmed at the moment of choice and pick the wrong option; or some might have different preferences and, as a result, choose courses of conduct that work at cross purposes. Without some method for correcting or guiding behavior, information asymmetries, cognitive incapacities and divergent preferences threaten to plunge joint ventures into chaos.

Second, improvisation of this kind might also fail to coordinate behavior due to problems of predictability. Since rational deliberation in shared activities is strategic, improvisers must be able to predict what their fellow improvisers will choose. Predictions, however, may be hard to come by. Participants cannot assume that others will do what they want them to do because group members might have different wants. Although participants might be able to predict behavior if they knew what everyone believes and desires, they will not typically have that sort of information, and even if they did, it might be very time consuming to figure out what they will do by calculating what it would sensible for them to do given every thing they believe and want.

The problems of predictability are especially acute when the group faces a coordination problem. Recall that in these strategic situations, the solutions to the games are arbitrary. When solutions are arbitrary, each player's preferences are determined exclusively by their expectations of

what the other players will do. For example, I may not care whether I get the fish and Henry the salad ingredients or he gets the fish and I the salad ingredients. He may be similarly indifferent. The right strategy for each of us, therefore, depends entirely on which strategy the other chooses. Unless we have some basis for predicting each other's choices, our attempt to coordinate our actions is likely to be thwarted.

Unconstrained improvisation is not a robust method for coordinating shared activities because it is appropriate only when there is a very high degree of trust and predictability among participants. When either breaks down, some form of advanced planning will be the preferred strategy. Thus, if I have worries about Henry's abilities or preferences, I should raise them with him prior to action. For if I can convince him that it would be best for him to choose one option and me another, or bargain to some form of compromise, we can settle on the same joint strategy and implement a good plan when the time for action rolls around.

Planning in the context of shared activities, thus, serves a crucial control function. It enables some participants to channel the behavior of others in directions that they judge to be desirable. The need to guide the behavior of the other members will be pressing whenever members have jaundiced attitudes toward each other's intelligence, character or knowledge or when their preferences significantly diverge. In such circumstances, participants cannot simply assume that others will be able to coordinate their behavior properly. They must use plans to direct their fellow participants to act in the way that they want them to or believe they should.

Planning in group contexts also alleviates problems of predictability. The adoption of shared plans by members of the group obviates the need for detailed knowledge about everyone's beliefs and desires. I don't have to *know* that Henry wants to choose some option in order to be able to expect that he will choose it. Notice here that the function of planning is not to improve choices, but rather to render them legible to others. By having a common blueprint that each of us accepts, each of us can reasonably forecast that the others will do their part. In these circumstances, it is better to settle for a decent plan than to hope for the best solution.

Complex, Contentious, Arbitrary

As we have seen, group planning is unnecessary for shared activities when it seems clear that the members of the group, if left to their own devices, will end up coordinating their behavior effectively. However, if participants harbor reasonable worries that order will not appear extemporaneously, or that it will be significantly defective, then they ought to formulate and adopt shared plans. Such fears will naturally arise in three kinds of scenarios: when the activities to be shared are complex, contentious, or possess arbitrary solutions.

In the absence of guidance, *complex* activities demand significant knowledge and skill, tax cognitive capacities, and consume precious intellectual resources. Completely improvised attempts at coordination are thus bound to lead participants to distrust their own judgments or those of their fellow group members. Plans aim to resolve these doubts and disagreements by greatly simplifying the decision-making procedure. Instead of having to arrive at an all-things-considered judgment about what to do, participants can focus on the same few variables and, as a result, make better choices, or at least ones that cohere well with those of others.

In the case of *contentious* activities, there is a threat that, without planning, some participants will choose poorly, or worse, act at cross purposes. The contentiousness of an activity might stem from its complexity, or from the simple fact that the members of the group have different preferences or values. In either case, it is crucial that potential conflicts be identified and resolved ahead of time. The function of planning here is to settle disputes correctly and definitively before mistakes are made and become irreversible.

Finally, the *arbitrariness* of many aspects of shared activities generates coordination problems which render the behavior of the other participants difficult, if not impossible, to predict. Plans pick one solution out of a multiplicity of options, enabling the group to converge on that solution and hence to coordinate its actions successfully.

To be sure, a shared activity may be *so* complex that planners may be unable to map out a sequence of events that will lead to the desired outcome. A standard critique of planned economies, for example, is that allocation decisions are so intricate that no central body can gather all the necessary information, process it correctly and optimally direct production and consumption.⁶ This does not however mean that planning plays no role in market economies. As I will show in the next chapter, the rules of property, contract and tort can be understood as general plans whose function is to create the conditions favorable for optimal order to emerge spontaneously. Rather than acting as visible hands directly guiding economic decisions, they provide market actors the facilities to carry out their own profit-maximizing plans so that overall economic efficiency will be maximized in the process.

Similarly, if a shared activity is too contentious, participants will be unable to agree on a common plan to order their affairs. Imagine, for example, trying to use the political process to distribute food, shelter, education, child care, sneakers, books, shampoo, laptops, iPods, dvds, beer, candy bars, paper clips and so on. Aside from being impossibly complicated, questions about optimal levels of production and fair distribution are simply too contentious to be resolved in a collective manner. The plans that structure market interactions, on the other hand, allow individuals who fundamentally disagree with one another to place values on goods and services and to engage in mutually advantageous

trades. The benefits that are unavailable through collective action can thus be had through the spontaneous interaction of group members following their own conceptions of the good life.⁷

4. SHARED PLANS AND SHARED AGENCY

In the last section, we saw why a group would want to converge on a common plan when engaged in a shared activity. We said that the value of planning stems not only from its ability to lower deliberation costs and compensate for cognitive incapacities, but also from its power to coordinate the participants' behavior. Insofar as the utility of an individual action is a function of the choices made by other participants, it is imperative that the behavior of the group members be channeled in the right direction and made predictable to one another. In complex, contentious and arbitrary environments, however, doubts and disagreements about the best way to proceed thwart the prospect of coordination through complete improvisation. Shared plans resolve these doubts and disagreements, harnessing and focusing the individual efforts of the participants so that they may accomplish together what they could not achieve separately.

Having argued for the importance of plans in joint ventures, we might ask how exactly groups have or share plans. For example, what makes our plan to cook dinner tonight *our* plan?

Clearly, when we speak of a group sharing a plan, we don't mean that the group has a collective mind which has adopted a plan. A plan is shared by a group only if each of the members of the group in some sense "accepts" the plan. Henry and I would not share a plan to cook dinner if both of us did not accept the plan to cook dinner.

By the same token, two people cannot be said to share a plan simply because each intends to engage in the same generic activity. I intend to cook dinner tonight and my neighbor intends to cook dinner tonight, but my neighbor and I do not share a plan to cook tonight. To say that a group has a plan to A is to say more (and, as we will soon see, sometimes less) than that each member of the group plans to A.

One reason that Henry and I can be said to share a plan, but my neighbor and I can not, is that Henry and I designed the plan for ourselves, and not for my neighbor. This suggests that a group shares a plan only if the plan was designed, at least in part, with the group in mind, as a joint activity constituted by our individual actions.

The requirement that shared plans be designed for members of the group does not however require that every member play a role in the design of the plan. One group member could take the lead and design the plan for others. In fact, someone who is not even a member of the group could take on this role. My wife could plan for Henry and me to cook dinner tonight for all three of us. Henry and I would then share a plan in part because it would have been designed with Henry and me in mind.

But simply designing a plan for a group is not enough for plan-sharing. For even though my neighbor might have designed a plan with me in mind, my neighbor and I do not yet have a plan unless I agree to it. In order for a group to share a plan, then, each member of the group must accept the plan. And “acceptance” of a shared plan does not mean simply that each member accepts *their* particular part of the plan. To accept a plan entails a commitment to let the other members do their parts as well. Thus, if our plan requires that I cook the fish and Henry make the salad, I am committed to acting in a manner consistent with your making the salad. If you need the big knife to cut the carrots, I must at some point during our cooking let you have it. The acceptance of a plan does not require that the participants actually know the full content of the shared plan; the commitment may simply be to allow others to do their parts *whatever they happen to be*.

Because a plan can be shared only if it is accepted by all participants, shared plans will rarely be complete. I can plan on Henry making the salad without having any commitment to let him use the big knife first. In this case, the shared plan will specify only who makes what, but not who uses the knife first. In fact, there may be no accepted plan apart from the commitment to engage in the joint venture, in which case the shared plan will be virtually blank. “We cook dinner tonight” can be its only content. Wherever the shared plan is unspecified in this way, participants may be required to design individual sub-plans in order to execute the plan itself. Unless these sub-plans come to be accepted by others, these parts will not be shared and may be contested some time in the future. If these problems are anticipated, prudence dictates that efforts be undertaken to resolve them ahead of time by negotiating and accepting new provisions to the shared plan.

Thus far, we have said that a group shares a plan only if the plan was designed, at least in part, with the group in mind and the group accepts it. It seems nevertheless that one more condition is required. Because a plan that is completely secret cannot be shared, it should be insisted that a shared plan be at least “publicly accessible,” namely, that the participants could discover the parts of the plan that pertain to them and to others with whom they are likely to interact if they wished to do so.

At the same time, it should also be noted that plan-sharing does *not* require that members of the group *desire* or *intend* the plan to work. Let’s say I want my house painted and hire my two sworn enemies, Dudley and Stephens, to paint my house. I offer \$1000 dollars to Dudley if Dudley does what I tell him to do. I offer Stephens the same terms. Dudley and Stephens both agree because they need the money. I then tell Dudley to scrape off all the old paint and Stephens to paint a new coat on the scraped surface. Despite the fact that Dudley and Stephens hate me and loathe the idea that my house will be freshly painted, and, as a result, do not intend that the house be painted, they nevertheless share a plan, namely, one that directs Dudley to scrape and Stephens to paint.⁸ They share a plan

because I designed the plan for them, it is accessible to them and they accept it.⁹

Acting Together

In the last section, we tried to explain why Henry and I shared a plan to cook dinner, but my neighbor and I did not. We said that in the former case, the plan to cook dinner was designed with Henry and I in mind, we accepted it and the plan was accessible to us. In the latter case, however, none of these conditions obtained. No plan was designed to enable my neighbor and me to cook dinner together. And since there was no such plan, we could not accept it and it could not be publicly accessible.

Suppose now that Henry and I cook dinner together. One might ask: why is it the case that Henry and I cooked dinner together but my neighbor and I did not? A plausible response is to say that Henry and I acted together because we shared a plan to cook dinner and this plan enabled us to cook dinner, whereas my neighbor and I did not share such a plan. Shared agency, i.e., acting together, is distinguished from individual agency, i.e., acting alone, by virtue of the plans of the agents. Even if my neighbor used my kitchen to cook and cooked it at the same time as Henry and I cooked, and even if we cooked the same food, our cooking was distinct from his cooking because we did not share a plan to cook with him and he did not share a plan to cook with us. Shared plans, we might say, bind groups together.

Shared plans are constitutive of shared agency because they explain how groups are able to engage in the activity.¹⁰ By appealing to them, group members are relieved, at least partially, from deliberating about proper action. Shared plans do the thinking for the group, enabling participants to know what they should do and what others will do. They not only coordinate the behavior of each participant, they organize their further planning, directing them to fill in their sub-plans in manner consistent with their own and other participants' sub-plans.

But while sharing a plan is necessary for shared activity, it is clearly not sufficient. For even if Henry and I shared a plan to cook dinner, we will not have cooked dinner together unless we acted on the plan and successfully carried it out. This suggests that a group intentionally acts together only when each member of the group intentionally plays their part in the plan and the activity takes place because they did so. Henry and I cooked together because we played our respective parts in the shared plan and, in so doing, managed to cook dinner.

In addition to sharing a plan, acting on the plan, and achieving it, it seems that two more conditions are necessary for a group to act together. First, the existence of the shared plan must be common knowledge. We could hardly be said to have acted together intentionally if it were not plain to each of us that we shared the same plan. Second, members of a group intentionally act together only if they resolve their conflicts in a

peaceful and open manner. If Henry and I disagree with one another about who gets to use various pots and knives and, instead of talking our problems out, we wrestle each other over, or hide, every piece of cooking equipment, our activity would be more competitive than cooperative. Force and fraud not only destroy trust, but they render shared intentional activities impossible as well.

5. REDUCING PLANNING COSTS

As we have seen, the function of shared plans is to guide and coordinate the behavior of participants by compensating for cognitive limitations and resolving the doubts and disagreements that naturally arise in strategic contexts. Indeed, shared activities are partially constituted by the acceptance of shared plans precisely because the existence of shared plans explains how agents can work together in complex, contentious or arbitrary environments.

However, many of the same reasons that make shared plans necessary for shared activities also make them costly to produce. If shared plans are needed to regulate behavior in complex and contentious environments, it is likely that they will be expensive to create ahead of time through deliberation, negotiation or bargaining. Fortunately, it is often possible to reduce these costs. As we will see, policies, customs and hierarchy are three ways in which shared plans can be forged without the members of the group having to engage in the time-consuming process of plan formulation and adoption.

Policies

Having enjoyed our collaboration, Henry and I decide to invite several of our friends over to cook with us. Cooking in this larger group turns out to be even more entertaining and, as a result, we start to make dinner together every week. We call ourselves the “Cooking Club.”

Initially, we find planning these culinary events the least fun part of the process. Each week we make many phone calls and send numerous emails to club members trying to work out the details of our get-together: the day, time and location of our dinner, what we will make, who is to get what, who is to cook what and so on.

Slowly, though, we start to learn ways to avoid having to consult each other on every issue. In particular, we begin developing “policies,” i.e., general plans. For example, instead of selecting the menu each week, we decide instead to follow the recipes set out in the Wednesday edition of the New York Times. This general plan radically cuts down on our deliberation and bargaining costs. We simply follow this sub-policy of our shared plan every week and know what each of us should and, hence, will do.

Policies have their downsides, however. While planning every week was tiresome, at least it allowed our choice to suit our then-current tastes. By deciding to follow the recipes in the Times, we tie our hands to courses of action that may be less than ideal. In this respect, adopting policies involves a tradeoff: planners must decide whether the risks of suboptimal outcomes by following a plan outweigh the costs of repeated deliberation and bargaining.

Custom

Another way planning costs are reduced in ongoing shared activities is through the development of customs. When we began the club, for example, we chose a different person's house to cook the dinner every week. Once we went through the entire club roster, we followed the same pattern again. Eventually, we begin to treat this pattern as the norm. In other words, the group regards the choice of venue for our cooking as settled and, thus, not something normally up for reconsideration. Likewise, because everyone always agreed to make three courses for dinner – appetizer, main course and dessert – we eventually take this pattern as the standard for our dinners and act accordingly.

These customs turn out to be quite beneficial. We do not have to deliberate, negotiate or bargain about these matters in order to apply our shared plan. The customs, in other words, coordinate our behavior spontaneously. My sub-plans about venue and menu are consistent with my friends' sub-plans about venue and menu because we always fill out our shared plans in the customary way.

It is tempting to say that our past practice has led us to adopt a "plan" for venues and menus. After all, we regard alternating houses and three-course meals as the right way to cook dinner together. But this temptation should, I believe, be resisted since our customs were not created *for the purpose* of settling questions about proper conduct but instead emerged spontaneously. Each of us independently found it advantageous to act in accordance with the pattern set by past practice and eventually took the matter as settled.

When customary standards arise in a non-purposive manner, I will not refer to them as "plans," but rather as "plan-like" norms. They are plan-like because they do what plans normally do: they economize on deliberation costs, compensate for cognitive incapacities and organize behavior between participants. Like plans, and unlike other norms such as the rules of morality and logic, they are created and sustained by human action. Moreover, they are also typically partial and hierarchical. Our custom to alternate houses does not specify the time that we are supposed to show up at each others' houses. And if we were to fill in this custom by setting a time, this decision would act just as a sub-plan, i.e., specifying the means by which we carry out the end of alternating houses.

Although some customary standards might not be plans, I will nevertheless consider them as eligible to be part of shared plans. Shared plans, then, need not contain only plans, but may incorporate plan-like norms as well. Customary standards are part of a shared plan just in case they are accepted by the members of the group and are seen as specifying the means by which they are to engage in the shared activity.

Introducing Hierarchy

While adopting policies and developing customs did lessen some of the burden associated with planning our dinners, we were nevertheless forced to engage in extensive deliberation and negotiation each week to set up our club meetings. In order to reduce the costs of planning more radically, we decided to let one person take charge of planning the whole meal for the rest of us. The “head chef” for the week would direct the “sous chefs” on what to make and buy, where, when and how to cook the food and so on. We decided to select our leader randomly: the head chef for the next week is the one who draws the longest straw at the end of each week’s dinner.

As expected, most of our planning problems disappeared. When I am head chef, I am able to plan the shared activity without having to worry about winning an argument, striking a deal or forging a consensus. Regardless of whether my friends agree with me, I can get them to do what I think they should do straightaway. That is, I can *order* them to do so. Similarly, when I am the sous chef, I need not enlist the others in filling out our shared plan. I can just sit back and take instruction from the head chef. To be sure, this scheme does not totally relieve me of responsibility for planning. The head chef never completely plans out my actions and hence I am required to fill in the gaps of the shared plan that apply to me.

When the head chef orders a sous chef to perform some action, we might say that he “adopts a plan” for the sous chef. By issuing the order, the head chef places the sous chef under a norm designed to guide his conduct and to be used as a standard for evaluation. Moreover, the head chef does not intend her order to be treated as one more consideration to be taken into account when the sous chef plans what to do. Rather, she means it to settle the matter in her favor. And because the sous chef accepts the hierarchical relationship, he will adopt the content of the order as his plan and revise his other plans so that they are consistent with the order. He will treat the order as though he formulated and adopted it himself and, as a result, will be disposed to apply and comply with it.

In setting up our hierarchy, therefore, we “vertically” divide our labor. Instead of everyone deliberating and negotiating with each other, we entrust one person with the responsibility to fashion the shared plan for us. When accepting the role of sous chef, club members thus surrender their exclusive power to plan. Put somewhat more precisely, when

accepting their subordinate position, the sous chefs *use* their power to plan to outsource various stages of planning to the head chef. Instead of formulating and adopting their own plans, they accept a plan to defer to someone else's planning. In turn, when one of us assumes the role of head chef, we agree to play the role of planner for other members of the group. Instead of simply planning our own affairs, we plan to formulate and adopt plans for others.

In this context, it is possible to see hierarchy as a major technological advance in behavioral organization. By dividing labor between those who plan and those who follow through on such plans, group members are no longer limited to arduous deliberations and unpleasant squabbling on the one hand and precarious attempts at improvisation on the other. When doubts or disagreements arise with respect to the proper way to proceed, superiors can resolve these conflicts quickly and cheaply by issuing orders and thus changing the shared plan that subordinates are required to follow. Leaders are useful, in other words, because they are efficient "planning mechanisms." They can simplify complexity, settle controversy and disambiguate arbitrariness without having to engage in costly deliberations, negotiations or bargaining.

Self-Regulating Shared Plans

But hierarchy is not only an efficient tool for producing shared plans; it is often the product of shared plans as well. In the cooking case, for example, part of our shared plan authorizes one member of the group to adopt plans for the others. Thus, the reason I become the *head chef* after drawing the longest straw is that our shared plan authorizes the longest straw drawer to do the groups' planning.

In a shared activity involving hierarchy, then, shared plans are self-regulating, that is to say, they regulate the manner of their own creation and application. Parts of the shared plan authorize certain members of the group to flesh out or apply the other parts of the shared plan. These "authorizations" are accepted when members of the group agree to surrender their exclusive power to plan and commit to follow the plans formulated and applied by the authorized members. Thus, when someone authorized by the shared plan issues an order, she thereby extends the plan and gives members of the group new sub-plans to follow.

The fact that someone adopts a plan for others to follow does not, of course, mean that, from the moral point of view, those others *ought* to comply. The plan might be foolish or evil and, thus, unless there are substantial costs associated with non-conformity, the subjects morally should not carry it out. However, if the subject has accepted the shared plan which sets out the hierarchy then, from the point of view of instrumental rationality, he is bound to heed the plan. For if someone submits to the planning of another, and yet ignores an order directed to

him, he will be acting in a manner inconsistent with his own plan. His disobedience will be in direct conflict with his intention to defer.

While acceptance of a subordinate position within a hierarchy creates rational requirements of obedience, it may of course be the case that participants were irrational for acquiescing to the shared plan in the first place. Their superiors may be ignorant, unethical or irresponsible. Nevertheless, there are often good reasons to defer. For example, others might know more than the subordinates do about what the group should do and can be trusted to point them in the right direction. As we have also seen, the complexity and contentiousness of shared activities not only increases the benefits of planning, but also its costs. By vertically dividing labor between those who adopt plans and those who apply and carry them out, participants are able to resolve their doubts and disagreements without having to engage in costly deliberations or negotiations. It should also not be overlooked that an individual might accept a subordinate role in a shared activity because they have no other viable option. They might desperately need the money or fear that they will be harmed if they do not. Even in cases of economic or physical coercion, once an individual forms an intention to treat the superior's directives as trumps to his or her own planning, he or she has transformed their normative situation and are rationally – if not morally – committed to follow through unless good reasons suddenly appear that force them to reconsider.

6. PLANNING FOR MASSIVELY SHARED AGENCY

While concentrating the power to plan in the hands of a few is often useful for small-scale shared activities, it is absolutely indispensable when large collections of individuals act together. On the one hand, the complexity, contentiousness and arbitrariness of shared activities grow with the size of the group participating, leading to a corresponding increase in the need for and cost of planning. Without economical methods for adopting and applying plans, it is unlikely that the members of the group will be able to organize themselves through sheer improvisation or group deliberation and bargaining. Even more importantly, hierarchy is necessary because of the need to hold members of the group accountable. If an activity is to be shared in a group of considerable size, those who are committed to the success of the activity must have some way of directing and monitoring those who fail to share their enthusiasm.

As we will see, the simple forms of planning and hierarchy we have been exploring are insufficient for these larger scale tasks. To manage instances of massively shared agency, it is imperative to divide labor horizontally, develop a dense network of plans and erect sophisticated planning structures so that the participants can navigate their way through unfamiliar and challenging terrain and others can chart their progress.

Plans and Alienation

The Cooking Club has been going strong for so long that one day one of us suggests that we open up a catering company. Why not make money doing something that we enjoy doing for free? We all find the idea appealing and thus decide to turn the Cooking Club into The Cooking Club Inc.

Our initial venture into business turns out to be hugely successful. Word of mouth spreads the news quickly and soon we can no longer meet the demand for catering services. We know that we must hire more workers to help us with our business. Expanding the business in this way, however, requires us to change the way that we run it. Because the new workers know little or nothing about the complexities of the catering business, we must provide them with detailed instructions if they are to be productive.

But lack of catering experience is not the only reason forcing us to plan for them. Unlike the founders and owners of The Cooking Club Inc, a large percentage of these workers are not committed to seeing the business prosper. Many are aspiring actors waiting for their big break and care only about picking up their paycheck. Relying on them to organize themselves, therefore, would be foolish. If they get paid as long as they merely appear to help, there is no reason to think that they will in fact be helpful. Given their alienation from the activity, they will not do what really needs to be done if doing it is too demanding.

The natural solution is for those who care about the success of the endeavor to direct the actions of those who do not. Having all read Adam Smith and knowing about the gains to productivity that the division of labor enables, we decide to assign to each worker separate roles, e.g., cook, dishwasher, waiter, driver, bartender, bookkeeper, and so on. The benefits of specialization, we anticipate, will be considerable: instead of teaching the staff how to perform every single task, we can simply train each to do one job well; because each worker only performs one job, they are able to perfect their skills; given that workers stay at their posts, they waste no time shifting and retooling from one task to another; and since each staff member is assigned a specific task, we are able to determine whom is responsible, and should be held accountable, for shortfalls in performance.

In order to divide the labor in the horizontal direction, we adopt policies that direct staff members to act, or not act, in certain ways. For example, the bartender policy states: "During a job, the bartender is to stand behind the bar and prepare the drinks that the guests request." In addition to these role-specifications, we also adopt company-wide directives, such as "Waiters, cooks and bartenders must wear the Cooking Club, Inc., uniform" and "All employees must wash their hands after using the restroom."

Since the policies allocating roles are highly general, we adopt further policies to help guide the staff in applying them. For example, we provide the bartender a book of drink recipes. These recipes stipulate the “right” way to mix various drinks. When a guest requests, say, a Bloody Mary, the bartender is required to carry out the bartender policy by using the Bloody Mary recipe as his guide. If the bartender does not use that recipe, he will not have performed his job correctly.

Call these types of policies *stipulations*. Stipulations do not demand that their subjects believe the stipulated propositions to be true; rather, they are merely required to *treat them as true* for the purpose of applying certain plans. Suppose that the Bloody Mary recipe uses mango nectar. The bartender need not believe that Bloody Marys should be made with mango nectar. Rather, he should regard the stipulated recipe as the right way to make a Bloody Mary only for the purposes of preparing drinks for the guests. Another stipulation is that the customer is always right. Regardless of whether the customer is actually justified on some matter, workers are required to treat what they claim as correct for the sake of doing their job.

In addition to stipulations, we promulgate *factorizations*. Factorizations specify the factors that should be taken into account when planning how to act. For example, we direct the staff to be cost-conscious. This plan directs the staff to give weight to the cost of various actions and adopt plans in part based on this consideration. Like stipulations, factorizations do not require that their subjects actually value the factors specified by the plan; rather, they merely direct them to *treat them as valuable* for the sake of doing their jobs.

In addition to plans that direct planning and action, we also introduce *permissions*. Permissions are best understood as “anti-directives”: they do not direct the staff to do, or not do, any action; rather, they inform their addressees that they are *not* required to perform, or refrain from performing, some action. Thus, the permission to take home leftover food instructs the staff members that they are not required to leave leftover food, which is useful to know if one is concerned that taking food home may be forbidden.

These general plans, stipulations, factorizations and permissions constitute sub-plans of the shared plan to engage in the catering business together. Their function is to guide and organize the behavior of the group. Instead of staff members having to design a shared plan themselves, each can simply appeal to the parts of the shared plan formulated and adopted for them. The promulgated policies also serve a crucial monitoring function. Once they have been adopted and disseminated, the lines of responsibility become clear, rendering it difficult for workers to shirk or blame failure on ignorance. If they do not perform their assigned role or carry it out in the manner specified by the plans, they can be held accountable for any omissions, mishaps or abuses that resulted from their waywardness.

Decentralized Planning Mechanisms

Assigning roles to the staff is necessary, but it is by no means sufficient. While the adopted directives set out the basic division of labor, most of the operational details are left unresolved. Moreover, given the staff's indifference to our success, we need a way of checking that they are indeed doing their jobs. Unfortunately, we cannot organize and oversee day-to-day operations because there are simply too many problems to solve and too much activity to supervise. We are able to allocate roles and set broad institutional objectives but our group is too slow and unwieldy to effectively run the day-to-day aspects of the business.

As a result, the owners develop a more elaborate vertical division of labor. First, we empower individuals who we deem trustworthy to adopt detailed plans for the day-to-day operations of the company. They are authorized to determine who to staff on which job, where trucks should be parked when at catering sites, when the soup gets served during the meals and so on. Second, we direct them to supervise the staff. They are, in other words, to "apply" company policies to the staff to see whether they are carrying them out properly. If they are not, we direct the supervisors to notify the staff members of the shortfall and respectfully insist that the job be done correctly.

By appointing these supervisors, we *decentralize* the process of group planning. Instead of direction coming exclusively from the center, multiple planning mechanisms are available to create and administer the company's shared plan on a distributed basis. When the supervisors create new plans for daily operations, they are engaged in decentralized plan-adoption; when evaluating staff behavior using existing company policies, they are engaged in decentralized plan-application. Because of decentralization, doubts and disagreements about the best plans to adopt, or the proper way to apply existing plans, need not make their way to us. Nor must we supervise every aspect of the business. Local supervisors who are close to the action can resolve conflicts and monitor behavior in an agile fashion.

To decentralize our planning hierarchy, we adopt various *authorizations*. Thus, the authorization which empowers supervisors to apply plans to others can be formulated as follows: "Supervisors have the power to apply those company plans that are directed to staff members." By accepting this authorization, staff members commit themselves, for the purpose of applying company plans, to treat as correct their supervisors' judgments about the applications of company plans.

We also adopt plans that specify how supervisors are to exercise their authorized powers. For lack of a better term, I am going to call plans of this sort *instructions*. One instruction, for example, requires the supervisor to issue a warning before he docks a staff member for failures to comply with company policies. This plan has the following form: "In

order to dock pay from a staff member, a supervisor must first issue a warning.” The instruction does not actually require the supervisor to issue a warning; rather, it specifies the proper procedure that the supervisor is to follow if he wishes to validly exercise the power to dock pay. Thus, if the supervisor fails to issue a warning, the worker cannot be denied wages.

Authorizations, instructions, stipulations and factorizations are special types of plans. Unlike the bartender directive or the smoking permission which regulate *action*, these types of plans guide *planning*. Authorizations specify *who* is to plan, while instructions, stipulations and factorizations specify *how* to plan. We might say that these plans are “plans for planning.” They constitute the self-regulating parts of shared plans which specify the manner in which the shared plan is to be formulated, adopted, applied and enforced.

Affecting Plans

One of the powers conferred on supervisors is to hire and fire employees. But this power is not the same as the power to adopt or apply any plan. If a supervisor fires a waiter for being rude to a patron, he is not directly telling anybody to do anything: he is simply letting the waiter go.

Of course, by firing the waiter, the supervisor affects the applicability of numerous company plans. For example, the bookkeeper is no longer required to pay the fired employee and other employees are not permitted to let him in the kitchen. We might say, therefore, that the authorization to fire employees involves the power to *affect* certain pre-existing company plans.

At the risk of some artificiality, I will consider the exercise of an authorization to affect plans to be a form of planning. For although affecting plans does not involve the creation of any new plans, it is the functional equivalent. When someone affects plans in an authorized manner, he generates the same normative consequences as if he adopted a new set of plans himself. Indeed, affecting pre-existing plans is typically a more efficient way of organizing behavior than adopting new plans. Instead of separately directing the bookkeeper not to pay the waiter, the employees not to let him back in the kitchen and so on, the supervisor can accomplish the same ends by simply firing the waiter.

Modernity and Massively Shared Agency

As we have seen, we respond to the challenge of managing a large group of inexperienced and unmotivated individuals by requiring them to hand over vast amounts of planning power to us. By accepting the shared plan, they not only assume certain roles, but transfer their powers to adopt and apply plans in case where their plans conflict with the planning of the supervisors. This dense horizontal and vertical division of labor channels the behavior of the staff in such a way that they eventually end up doing

what we want them to do. The beauty of the scheme is that the workers themselves need not care a wit about helping us; their interest can lie simply in earning enough money to make it to the next audition. Nor do they have to understand how the whole enterprise hangs together. As long as they do what they are told, our business prospers.

That individuals can be made to work together in pursuit of ends that they do not value is critically important in understanding how the modern world is possible. For the world that we encounter in day-to-day life is distinguished by the enormous scale of social life. Business corporations, consumer cooperatives, trade unions, research universities, philanthropic organizations, professional associations, standing armies, political parties, organized religions, governments and legal systems, not to mention the collaborative ventures made possible by the digitally networked information and communication technology, such as Wikipedia, massively multi-member on-line games (MMOGs), open-source software and the World Wide Web itself, all harness the agency of multitudes in order to fulfill certain objectives. The modern world, we might say, is one defined by “massively shared agency”—the agency of crowds.

Because the modern world is also characterized by diversity, it is extremely unlikely that large-scale ventures can be staffed with individuals who are all committed to the same goals. The Cooking Club Inc., for example, simply could not find enough truly dedicated people to staff our services. Ultimately, we had to rely on others who were willing to do what was demanded of them but no more. In the modern world, alienation and massively shared agency almost always go hand in hand.

As we have seen, shared agency is indeed possible in the face of alienation. In order for a group to act together, they need not intend the success of the joint enterprise. They need only share a plan. That plan, in turn, can be developed by someone who does intend the success of the joint activity. As long as participants accept the plan, intentionally play their parts, resolve their disputes peacefully and openly, and all of this is common knowledge, they are acting together intentionally.

To be sure, some participants may be so apathetic, lazy, selfish, misguided, rebellious, or, in some cases, honorable that they will not be committed to acting on their part of the plan or letting others do likewise unless they are forced to do so. In such cases, the only alternative is to direct others who do accept the shared plan to enforce the group policies through sanctions. At least in these kinds of cases, plan-enforcement is a distinct stage of group planning. For only by enforcing the shared plan can some participants bring others to accept it as well and dispose them to act accordingly.

Although alienation does not confound the possibility of shared agency, the case of The Cooking Club, Inc. illustrates that its existence presents difficult logistical problems for planners. Because alienated participants are not usually committed to the success of the joint activity, it is likely that they will have to be given detailed guidance on how to act.

It may also be necessary to create hierarchical structures so that conflicts are resolved and performance monitored. Finally, those in supervisory positions might need to be authorized to enforce the group's policies through the imposition of sanctions. The task of institutional design, in other words, is to create a practice that is *so thick with plans, plan-adopters, -affecters, -appliers and -enforcers that alienated participants end up acting in the same way as non-alienated ones*. The fact that activities can often be structured so that participants intentionally achieve goals that are not *their* goals accounts for the pervasiveness of massively shared agency in the world around us.

7. LIVING TOGETHER

We began this chapter by exploring individual planning and why we need it. We saw that human beings plan their individual actions because they typically pursue ends that can only be achieved by taking several, sometimes myriad, different actions over time sequenced in just the right order. Our desire to achieve complex ends outstrips our capacity to deliberate continuously and arrive at the optimal choice for every moment. We compensate for this cognitive failing by thinking through the best course of action in advance, settling on it, and then relying on this judgment when the time comes to carry it out.

We then proceeded to explore why and how small groups plan their shared activities. Aside from the deliberative demands that complex activities place on us, we saw that shared activities require constituent action to be coordinated in certain ways. When faced with complex, contentious or arbitrary activities, it is unlikely that completely improvised attempts at ordering will result in synergistic patterns of behavior. Group planning is an improvement over simple improvisation insofar as it enables participants to control behavior and render it predictable to other. By having a common blueprint to guide them, members of groups need no longer guess what part they should play.

While shared plans are often essential to the success of shared activities, we also saw that they are costly to produce. We, therefore, went on to examine several strategies that participants normally use to reduce their planning costs. Adopting policies enables participants to guide their conduct over a whole class of cases; developing customs permits groups to take advantage of plan-like norms in order to settle questions about how to act without anyone formulating or adopting them for the group; and consolidating and concentrating planning power in the hands of a few circumvents the need for the many to deliberate, negotiate or bargain about how to conduct their shared activity.

We also considered the challenges posed by massively shared agency. We saw that the complexity, contentiousness and arbitrariness of shared activities tend to increase with group size to the point that planning and hierarchy becomes not only desirable but absolutely indispensable.

But massively shared agency brings with it a pressure for planning not typically present in the small case. As a group enlarges, the odds that some members will be alienated from the joint activity grow. Developing a dense network of plans and establishing decentralized plan-adopters, affecters and appliers are essential to supply alienated participants with correct instructions for how to proceed as well as standards for holding them accountable.

I would like to end this long discussion of planning by noting one other occasion in which members of groups plan for one another. Most roommates, for example, have policies about how they are to behave in their shared dwelling. These policies usually prohibit playing loud music late at night, require certain cleaning duties and responsibilities, specify who must buy communal items such as toilet paper, butter and beer, identify the proper place to put the key when they leave the house, and so on. While some of these plans regulate shared activity (e.g., stocking the house with essentials), others concern solely individual pursuits (e.g., playing music in one's room late at night).

There is no mystery about why plans are needed to regulate individual actions in communal settings. When people occupy the same space and share a common pool of resources, certain courses of action will result in clashes between individual pursuits, while others will avoid them. Planning is often necessary to ensure that those who live together do not undermine each other's ends.

As with cases involving individual and shared agency, plans that regulate individual pursuits in communal contexts aim to harness the benefits of thinking ahead. First, plans enable the group members to figure out the best ways to avoid conflict and hence eliminate the need to deliberate at every turn about how to steer clear of trouble. Second, they allow group members to anticipate possible mistaken choices that negatively affect others and to prevent them before they happen.

As we saw with shared activities, plans are also useful in communal settings because they are capable of coordinating behavior in complex, contentious and arbitrary environments. Social life presents numerous ethical quandaries about personal and social rights and responsibilities. People not only have doubts about the proper way to live together but, more perilously, often find themselves at odds with one another about how such doubts should be resolved. The contentiousness of living together, let alone its complexity, increases the costs of deliberation, negotiation and bargaining and threatens to generate additional emotional and moral costs should the parties fail to talk through their problems.

Plans are vital for groups because they are capable of resolving many of the ethical problems of communal life. Members of the group who live together and face conflict need not litigate every dispute, disagreement or perceived act of disrespect. Nor need they try to overpower or deceive each other in order to circumvent the difficulty. They may rely instead on plans that were adopted in anticipation of the

conflict. Prior planning allows the community members to treat questions of fairness and what they owe to each other as settled, as matters not up for reconsideration. And in this way they are not only able to economize on costs and increase predictability of behavior, but also to facilitate an ethic of respect among the entire community.

VII

THE MAKING OF A LEGAL SYSTEM

1. THE IDEA OF SOCIAL PLANNING

The 20th Century was not very kind to the activity often referred to as “social planning.” Any list of social engineering projects of the past hundred years tends to read like a veritable Who’s Who of Unmitigated Human Disasters, e.g., the collectivization of Russian agriculture after the Bolshevik revolution, the command economy of the Soviet Union, the Great Leap Forward, the deurbanization of Cambodia under the Khmer Rouge, the villagization of Tanzanian farmers after independence, the totally planned city of Brasilia, and so on.¹¹ At the very least, it seems safe to say that the social planners responsible for these tragedies – Lenin, Stalin, Mao, Pol Pot, Julius Nyerere and Le Corbusier, have fared far worse in history’s estimation than their critics – Karl Popper, Friedrich Hayek, George Orwell, Jane Jacobs, and, depending on your politics, Rosa Luxemburg and Ronald Reagan.

But, in truth, there is no reason that these notorious large-scale public projects should be taken to represent and thus discredit the practice of social planning in general. As we have seen, planning is an excellent, often indispensable, method for guiding, coordinating and monitoring behavior in social settings. What the above list *does* bring into disrepute is a very specific mode of social planning. What distinguishes these disastrous political experiments is the hubristic and coercive use of an untested ideology to radically transform communities purely through directives issued from the center. These social planners conducted themselves as experts whose monopoly on superior scientific, technological and ethical knowledge entitled them to ruthlessly impose their vision of society on everyone else.

Most social planning, however, is not revolutionary, centralized, top-down and directive in nature. In fact, most attempts at group planning in general, and social planning in particular, combine centralized and decentralized mechanisms for progressive and conservative ends; use bottom-up, as well as top-down, practical reasoning; and rely on authorizations in addition to directives. As Hayek himself complained, socialists hijacked the term “social planning” to suggest that socialism is “only rational way of handling our affairs.”¹² But as Hayek reminded us, liberals engage in social planning as well. “The dispute between the modern planners and their opponents ... is not a dispute about whether we ought to employ foresight and systematic thinking in planning our common affairs. It is a dispute about the best way of doing so. The question is whether for this purpose it is better that the holder of coercive

power should confine himself in general to creating conditions under which the knowledge and initiative of individuals is given the best scope so that *they* can plan most successfully; or whether a rational utilitisation of our resources requires *central* direction and organization of all our activities according to some socially constructed ‘blueprint.’”¹³

Hayek’s point here is uncontroversially right. Socialism is hardly unique in advocating that the state engage in social planning. With the exception of certain extreme forms of anarchism, *all* political theories do so to some degree. What distinguishes these various theories is *how* they understand the planning process: to whom they allocate planning authority, the moral ends of planning and which activities should be subject to social planning in the first place.

In this chapter I will begin to develop the central claim of the book, namely, that legal activity is a form of social planning. Legal institutions plan for the communities over which they claim authority, both by telling members what they may or may not do, and by identifying those who are entitled to affect what others may or may not do. Following this claim, legal rules are themselves generalized plans, or plan-like norms, issued by those who are authorized to plan for others. And adjudication involves the application of these plans, or plan-like norms, to those to whom they apply. In this way, the law organizes individual and collective behavior so that members of the community can bring about moral goods that could not have been achieved, or achieved as well, otherwise.

In order to motivate these claims, I begin with the oldest trick in the book. I return to the Cooking Club narrative I began in the previous chapter, drop the club members into the “state of nature” and describe their various reasons for creating a legal system. As I show, communal life generates a need for social planning. Those who live together must be able to organize shared activities, solve coordination problems, settle disputes and ensure that individual pursuits do not thwart one another. As the group attempts to cope with these pressures, they develop simple social planning mechanisms. The success of these mechanisms, nevertheless, inevitably leads to population growth and hence the need for further planning. In order to meet this increased need, simple techniques give way to more complicated and efficient structures of planning. The end result is the creation of a highly sophisticated planning organization, otherwise known as a legal system.

I should emphasize, of course, that I am not making an empirical claim about the evolution of legal systems or why they were originally created. My claim is conceptual: I use the state of nature narrative as an analytical device for developing a more general account about the identity and foundations of law. By examining the reasons that one community has for developing legal institutions, I hope to show that these pressures are hardly unique to this particular fictional setting. Any human community of modest size will experience similar needs for social

planning which can only be met by highly sophisticated technologies of plan adoption and application.

My aim in what follows is also to build on the discussion in the previous chapter by demonstrating that *technologies of planning*, even the highly complex ones that are mobilized by the law, *can be constructed through planning alone*. In other words, to build or operate a legal system one need not possess moral legitimacy to impose obligations and confer rights: one need only have the ability to plan. The existence of law, therefore, reflects the fact that human beings are planning creatures, endowed with the cognitive and volitional capacities and dispositions to organize their behavior over time and across persons in order to achieve highly complex ends.

2. PRIVATE PLANNING

The Cooking Club Inc. eventually becomes so successful that Wall Street approaches us with an offer to take the company public. Unable to resist the lure of obscene wealth, we agree to turn our business over to the public markets. We will still participate in the business at the management level but our cooking days are over.

As it turns out, however, the thrill of being executives in a multinational corporation nevertheless proves to be extremely short-lived. None of us went into the catering business in order to push paper in a corporate office. We soon decide to sell our shares, use the money to buy an island in the South Pacific and start a new community. Alienated by our brief experience of corporate life, we plan to live off the land, treat each other with equal concern and respect, eschew coercive means of social control, and live happily ever after.

The island that we purchase, and rename “Cooks’ Island,” appears to be an ideal location for hunting and gathering. It is inhabited by wild boar, deer and game birds, has a wide variety of fruit trees, numerous species of wild grains and cereals, and a natural lagoon filled with fish, crab and edible seaweeds. We move to Cooks’ Island in the spring, finding plenty of food to eat and abundant materials with which to build huts.

From the outset, small-scale group planning is crucial to our ability to live peacefully and productively together. The main features of our lives – building shelter, collecting fuel, finding food, preparing meals, raising children, playing and entertaining – are shared activities and so we need a way to organize them. The hunters among us must decide where to go, what to hunt, who should flush, shoot, gut and so on. We will also need designated caretakers to look after the children while the hunters are on an expedition. Since everyone has plenty of time on their hands, group members are free to negotiate with one another about who will perform which tasks. In instances of group planning such as this, no one has the authority to tell everyone else what to do: each individual decides which

course of action to take and then finds ways to coordinate his or her chosen course of action with those who wish to join them.

Group planning is not always necessary, however; in some cases, order is spontaneously generated. Coordination problems such as where to dispose of our refuse or which side of the road to ride our bikes on are usually solved through the emergence of coordination conventions. We find others disposing of their refuse in a certain location or riding on one side of the road and we simply follow suit. No one plans for this result – it simply happens through “spontaneous ordering.”

In this respect, we might say that Cooks’ Island is an *unplanned* community. All issues concerning how members should act are resolved solely through the unilateral decisions of individuals or small sub-groups. Questions about which individual and shared activities ought to be required, prohibited and regulated are not resolved on a community-wide basis, either through unanimous consensus or exercises of authority, but rather exclusively through private ordering. In other words, while there is plenty of *group* planning, there is no *social* planning.

Internalizing Costs

Notwithstanding the absence of social planning, members of our unplanned community are able to work and live together in harmony throughout the spring, summer and fall. Winter, however, is a different story. Many of the animals hibernate for the winter and the game birds migrate north. The fish move farther offshore in order to take advantage of richer feeding grounds. The fruit trees bear fruit only in the spring and summer and the wild grains refuse to germinate in the winter. Each of us anticipates this shortage to some extent by privately storing smoked meat and surplus grain, but there is not nearly enough food to feed us all. We are ultimately able to survive only by importing food from the mainland.

After this first difficult winter, we all recognize that hunting and gathering is not a sustainable way of life and that community-wide action is necessary if we are to survive on the island. In keeping with this, the community decides to pool its resources and buy domesticated grains and livestock from the mainland. Together we clear large portions of the island to plant the domesticated strains of grains and cereals and graze our newly purchased sheep, goats and cows. We abandon hunting and gathering and take up farming and ranching instead.

Before we are able to embark on this new agricultural lifestyle however, we have to make an important choice. Until now, the island has been treated as common property. Everyone was entitled to harvest the game animals, lagoon fish, wild berries, fresh water and hardwood timber and use them in any way they saw fit. Now we must decide whether to maintain common ownership of resources, holding the new livestock, crops and pastures as joint owners and engaging in a shared activity of farming, or whether we should instead create a system of private property

in which a share of animals, seeds and land is allocated to each individual in order that they will be able to grow food for themselves.

Since all of the inhabitants of Cooks' Island have taken basic economics, we know what economists would advise us to do in this situation. In a collective property regime, there is always a danger of free-riding. If each islander must work the fields in order to produce food for all, the economist warns, each are likely to be tempted by the following calculation: to do my fair share is to work very hard; if I do not do my fair share, there will be a little less to eat; I would rather eat a little less and not work very hard than eat a little more but work very hard; therefore, I ought not do my fair share. But, of course, if each islander reasons in this way and acts accordingly, everyone will starve.

The economically efficient decision is to switch from a system in which each enjoys the benefits of other's labor to one in which each gain only from their own efforts. In economic terms, the socially optimal decision involves instituting an arrangement whereby each "internalizes" the costs and benefits of their actions. If I benefit from my labor alone, and not from anyone else's, then I have no incentive to be lazy. For if I do not work, I will be the one to starve.

In an effort to make good on this economic insight, we assemble together and engage in another act of social planning. We agree as a community to allocate the newly arable land, seeds and livestock on the basis of family size. The larger the family, the greater the share received. This allocation is accomplished through the adoption of stipulations of the form: "For the purposes of complying with island land-use policies, Family X is to be deemed the owner of Plot A, Seed lot B and Livestock lot C."¹⁴ In addition to these specific plans, we also adopt several general policies that govern the use of the allocated land. For example, one policy permits owners to use and enjoy the property as they see fit, while another forbids a non-landowner from taking the grain or livestock of another without the owner's consent. A third policy permits anyone to cross another's field in order to reach the village square, but a fourth one requires that anyone who does so must compensate the landowner for the destruction of any crops or injury of any livestock on that property.

A principal purpose of these plans is to force each member of the group to internalize the costs and benefits of their actions. The first policy, for example, permits owners to enjoy the fruits of their labor, while the second one attempts to avert free-riding by directing non-owners not to benefit from the agricultural labor of owners without their permission. The third and fourth policies permit non-owners to use another's land when the alternatives are particularly costly, but incentivize them to exercise reasonable precautions by requiring them to pay for what they damage. Insofar as a major aim of these policies is to prevent free-riding, it is imperative that the policies govern the activities of the whole community. That is to say, in order to be effective, they must be *social* plans.

We also adopt plans that allow for the transfer of property rights. A fifth policy, for example, authorizes owners to alienate their property, whereas a sixth instructs the owners of movable property as to how to exercise this power. It states that in order for ownership to pass, there must be physical delivery and acceptance, unless the parties agree otherwise. These policies, in other words, confer power on owners to affect the previous four policies. By alienating property in the proper manner, owners affect who falls within and without these pre-existing policies and hence who may use and enjoy the alienated property.

Planning for Spontaneous Order

Fortunately, our shared plan fulfills its intended purpose. The new property regime leads to a substantial crop yield and livestock supply and provides the community with plenty of food to eat and store for the winter. In fact, each of us has more goods than we can possibly use. And, as a result, markets emerge in which the islanders trade their surplus goods.

Not surprisingly, the participants in these markets are able to adjust their production to aggregate demand and their consumption to aggregate supply without engaging in social planning. Our group does not, for example, decide in advance that the cheese-makers will make more or less cheese. The cheese-makers instead take their cues from the demand they find in the market. If more cheese is demanded, more cheese is supplied. If the amount of cheese demanded cannot be met, the cost of the cheese goes up until the amount of cheese demanded equals the amount of cheese available.

To say that market decisions about supply and demand proceed in the absence of social planning should not be taken to mean that the existence of the market itself is independent of social planning. For the regime of private property that makes our market possible – the allocation of ownership rights to members of the group, the policies for how to treat the items allocated and the power to affect these policies through voluntary exchange – is the product of plans developed *by us for us*. That group members can only procure one set of goods by trading them for another, and cannot simply appropriate them without permission, is a requirement imposed by our shared plan. This plan seeks to boost agricultural output by creating the conditions favorable for spontaneous order to emerge.¹⁵

Nevertheless, as more goods are produced and traded, the possibility of economic loss through mishaps grows. These accidents raise numerous questions of responsibility. Let's say that my cow wanders from my pasture and eats some of your crops. Am I responsible for the damage? Or should you bear the costs of not having built a fence?

Even though our group has no shared plan about liability in these circumstances, the various parties find themselves able to resolve the conflict through private deliberation and bargaining. Each person relies

on their sense of fairness and neighborliness in determining how to reach a just settlement. For example, neighbors normally choose to overlook small damage to their crops caused by grazing livestock. Because they are aware that their own animals probably inflict similar damage to their neighbor's property, they figure that it all evens out in the long run. As for large losses, which tend not to net out over time, livestock owners generally feel obligated to compensate their neighbors for the damage.¹⁶

3. THE SUPPLY AND DEMAND FOR PLANS

The Need for Social Planning

Although our shared plan solves important social problems, it generates new ones as well. For it turns out that the move to a system of private property exponentially increases the range of matters over which we can quarrel. When property was held in common, everyone was permitted to plan their own actions on any aspect of island life. The land, water, air, animals, fruit, grains and so on were freely available to all. The private property regime changed all that, rendering previously abundant resources instantaneously scarce. Our shared plan has rendered most of the land, and the goods it yields, inaccessible to almost everyone. Moreover, while the incentives to create and innovate that were generated by the new property regime increase the overall production of goods, the prevailing system of ownership dictates that these goods are under the exclusive control of only certain members of the group.

As a result of our shared plan, therefore, questions of rights become extremely urgent and, at the same time, increasingly contentious as well. For we now have an incentive to dispute which objects we own and what we are permitted to do with them. Thus, we bicker about whether islanders are authorized to acquire new land, and, if so, how they can or should do so. Owners who live upstream assert the right to use the water for irrigation, even if this means that there is less water for downstream farmers to use. Downstream farmers hotly contest this claim. Those who find freshwater on their property assert exclusive control over this precious resource, while the bulk of the group denies their right to do so. The merchants claim that the islanders should have the ability to condemn private property for public purposes, such as building roads, while another portion of the group, mainly the farmers, vehemently rejects this position.

Our private property regime has not only rendered our communal life more contentious but also greatly increased its complexity. For in an effort to innovate economically, we start to unbundle the property rights allocated to each family under the original shared plan and recombine them with various promises to create new packets of claims and duties. In this way, gifts, leases, easements, bailments, consignments, life-estates, loans, assignments, mortgages, partnerships, trusts, wills, negotiable instruments and other types of contractual and property arrangements, are

formed. But, while successful in certain respects, these new measures once again raise further questions about the content and scope of the normative relationships created. For example, if a farmer enters into an agreement with, say, the baker to supply a certain quantity of wheat at harvest and then, due to bad weather, the crop fails, does the farmer owe the baker damages and, if so, what kind? If one farmer sells a cow to another but, unbeknownst to both, the cow is barren, can the purchaser get his money back? And what if someone builds a hut for someone else but fabricates the roof of that hut out of an inferior material that makes it less durable than the other huts in the village? Can the purchaser insist on a better roof? Can an unsecured creditor foreclose on property if the debtor becomes bankrupt? In what circumstances can a tenant refuse to pay rent and, if those circumstances do not obtain, at what point can a landlord evict a tenant?

Not only have our private transactions become more complex and contentious, but our public projects have done as well. For example, our new economy has generated much material prosperity but it has also skewed its distribution. As a result, we are eager to institute a program of income redistribution. Unfortunately, however, this particular shared activity turns out to be too complicated and presents too many coordination problems for us to be able to bring them about exclusively through improvisation or planning in small groups. Furthermore, while we all agree on the broad outlines of the redistributive program, we disagree about its precise implementation, thus increasing the costs of resolving our conflicts through private deliberations and negotiations.

It is important to note at this point that the doubts and disagreements that arise on Cooks' Island are entirely sincere. Each of us is willing to do what we ought to do – the problem is that none of us knows or can agree about what that is. Customs cannot keep up with the evolving conflict because they develop too slowly to regulate rapidly changing social conditions and are too sketchy to resolve complex disputes and coordinate large-scale social projects. While private negotiation and bargaining are able to quell some conflicts, this process can be very costly, not only in terms of time and energy but emotionally and morally as well. With many more ways to interfere with one another's pursuits and many more goods to fight over, there is a danger that disputes will proliferate and fester, causing the parties to refuse to cooperate in the next communal venture or, worse, to become involved in ongoing and entrenched feuds. Some projects, such as income redistribution, are so complex, contentious and arbitrary that they are simply not feasible through private planning alone.

To compensate for this failure of private ordering, we revive our earlier experiment in social planning and regularize it as an ongoing shared activity. We get together several times a week to discuss how best to handle the social issues that arise within our group. We discuss not only how to structure our interactions in the market but also how to collect and redistribute wealth, educate our children, protect ourselves against

droughts, hurricanes and wild animals, and so on. Though our deliberations and negotiations are sometimes long and protracted, dealing as they do with weighty matters of political morality and group morale, we eventually settle on plans that we believe will enable us to live together on the island for the foreseeable future.

Failure of Consensus

The disadvantages of social planning via consensus nevertheless become apparent very quickly. Not only is it time-consuming and emotionally draining, but it is extremely unstable. For the plans are useful only so long as they are accepted by almost everyone. As soon as people start to reconsider their wisdom, the plans lose their ability to guide behavior and settle conflict and the group must start deliberating and negotiating once again.

Eventually, as the island economy booms and its population multiplies, this consensual method of governance becomes intolerable. Economic prosperity makes it possible to sustain a greater number of people, and family size increases along with this new capacity. At the same time, people from the mainland immigrate to the island in search of the good life. More people engage in more interactions, more interactions lead to more doubts and disagreements, and more doubts and disagreements generate higher planning costs.

The increase in conflict is not only a function of the increased number of interactions. Population growth also entails a more intricate division of labor, with group members engaging in ever more complex activities (the cheese-maker, for example, recently hired ten workers and incorporated his business, with the bread-maker and meat-smoker as the chief equity investors). Simply relying on untutored judgments of fairness and neighborliness tends to be a poor method for resolving the complicated disputes that arise from these arrangements. And the expansion of the population makes it increasingly unlikely that these untutored judgments will be shared among the contestants or that losses from any particular interaction will balance out in the long run. Social life has become extremely complex and more contentious as well. Community-wide shared activities are less and less amenable to large doses of improvisation.

On the one hand, then, population growth enhances the need for plans in order to guide and organize the behavior of the islanders. This *increased* demand for shared plans, however, coincides with a corresponding *decrease* in our community's ability to supply them through consensus. We simply can no longer get everyone to agree to particular solutions to many social problems and, when we can, the time and expense incurred in the course of forging a consensus is enormous.

Nor do our difficulties end here. Experience has shown us that the mere existence of shared plans is not a panacea. For in order for a plan to

resolve doubts or disagreements, the relevant parties must agree about how to implement it. Not only has the number of disputes unregulated by prior plans increased, so have the number of disputes about the application of prior plans. Yet, as the demand for adjudication has increased so, once again, has the cost of supplying it. If the parties must agree on the application of plans in order to settle their disputes, the expense of conflict resolution will rise with the quantity and complexity of the disagreements.

The Solution: Hierarchy

Recognizing that our need for plans and adjudication exceeds our ability to generate them, the inhabitants of Cooks' Island converge on the idea of *hierarchy*. At one of our weekly meetings, we decide to divide the social labor vertically by outsourcing various stages of social planning to a small group of trusted islanders. First, we identify three people who will be the chief plan-adopters for the island. They will act together as a unit to develop social plans for our community. Second, we identify three islanders who serve as plan-appliers for the Island. When a dispute arises between islanders as to the proper application of some social policy and one of the parties wants the dispute resolved, he or she may ask one of these appliers to do so. The plan-applier will determine by herself which course of action is required by the island policies and her decision will be binding on both parties.

In addition to this vertical division of labor, which delegates the social planning to a small group of islanders, our newly hierarchical approach also distributes the planning labor horizontally among those few designated social planners. The new plan authorizes three people to *adopt* plans for the islanders and three other people to *apply* these norms whenever a conflict arises and their services are sought. Our hierarchy, in other words, involves a separation of planning powers. Moreover, the plan centralizes plan-adoption, but decentralizes plan-application: the adopters must act together in order to make plans, but the appliers can resolve conflicts on a solo basis. Only one body can adopt plans, but multiple bodies can apply them.

Notice that the plan which establishes the hierarchy for the island is a shared plan: it has been designed for the social planners, it is accessible to them and they accept it. This shared plan regulates the activity of social planning. It guides and organizes the behavior of the social planners so that each knows which part they are to play in the shared activity. *It is a shared plan for social planning.*

Notice further that since the shared plan was designed for the handful of social planners; it is they who share the plan, not the islanders as a whole. This means that it is not necessary for the community to accept the shared plan in order for it to obtain—though, as a matter of fact, we do accept the plan. Since we consider the social planners to be morally legitimate, we plan to allow the adopters and appliers to adopt and apply

plans for us. For this reason, we consider the shared plan to be the “master plan” for the group.

As expected, the master plan does solve many of the planning difficulties we encountered earlier. Now social policies can be adopted simply by the adopters proclaiming that such-and-such shall be the case – no one has to agree with the wisdom of the policies themselves. As a result, we have an agile protocol for guiding and organizing our community’s behavior and for resolving any disputes that might arise. When one of the designated appliers determines that a policy has been satisfied or violated, his or her judgment does not need to be regarded as wise or right by anyone involved. It is binding simply because it was applied by the authorized individual.

The newly adopted plans prove to be durable as well. In our revamped system, the persistence of plans does not depend on whether members of the community accept their wisdom. As long as they are approved in accordance with the requirements of the shared plan, the plans will be deemed binding, both by the planners and the islanders generally, and followed accordingly.

The Office

Unfortunately, our saga does not end here. For while the islanders find that their original master plan drastically reduces the costs of social planning, it is nonetheless a crude prototype which suffers from several significant flaws.

It turns out that the plan’s *most* significant flaw is that it is limited to a particular set of individuals. The plan specifies those who are currently authorized to plan for the community (i.e., “Bob, Ted and Jane have the power to adopt plans for the residents of Cooks’ Island”) but says nothing about who is to succeed them if they step down, become physically or mentally incapacitated, or die. As a result, when one of them vacates their post, we have to deliberate again about whether we want hierarchy and, if so, who should possess the power to plan for others. What we learn from this frustrating duplication of effort is that it would be far more convenient to devise policies that create an abstract structure of control and specify in impersonal terms who should occupy which role at which time. In this way, when one planner quits, another person can assume his or her structural role but the master plan itself will remain unaffected.

With this in mind, we develop new policies that define various structural roles and identify their occupants in less personal, and more general, terms. With respect to the adopters, for example, three new policies are adopted: an authorization of the form: “Adopters have the power to adopt plans for residents of Cooks’ Island”; a directive of the form: “Appliers are required to apply the plans adopted by adopters in cases that arise before them”; and a stipulation of the form: “A person shall be deemed an ‘adopter’ if and only if they have lived on the island

for more than a year and receive the most votes in the latest island election.”

This hierarchical structure establishes what is otherwise known as an “office”: an indefinite, non *ad hoc* position of power. The office of “adopter,” for example, carries with it various rights and responsibilities, all of which persist over time and attach to whoever happens to occupy the office at the given moment. Because the authorization to adopt plans does not single out planners by name, but rather defines a class of individuals who meet the appropriate qualifications (e.g., they were elected by a majority of islanders), the master plan does not need to be amended each time a new person seeks to acquire the power to plan for the island. As long as that person satisfies the impersonal qualifications associated with the office, they will immediately inherit all of its powers as determined by the relevant parts of the master plan.¹⁷

And we soon discover another advantage of offices. Instead of requiring succeeding adopters to readopt every plan adopted by their predecessors, the master plan is amended to mandate that the policies of past holders of the office are to be followed whenever a new occupant takes over, unless and until the new occupants change the policies in question. In this way, the plans of previous office holders acquire a normative inertia that renders them even more durable.

Institutionalizing Plans

But while the introduction of offices does depersonalize our hierarchy, we soon find that it is still not impersonal enough. For example, when the adopters wish to adopt a plan, they must gather us in the village square and issue proclamations of the form: “We hereby direct all farmers to erect fences on their land no lower than 2 meters” or “In order to discharge debts, debtors must use clamshells.” These proclamations are commands that direct the community to comply with the proclaimed policies.

Needless to say, this face-to-face mode of social planning proves to be tedious for all and so the adopters eventually switch to a less intimate system. This process involves writing down their plans and deliberating on the proposed edict before them. If at the end of the deliberations all of the adopters form an intention to settle the matter in favor of following the order, then the edict is valid. The document is then posted in the village square as a written record for all to see.

The problem with *this* method is that it ties the validity of an edict to the private mental state of the plan-adopters. An edict is binding only if the adopters *intend* for the edict to settle the matter in its favor and for others to guide their conduct by it. In many cases, however, the islanders have legitimate concerns about whether the posted decree truly represents the will of the planners, e.g., when the plan was long and complex, or contained provisions that led to absurd recommendations. These doubts

lead some residents to question whether they are in fact required to follow certain of the posted orders.

In an effort to make the plans more robust, the master plan is amended to include new provisions that specify the formal conditions for the exercise of planning power. For example, in order to enact a new plan, the master plan merely requires that a majority of the adopters say "Aye" when polled. The master plan, in other words, does not mandate that the plan-adopters *intend* that others follow the plan. Instead, plan-adoption has become *institutionalized*: the adopters' votes have normative significance for the islanders regardless of the specific intentions with which they were carried out.

By institutionalizing our social planning in this way, the governance system of the Island attains a very high degree of impersonality. Not only can those who hold the office of adopters adopt plans outside the earshot of the islanders, they can do so regardless of the intentions with which they performed their actions. As long as they follow the procedures set out in the master plan, their actions will have binding force. And so the islanders no longer need to divine the intentions of the planners in order to know which plans they must follow.

4. LAW AS SOCIAL PLANNING

Sanctions?

At this point, it seems safe to say that Cooks' Island has developed a legal system. The planners are the legal officials; the plan-adopters are the legislators; and the plan-apppliers, the judges. The master plan is the constitution that defines their offices. The plans created and applied by these officials pursuant to the shared plan are the laws of the system: the policy directives are the duty-imposing rules and the authorizing policies are the power-conferring ones. Finally, the islanders all act according to plan. They are law-abiding citizens.

Some might object and deny that Cooks' Island has a genuine legal system because its plans make no provisions for penalties in cases of disobedience. But as I argued in Chapter Three, it seems to me a mistake to consider sanctions to be a *necessary* feature of law. There is nothing unimaginable about a sanctionless legal system; in fact, we have just imagined one. The Cooks' Island legal system makes no provisions for sanctions but it has a constitution, a legislature and judges. It has norms that confer powers and impose duties. It maintains order, redistributes wealth, protects the moral rights of parties, provides facilities for private ordering, solves coordination problems and settles disputes. This legal system is sanctionless not because it could not impose sanctions; after all, to impose sanction merely requires that certain types of plans be adopted. Rather, sanctions are not imposed because its designers did not think them necessary. The islanders all accept the legitimacy of the group plans and,

as a result, abide by them. And when they make mistakes, they voluntarily make amends. Sanctions would simply be otiose in such a setting.

Sanction-oriented theorists often discount the possibility of sanctionless legal systems because they cannot imagine *why* such a legal system would exist. What would be the point of a community having law if its members are willing to listen to the existing social or moral norms regardless? The story of Cooks' Island rebuts this concern. The islanders' decision to develop a legal system was not motivated by distrust. They had no problem relying on one another to follow the policies they created. Their problem was that they could not create enough policies to follow.

As we saw, the Cooks' Islanders were motivated to develop a legal system as part of their effort to break a potentially destructive dynamic. On the one hand, population growth on the island led to an increase in the need for policies to guide and coordinate behavior. Yet, the same growth also amplified the cost of producing and applying such policies. At some point, the costs associated with improvisation, spontaneous ordering, private bargaining and communal consensus became so great that the demand for policies outstripped the island's ability to supply them. In an effort to radically reduce the costs of planning, the islanders were led to construct a hierarchical, impersonal and shared form of social planning. In doing so, they were able to adopt and apply the policies they needed and thereby solve the moral problems that more expensive or risky methods of planning could not.

The Circumstances of Legality

The residents of Cooks' Island may be atypical in their level of communal spirit and moral virtue, but their social problems are hardly unusual. For it is plausible to suppose that any modestly sized community will face similar questions about ownership, contractual obligations, duties of care to one another, proper levels of taxation, limitations of public power, legitimacy of state coercion and so on. Moreover, like the islanders, it will find that resorting exclusively to non-legal forms of planning is an inefficient or inadequate way of resolving these questions. To settle the doubts and disagreements of its members in a cost-effective manner, or even at all, requires sophisticated techniques of social planning such as those provided by legal institutions.

I am going to refer to the social conditions that render sophisticated forms of social planning desirable as the "circumstances of legality."¹⁸ The circumstances of legality obtain whenever a community has numerous moral problems whose solutions are complex, contentious or arbitrary. In such instances, the benefits of planning will be great, but so will the costs and risks associated with non-legal forms of ordering behavior, such as improvisation, spontaneous ordering, private agreements, communal consensus or personalized hierarchies. Indeed, the costs and risks of non-

legal planning may be so large as to be prohibitive. Communities who face such circumstances, therefore, have compelling reasons to reduce these associated costs and risks. And in order to do so, they will need the sophisticated technologies of social planning that only legal institutions provide.

Although the circumstances of legality emerged on Cooks' Island due to its system of private property, we can easily imagine similar conditions cropping up in a system of common ownership as well. In fact, in a collectivist regime where mass mobilization of the community is needed to produce the necessities of daily life, the value of sophisticated technologies of social planning will be especially great. Massively shared activities, as we saw in the previous chapter, can only be managed through the development of a dense network of plans and planners. Relying solely on non-legal methods to coordinate collective action on such an immense scale will eventually prove to be inferior to legal forms of social planning.

Once we recognize the extent to which even modestly sized human communities require sophisticated methods for guiding, organizing and monitoring conduct, we can begin to see legal institutions in a new light. According to what I will call the "Planning Theory of Law," legal systems are institutions of social planning and their function is to compensate for the deficiencies of alternative forms of planning in the circumstances of legality. Legal institutions are supposed to enable communities to overcome the complexity, contentiousness and arbitrariness of communal life by resolving those social problems that cannot be solved, or solved as well, by non-legal means alone.

Of course, the aim of the law is not planning for planning's sake. If legal systems were merely supposed to adopt and apply plans regardless of method or content, the task would be better served by flipping a coin. Rather, the law aims to compensate for the deficiencies of non-legal forms of planning by planning in the "right" way, namely, by adopting and applying morally appropriate plans in a morally legitimate manner. Legal systems are improvements over alternative forms of planning, and hence serve their function, whenever the total reduction in the costs of planning more than offsets any increase in the moral costs engendered by the switch. The task of institutional design, therefore, is to ensure that the legal process does not render mistakes so likely, or use methods that are so unsavory, that the moral benefits of switching to law vanish. Indeed, as we will see in the next chapter, legal systems are sometimes designed to *increase* the cost of adopting plans so as to *decrease* the risk that bad plans will be adopted. The worry here is that social planning may be *too* easy and, thus, overly responsive to the momentary passions of the electorate or self-interest of politicians. Throwing some sand in the gears may slow down the legal process enough to improve its ultimate reliability.

The Planning Theory of Law's central claim – that the law is first and foremost a social planning mechanism – is supported by two considerations. First, it explains why we consider law to be valuable. It

is, for example, a widely shared assumption of political theories that agree on virtually nothing else that the law is an indispensable social institution in the modern world. This belief is of course entirely sensible when we consider the benefits and costs of various methods of planning in such settings. Given the complexity, contentiousness and arbitrariness of modern life, the moral need for plans to guide, coordinate and monitor conduct are enormous. Yet, for the same reasons, it is extremely costly and risky for people to solve their social problems by themselves, via improvisation, spontaneous ordering or private agreements, or communally, via consensus or personalized forms of hierarchy. Legal systems, by contrast, are able to respond to this great demand for norms at a reasonable price. Because of the hierarchical, impersonal and shared nature of legal planning, legal systems are agile, durable, and capable of reducing planning costs to such a degree that social problems can be solved in an efficient manner.

On the other hand, when the net savings in planning costs engendered by the switch to law are low or non-existent, we tend not to judge legal institutions to be necessary or even desirable. For example, it would be quite odd to pity hunter-gatherers for not having law. Hunter-gatherers, after all, need very few rules. Since they do not cultivate land or domesticate animals, they have no need for a fixed system of property rights in real property. Moreover, because the food collected is perishable, hunter-gatherers accumulate little or no surplus and hence have no need for rights in personal property as well. Rules for the voluntary transfer of property and compensation for their damage are similarly useless to them. Not only are the benefits of social planning low but so too are the costs. For when hunter-gatherers require rules, they can either rely on custom or create them straightaway. They can deliberate among themselves about how they ought to live and arrive at some consensus or, failing that, the discontents can separate and merge with other groups. Determining whether the rules have been broken is easy both insofar as there are very few of them and collective deliberations are possible. And when a rule has been violated, communal responses are not difficult to organize. In short, hunter-gatherers do not need law because they do not face the circumstances of legality and, hence, have no need for sophisticated technologies of social planning.

Attributing a planning function to law is further supported by the observation that legal systems can be criticized not only for being evil, but also poorly designed. The Articles of Confederation, for example, were roundly condemned for their inability to regulate interstate commerce, impose taxes, raise an army, establish a system of federal courts, etc. They were also assailed for their amendment procedures which required unanimous consent to change any article. In other words, the former colonists considered their legal system defective precisely because it was an inadequate response to the circumstances of legality. Confederation following Independence generated so much complexity, contentiousness

and arbitrariness that the system could not meet the nation's new demands for social planning.¹⁹

5. LAW AS A UNIVERSAL MEANS

According to the Planning Theory then, the function of law is not to solve any *particular* moral quandary. Its function is to solve a higher-order problem, namely, the problem of how to solve moral quandaries *in general*. A community needs law whenever its moral problems – whatever they happen to be – are so numerous, and their solutions so complex, contentious or arbitrary, that non-legal planning is an inferior way of guiding, coordinating and monitoring conduct.

In this sense, laws, like intentions, are “universal means.”²⁰ Just as there are no specific ends that intentions are supposed to serve, there are no substantive goals or values that laws are supposed to achieve or realize. They are all-purpose tools that enable agents with complex goals, conflicting values and limited abilities to achieve ends that they would not be able to achieve, or achieve as well, without them.

The Problem of Bad Behavior

It is worth noting at this point that the Planning Theory contrasts sharply with a more conventional view of the law, famously expressed by James Madison in Federalist 51, when he wrote that “if men were angels, no government would be necessary.” Following this popular view, the function of the law is to solve a particular social problem, namely, *the problem of bad character*. Legal institutions are created not as general purpose technologies of social planning but rather as antidotes to the infirmities of human nature that inevitably lead people to transgress existing social or moral norms.

The problem of bad character is perhaps the most salient theme running through the classical social contract theories of the early modern period. Hobbes argued, for example, that the state of nature is a state of war because men are greedy and vain.²¹ In their desire to dominate others as well as protect themselves, they inevitably disregard their covenants of non-aggression and launch preemptive attacks against those who might attack them first. Locke also thought that individuals in the state of nature would act aggressively.²² Unlike Hobbes, however, he did not think they would do so out of callous disregard for the natural law but rather as a result of self-deception. Since people are often biased in their own favor, each side in a dispute will judge themselves justified and hence be unwilling to yield. Unable to settle their conflicts peacefully, individuals in the state of nature will resort to violence, leading to destructive cycles of feuding. Similarly, Hume believed that, in the absence of government, people will tend to ignore the principles of justice.²³ Hume attributed this non-compliance largely to irrationality: people often heavily discount the

future and seek to maximize short-term benefit over long-term gain. As a result, they routinely fail to recognize the benefits of abiding by the principles of justice.

For each of these three political theorists, the state of nature is undesirable because human nature is corrupt in some way. And, correspondingly, the law is a necessary social institution precisely because it compensates for the infirmities of human nature by ensuring that individuals abide by the existing norms. For Hobbes and Hume, the law secures compliance by threatening coercion and sanctions whereas, for Locke, it acts as a third party adjudicator, providing impartial resolution of disputes for those who consent to its authority.

In view of our earlier discussion however, this general account of the law is obviously flawed. For if the principal aim of law is to solve the problem of bad character, we would expect law to be deemed unnecessary in situations where everyone has good character. But as we saw in the last section, legal institutions can be highly desirable even though every one in the community is willing to abide by the existing norms. The residents of Cooks' Island, for example, were committed to following their shared plan, and acting in accordance with morality. The complexity and contentiousness of these normative questions, however, rendered that task difficult and costly to accomplish. Moreover, the complexity and arbitrariness of many large-scale shared activities rendered it impossible for members of the group to do what they had good ethical reasons to do. Even when they knew what moral problems they ought to rectify, they could not figure out how to coordinate their behavior so as to resolve these problems. Their sterling characters did not, in other words, diminish their need for law.

The Plan is Mightier than the Sword

Following the Planning Theory, it is a mistake to suppose that the function of the law is to solve the problem of bad character or any other particular social or moral quandary. As a contingent matter, of course, the law serves a number of important social aims. It builds roads and bridges, educates the population, finances and organizes communal self-defense, sets up markets, regulates imports and exports, controls the money supply, standardizes weights and measures, collects and redistributes wealth, arbitrates and mediates disputes, constitutes national identity and so on. It also ensures that people listen to its rules. Indeed, it would be absurd to deny that, in the modern world at least, social deviance caused by vicious character is one of the reasons why law is an indispensable social institution. It is indeed likely that life would be poor, nasty, brutish and short without legal systems maintaining order through threats of coercion.

The essential point, however, is that whenever the law properly addresses a particular social problem, it does so because, given current

social conditions, alternative methods of planning are somehow deficient. Thus, when the law is needed to combat bad character, it is because, and only because, *coercion in the absence of sophisticated forms of social planning would be expensive, ineffective or dangerous*. To appreciate this, it is important to remember that law enforcement is a shared activity and, in modern states, a quite massive one. In some cases, the roster of officers who are engaged in enforcement – police, judges, magistrates, bailiffs, clerks, wardens, guards and lawyers – contains millions of names. It is hard to imagine such groups acting together in such complex and contentious environments without an extensive network of social plans to regulate their behavior. These inherent problems can be further aggravated when alienated participants are asked to play roles that they cannot be trusted to perform in the absence of guidance and monitoring. Add to these difficulties the enormous number of coordination problems that arise in such large scale contexts and it becomes evident that improvisation, spontaneous ordering, private planning and simple forms of social planning are not adequate to guide, organize and monitor such activities.

Kelsen once described law as an “organization of force.”²⁴ Although I disagree with this claim that the law necessarily uses force, I agree that, when the law does use force, it is always organized. Both to maximize its effect and control its power, the law organizes a coercive response to social deviance through an interlocking set of social plans. The master plan which regulates all official conduct controls the procedural aspects of coercion: it selects those whose role it is to enforce the law and the procedures that they must follow in order for coercion to be permissible. The directives that are created pursuant to the master plan concern the substantive aspects of coercion: they identify those actions that warrant a coercive response. Legal officials know which behaviors to punish because other officials have issued directives informing members of the group of their rights and obligations. When enforcement personnel follow this dense network of social plans, they are able to act collectively to subdue members of the community. Thus, despite the fact that legal officials are almost always a small minority of a population, the shared agency made possible by social planning harnesses and magnifies their power, thereby enabling them to enforce the will of the law.

6. THE PRIMACY OF SOCIAL FACTS

As we have seen, the Planning Theory not only maintains that legal activity results *in* planning; it maintains that it results *from* planning as well. Legal institutions are structured by shared plans that are developed for officials so as to enable them work together in order to plan for the community. These norms set out the vertical and horizontal divisions of social labor, specifying who is authorized to formulate, adopt, affect, apply and enforce the plans and instructing them about how to engage in

these various stages of social planning. These shared plans can be thought of as *the law's plans for planning*.

Once we recognize the central role that shared plans play in the law, we can begin to address the question with which we began the last chapter. Why might one claim – as legal positivists do – that law and morality do not share the same basic ground rules? Why is the determination of legal validity a matter of a sociological, rather than moral, inquiry?

I hope that my answer to these questions is now apparent: namely, that the fundamental rules of a legal system constitute a shared plan and, as we have seen, the proper way to determine the existence or content of a shared plan is through an examination of the relevant social facts. A shared plan exists just in case the plan was designed with a group in mind so that they may engage in a joint activity, it is publicly accessible and accepted by members of the groups in question. As a result, if we want to discover the existence or content of the fundamental rules of a legal system, we must look only to these social facts. We must look, in other words, only to what we think, intend, claim and do round here.

Notice further that the existence of the shared plan does not depend on any moral facts obtaining. The shared plan can be morally obnoxious: it may cede total control of social planning to a malevolent dictator or privilege the rights of certain sub-groups of the community over others. Nevertheless, if the social facts obtain for plan sharing, then the shared plan will exist. And if the shared plan sets out an activity of social planning that is hierarchical and highly impersonal and the community normally abides by the plans created pursuant to it, then a system of legal authority will exist as well.

The crucial point here is that the determination by social facts is not some necessary, but otherwise unimportant, property of shared plans. Shared plans must be determined exclusively by social facts *if they are to fulfill their function*. As we have seen, shared plans are supposed to guide and coordinate behavior by resolving doubts and disagreements about proper action in complex, contentious and arbitrary environments. If a plan with a particular content exists only when certain moral facts obtain, then it could not resolve doubts and disagreements about the right way of proceeding. For in order to apply it, the participants would have to engage in deliberation and bargaining that would recreate the problem that the plan aimed to solve. The logic of planning requires that plans be ascertainable by a method that does not resurrect the very questions that plans are designed to settle. Only social facts, not moral ones, can serve this function.

The purpose of the master plan, we have said, is to guide, organize and monitor the shared activity of legal officials. It seeks to overcome the enormously complex, contentious and arbitrary problems associated with arranging a system of social planning. Because reasonable (and unreasonable) people can have doubts and disagreements about which social problems to pursue and who should be trusted to pursue them, it is

essential to have a mechanism that can settle such questions, creating a mesh between legal officials and leading them all in the same direction. To seek to discover the existence or content of such a mechanism by looking to moral philosophy, as the natural lawyer recommends we do, would frustrate the function of the master plan. It would require members of the community to answer the very sorts of questions that the master plan aimed to circumvent.

This objection to natural law theory might helpfully be put in a slightly different way. Consider a theory called “plan positivism.” Plan positivists believe that the existence and content of plans never depends on moral facts. Plan positivism is uncontroversially true. Everyone – even natural lawyers – accepts that people can have morally bad plans. Terrorist plots, for example, exist even though they should not be carried out from the moral point of view; rather, they exist just because terrorists share certain plans. Indeed, as I argued in the last chapter, plan positivism must be true if plans are to fulfill their function. Plans can do the thinking for us only if we can discover their existence or content without engaging in deliberation on the merits.

Since plan positivism is uncontroversially true, natural law theory will be false *if the fundamental rules of a legal system are plans*. Since the Planning Theory claims that all law is grounded in shared plans, it follows that if the Planning Theory of Law is true, then natural law theory cannot be.

My objection to natural law theory, therefore, stands or falls based on the Planning Theory of Law. For once it is conceded that legal systems are grounded in shared plans, natural law theory can be maintained only on pain of contradiction. I hope that I have given the reader at least some reason to think that the Planning Theory is a reasonable jurisprudential theory. And if the reader has been convinced, he or she will be logically compelled to accept some form of legal positivism as well.

7. THE POSSIBILITY OF LEGAL AUTHORITY

I argued in Chapter Two that a crucial test for any jurisprudential theory is its ability to solve what I have called the Possibility Puzzle. Recall that the Possibility Puzzle purports to show that legal authority is impossible. On the one hand, legal authority must be conferred by legal norms; yet, on the other, legal norms must be created by legal authority. From these two assumptions, we get a classic chicken-egg paradox. Any time we try to establish a claim of legal authority, we either enter into a vicious circle (the authority created the norm which conferred the power on the authority to create that very norm) or an infinite regress (the authority got his power from another authority, who got his power from another authority and so on).

In this final section, I want to show that the Planning Theory does indeed provide a convincing solution to this apparent paradox. Before I

can do so, however, I need to say a bit more about how legal authority is generated by the plans of a legal system.

The Ability to Plan

According to the Planning Theory, someone has legal authority only if they are authorized by the master plan of a particular legal system. But while authorization is necessary for legal authority, it is clearly not sufficient. The reason is simple: if legal authority entails the ability to plan for others, as the Planning Theory claims, then the norms adopted and applied by legal authorities must be plans. Plans, as I have argued, are special kind of norms. They are not only positive entities that form nested structures, but they are formed by a process that disposes their subjects to comply. As a result, unless the members of the community are disposed to follow the norms created to guide their conduct, the norms created will not be plans.

Thus, being authorized to plan for others does not entail that one actually has the *ability* to plan for others. A group of poor deluded souls can share a plan authorizing one of them to plan for the withdrawal of US forces from Iraq. But the one authorized will not have the ability to adopt a plan for US forces because he cannot dispose them to act in accordance with his directives.

The disposition instilled by the legal process is obviously not a brute causal one; it must be tied in some way to human motivation. Legal authorities have the ability to plan for others, in other words, only if they are able to motivate their subjects to obey under normal conditions. Of course, not every official with legal authority need be able to instill this disposition; rather, they must be part of a group of individuals who, by planning together in accordance with the master plan, can dispose most of the people to comply with most of the plans most of the time.

When members of the community consider legal authorities to be morally legitimate, encouraging compliant behavior will be relatively straightforward. By designating a standard as *the* standard to be used to guide and evaluate conduct, their subjects will take themselves to have reasons to defer and, in the normal course, will obey. If members of the community are less “cooperative,” legal authorities can dispose them to comply through various forms of intimidation. When these threats are strong and credible enough, even those who do not accept the law’s moral authority will be nevertheless motivated to follow the adopted plans.²⁵

Legal Authority and Planning Authority

As we have seen, the Planning Theory claims that a body has legal authority in a particular legal system when two conditions are met: (1) the system’s master plan authorizes that body to plan for others; and (2) the members of the community normally heed all those who are so authorized.

Legal authority will be possible, therefore, just in case it is possible for both of these conditions to obtain. Let us, then, consider each of these two conditions in turn.

First, is it possible that a shared plan authorizes some body to plan? Of course it is! As we have seen, shared plans exist when certain social facts obtain. A shared plan can authorize some person to plan for another just in case some person or persons designed the plan, at least in part, for a group, part of the plan authorizes some body to plan for another, the plan is publicly accessible, and the members of the group accept it. When these conditions obtain, a shared plan will be created and will authorize some to plan for others.

On this account, the question of how the Queen of England is the Queen is no more perplexing than the question of how my friend can be the head chef during a meeting of our club. My friend is authorized to adopt plans for club members because we have committed ourselves to defer to him. Similarly, the Queen is legally authorized to plan for others because British legal officials have accepted a shared plan which authorizes her to play a certain role in adopting legal policies and, hence, requires them to defer to her in the circumstances specified in that plan.

In the end, shared plans are able to authorize legal officials to plan for others because human beings are planning agents and are capable of guiding and organizing their actions both over time and across persons. Not only can we figure out how others should act in order to achieve some complex goal, but we can form intentions to do what we are instructed to do. In other words, we are able to create law because we are able to create and share plans.

Notice that the Planning Theory is able to secure the existence of fundamental legal rules without generating vicious circles or infinite regresses. Legal officials have the power to adopt the shared plan which sets out these fundamental rules by virtue of the norms of instrumental rationality. Since these norms that confer the rational power to plan are not themselves plans, they have not been created by any other authority. They exist simply in virtue of being rationally valid principles. Instrumental rationality, therefore, constitutes the normative bedrock on which every legal system rests.

Having shown that shared plans authorizing bodies to plan for others are possible, we should consider the second condition of legal authority, namely, whether those so authorized can motivate their subjects to comply under normal conditions. As we have just seen, there is nothing perplexing about this condition obtaining either. Members of the group might all accept a general policy to obey the law or deem those in authority to be morally legitimate. In such cases, the adoption of plans by legal officials will induce a rational requirement for those individuals to comply. Even when members of the group are not predisposed to conform to the law, the commitment of officials to carry out parts of the shared

plan that direct punishment in case of disobedience may be sufficient to motivate ordinary citizens to obey.

We can see, therefore, that legal authority is possible because certain kinds of agents are capable of (1) creating and sharing a plan for planning and (2) motivating others to heed their plans. Legal systems are possible, in other words, because certain states of affairs are possible, namely, those that underwrite the existence of a legal system's master plan and those that account for the disposition of the community to comply with the plans created under normal conditions.

It might be thought that this solution is one that the legal positivist would wholeheartedly endorse since it maintains that the authoritative status of any official does not derive from any moral fact. Even if the shared plan is morally odious and the citizens are as servile as sheep, the authorized individual will have legal authority within the existing system.

Upon closer inspection, however, the Planning Theory poses problems for legal positivism. For according to the traditional formulation, legal positivists claim that the existence and content of the law rests on social facts alone. Yet, according to the Planning Theory, the existence of the law actually depends on a normative fact as well, namely, the fact that planning agents have the rational authority to give themselves plans. To be sure, the rational authority of planners is not a form of moral authority: terrorists, as we have said, have the rational authority to create evil plots among themselves. But the fact that planners are so empowered is nonetheless a normative fact because it affects, in a very important sense, what individuals ought to do.

It turns out, then, that the conventional debate between legal positivism and natural law theory has presented legal theorists with a false choice. It is true neither that the law ultimately rests on social facts alone nor that it rests on social *and* moral facts. The Planning Theory shows that there is a third alternative, namely, that legal facts ultimately rest on social and non-moral normative facts. To be sure, this third alternative is clearly within the spirit of legal positivism, because it denies a role to moral facts in determining the existence or content of the law. But this version also departs from the letter of legal positivism by recognizing that the law does indeed rest on irreducible norms of practical rationality.

The Legal Point of View

Natural lawyers might respond to this defense of a moderate version of legal positivism in the following way: they might concede that moral facts are not needed in order to establish a social planning system. Planning agents have the rational authority to plan their social planning and the resulting shared plans are indeed created by a combination of social and non-moral normative facts. Yet, they will surely deny that these facts are sufficient *for establishing legal authority*.

Someone possesses legal authority, natural lawyers will point out, only if they have the ability to impose legal obligations. But how can anybody have this normative power simply because certain of their cronies authorize them to have such a power and members of the community acquiesce? To generate relations of legal authority, the objection goes, a group needs more than social and non-moral normative facts – it needs *moral facts* as well. In the case of Cooks' Island, it is plausible to suppose that the adopters and appliers possessed legal authority because their planning system is morally legitimate. After all, it was set up in a democratically appropriate manner and designed to serve moral aims. Their plans do confer legal authority and impose legal obligations. Once we remove the assumption of moral legitimacy, however, the claim of legality becomes highly dubious.

In order to address this objection, we should examine claims of legal authority in more detail: what are we imputing to someone when we say that they have “legal authority”? One possibility is that we are imputing a type of moral authority. On this reading, the word “authority” means the same as it does in moral contexts, namely, the power to impose moral obligations, and the word “legal” functions as an adjective, identifying this kind of moral power. We are saying, then, that the person in question has moral authority in virtue of being an official in a legal institution. Call this the “adjectival” interpretation.

It appears that the natural law objection interprets claims of legal authority adjectivally. For on the adjectival interpretation, legal authority entails moral authority and since morally illegitimate shared plans do not confer moral authority they cannot confer legal authority. To be sure, this implication would be devastating for the Planning Theory if we were compelled to accept the adjectival interpretation of legal authority claims as the only possible one. I do not however think we are. The problem with the adjectival interpretation is that it ties legal authority too tightly to moral authority and thus precludes the possibility of morally illegitimate governments. As I argued in first chapter, not only are such cases possible, but they are actual. Stalin, Hitler and Mao, to use three paradigmatic examples, all had legal authority but were morally illegitimate. The adjectival interpretation, therefore, does not permit us to make the claims about legal authority that we customarily make.

Fortunately, there is another interpretation that does permit the ascription of legal authority to morally illegitimate bodies. The key here is to recognize that, although the term “authority” in legal authority claims refers to a moral power, the word “legal” often functions as a *qualifier* rather than a modifier. When we ascribe legal authority to someone, we are not necessarily imputing any kind of moral authority to them. To the contrary, we are qualifying our ascription of moral legitimacy. We are saying that, *from the legal point of view*, the person in question has morally legitimate power. Similarly, to say that one is legally obligated to perform some action need not commit the asserter to affirming that one is

really obligated to perform that action, i.e., has a moral obligation to perform that action. The statement may be understood to mean only that from the legal point of view one is (morally) obligated to perform that action.

To understand the discourse of legality, therefore, one must recognize that it typically performs a *distancing* function. It enables us to talk about the moral conception of a particular legal system without necessarily endorsing that conception. Sure, one might say, sodomy is wrong from the legal point of view. But this assertion does not imply anything about whether sodomy is in fact wrong. In some cases, the word “legal” registers our agnosticism: we do not know or care whether the law’s normative judgments are correct—we are simply reporting these judgments—in effect, bracketing them off in a special kind of invisible inverted commas. At other times, the word “legal” signals our alienation from the legal point of view. Sodomy, we can say, is legally wrong, but it is wrong *only legally*. From our own point of view, sodomy is morally permissible.

What, then, is the legal point of view? It is not necessarily the perspective of any particular legal official. No official may accept the law’s conception of itself. The legal point of view, rather, is the perspective of a certain normative theory. According to that theory, those who are authorized by the norms of legal institutions have moral legitimacy and, when they act in accordance with those norms, they generate a moral obligation to obey. The legal point of view of a certain system, in other words, is a theory that holds that the norms of that system are morally legitimate and obligating. Thus, communism is the point of view of communist legal systems, individualism the point of view of laissez-faire capitalist systems, democratic theory the point of view of democratic systems and so on.

The normative theory that represents a system’s point of view may, of course, be false from a moral perspective. That is, the legal point of view may not coincide with the true moral point of view. Those authorized by legal institutions to act may be morally illegitimate and their actions may generate no moral obligations to obey. The point of view of a particular legal system may be like the phlogiston theory of combustion: a scientific theory that aimed to be true but missed the mark. In short, the legal point of view always *purports* to represent truly the moral point of view, even when it fails to do so.

Given this interpretation of legal authority and obligation claims, we can easily see how morally illegitimate shared plans can confer legal authority. For to ascribe legal authority to a body in a particular legal system is to assert that, from the point of view of that legal system, the body in question is morally legitimate.²⁶

- (1) X has legal authority over Y in system S \leftrightarrow From the point of view of S, X has moral authority over Y.

The point of view of that legal system will ascribe moral legitimacy to a body just in case its norms confer legal power on that body. Since on the Planning Theory the legal norms that confer legal authority are subplans of the system's master plan, the legal point of view will ascribe moral legitimacy to a body when its master plan authorizes that body to so act.

- (2) From the point of view of S, X has moral authority over Y
 \leftrightarrow The master plan of S authorizes X to plan for Y.

It follows from (1) and (2) that a body will have legal authority in a particular legal system just in case the system's master plan authorizes that body to so act.²⁷

- (3) X has legal authority over Y in system S \leftrightarrow The master plan of S authorizes X to plan for Y.

We can see thus that the truth of legal authority claims is not dependent on the moral legitimacy of a system's master plan. A body has legal authority in a system because the master plan authorizes it. And, as we have seen, master plans are capable of authorizing someone to plan for another just in case the members of a group exercised their planning authority in the right sort of way.

8. COMPARISONS

At this point, I want return briefly to our criticisms of previous positivist attempts to solve the Possibility Puzzle in order to see if the planning solution escapes these same objections. The first attempt I discussed was Austin's attempt to ground all legal systems in habits of obedience. Recall that we made three principle objections to this reduction. First, habits are not normative activities. When someone habitually engages in some activity, he neither subjects his conduct to the evaluation of a norm and nor describes his actions using normative terminology. For this reason, it would be bizarre to describe a judge's hewing to the law as an instance of *habitual* obedience. The judge not only evaluates his own law-applying actions by norms, but expresses his own actions with normative-loaded words like "ought," "must," "right," and "wrong." Second, habits of obedience cannot account for the continuity of legal authority. Once Rex I dies, there will necessarily be a lag time before a new habit of obedience to his successor develops. Yet, the successor, Rex II, has legal power from the time of the death of Rex I and need not wait until a habit of obedience forms. Third, Austin's reduction to habits cannot account for the persistence of law. The laws made by Rex I carry over into the realm of Rex II even though Rex I can

no longer be the object of habitual obedience and Rex II has not reenacted his laws.

Because the planning model replaces habits with plans, it has no problem explaining the normative nature of legal activity. When a judge applies the law, he is following a plan directing him to do so. He takes the plan as settling for him how he ought to proceed. It is natural, therefore, for him to express his actions using normative terminology. For when one follows a plan, one takes the plan as determining what one "ought" to do and what it would be "wrong" not to do. Likewise, plans neatly account for the continuity and persistence of law. The shared plans which underwrite legal systems establish offices and accord power to different inhabitants as soon as they fulfill the requirements for accession to their offices. They need not wait for habits of obedience to kick in. These same plans accord validity to the planning of past inhabitants, hence ensuring the persistence of their plans.

Hart tried to solve the Possibility Puzzle via his practice theory of social rules. On the practice theory, all social rules are practices. The rule of recognition can thus be brought into existence by being practiced, because it is a social rule and social rules are just practices. We saw, however, that the practice theory of social rules is terribly flawed. First, it seems wrong to conceive of rules as practices, rather than merely created by them. Rules are abstract entities, whereas practices are concrete events that reside within space and time. Second, all social rules cannot be practices because moral rules can exist and not be practiced. Moral rules would lack their critical power if one could only apply them to groups that already accepted them. If all social rules are practices, then Hart has not given us an explanation of how the social rule of recognition can be brought into existence simply by being practiced.

Following the Planning Theory, however, shared plans are not identical with practices. They are abstract entities that members of a group accept and are thus the objects of their intentions. Furthermore, the Planning Theory does not claim that all social rules are shared plans. It merely claims that the social rule that lies at the foundation of law is a shared plan. It is a shared plan because its function is to guide and organize the planning behavior of the group of officials. And since the norm is a shared plan, it must be practiced by that group if it is to exist. They must accept it and it must dispose them to act accordingly.

The final positivist theory we explored was Postema's and Coleman's coordination convention account. On this account, the rule of recognition is like a rule of the road: it is a coordination convention and coordination conventions are behavioral regularities. The rule of recognition can thus be brought into existence simply by being practiced. This theory, however, suffers from a fatal flaw: that is, it places too many restrictions on the type of motivations that judges must possess. First, coordination conventions only arise when the players deem the game to have arbitrary solutions. However, there is no justification for insisting

that most officials in every conceivable legal system must treat their system of government like the rules of the road, as just one out of several acceptable alternatives. Second, coordination conventions require that the players take the fact that others are acting in a similar manner as a *reason* for them to act similarly. This seems to exclude cases of alienated players, those who participate in the practice for no other reason than that it is in their self-interest to do so.

The Planning Theory, by contrast, does not insist that officials have any sort of beliefs or motivations. They may think that there is one way, and only one way, to structure government and accept the shared plan precisely because they think that the plan has hit on this best way of engaging in social planning. Or they may think that the shared plan sets out the worst blueprint for social planning but accept it purely for reasons of self-interest. As long as officials accept the plan, play their part in it, resolve their disputes peacefully and openly, and all this is common knowledge, it does not matter why they are doing so. A shared plan will exist and their planning pursuant to it will be a shared activity.

Finally, the Planning Theory also avoids an important objection that bedevils each of the other jurisprudential theories we have seen. There is nothing in the accounts of Austin, Hart, Postema or Coleman to suggest that legal activity is a shared activity. But as I have argued, judges are not simply independent arbitrators who decide disputes for their clients. Nor are legislatures norm-producing institutions unconnected to courts. Each is part of, and plays a role in, the collective activity of communal governance. As we have seen, the Planning Theory treats legislators and judges as engaged in a shared activity of social planning. The plan that governs the conduct of legislators and judges is shared by them and hence binds them together into a group. Officials each play the planning role assigned to them in the expectation that, by doing their parts, they will bring about the political objectives of the collective legal enterprise.

9. FLIPPING THE SOUL-STATE ANALOGY

In The Republic, Plato attempted to reveal the nature of justice in the soul by constructing a just state and extrapolating from the latter to the former. He reasoned that it is easier to discern the properties of an entity when it is depicted on a very large canvass. “[I]s not a State larger than an individual? ... Then in the larger the quantity of justice is likely to be larger and more easily discernible. I propose therefore that we enquire into the nature of justice and injustice, first as they appear in the State, and secondly in the individual, proceeding from the greater to the lesser and comparing them.”²⁸

In the last two chapters, I have proceeded in precisely the opposite direction. I began by investigating how individuals govern their own individual actions, moved to the case of small shared activities then to instances of massively shared agency, back to households and, finally, on

to political government. My strategy, in other words, has been to explore familiar forms of agential management, i.e., individual and small-scale planning, and then extrapolate to lesser understood cases, i.e., massively shared agency and political government.

Thus we began with the simple fact, emphasized by Bratman, that human beings are planning creatures: that is to say, we not only have the capacity to act purposively, something we share with non-human animals, but also, the ability to form and execute plans. We saw that our need for planning arises from the complexity of our goals, the limitations of our abilities and the pluralism of our values and preferences. In order to achieve complicated goals, we must assemble long sequences of actions correctly; in order to act together, we must be able to coordinate many of our actions with one another. Deliberation at the time of action will often be too costly or risky due to the complexity, contentiousness or arbitrariness of the choice situation. Planning guides and organizes our behavior so that we can accomplish ends that we would not otherwise be able to accomplish, or accomplish as well.

Just as social creatures with complex goals, conflicting values and limited abilities require a special kind of psychology, we also saw that they need a special kind of social structure in order to flourish. In any community of modest size, the moral problems and opportunities that arise will not be soluble exclusively through improvisation, spontaneous ordering, private agreements, social consensus, personalized hierarchies or some combination thereof. Only a sophisticated mechanism of social planning, like the hierarchical, impersonal and institutional one exemplified by the law, can address these issues in an effective and efficient manner. Laws, we might say, play the same role in social life that intentions play in individual and shared agency: they are universal means that enable us to coordinate our behavior intra- and inter-personally.

By reversing Plato's course and moving from the small to the large, I have tried to show that political government is continuous with the intimately familiar activities we understand so well. Far from being an esoteric activity completely divorced from ordinary experience, I have argued that the law is simply a sophisticated apparatus for planning in very complex, contentious and arbitrary communal settings.

But the planning model of law that I have developed here not only allows us to understand the unity of legal systems; it accounts for their diversity as well. As the last two chapters have emphasized, planning is not a monolithic activity. It can be centralized or decentralized, top-down or bottom-up, progressive or conservative, directive or authorizing, coercive or non-coercive, dictatorial or democratic and so on. The institutional arrangements characteristic of a legal system typically reflect judgments about the best modes of social planning, and these judgments in turn reflect underlying attitudes about the proper moral aims of law and who can be trusted to attain them.

By returning to this soul-state analogy throughout the rest of the book, I intend to draw on our practical knowledge of planning in order to uncover and explore the identity of law. I will argue that many vexing jurisprudential puzzles can easily be solved by recognizing that legal activity is planning activity and that laws are just plans, or plan-like norms. Most importantly, I will try to make good on the promise made in the first chapter, namely, to show how a philosophical investigation into the nature of law can help us understand the practice of legal interpretation. For if the planning model is right, the interpretation of law is nothing but the interpretation of plans. Knowing how to interpret a plan does not require a Ph.D. or J.D. – it is an activity that we do hundreds of times a day. We know how to do this simply because we are planning creatures.

¹ Michael Bratman, Intention, Plans and Practical Reason, 28-30 (1987).

² Id at 32-35.

³ Id at 30-32. It should be noted that the pragmatic justification for the requirements of rationality, as well as the content of the requirements themselves, set out in the text is the subject of philosophical controversy. For alternative accounts, see John Broome, "Normative Requirements," 12 *Ratio* 398-419 (1999); R. Jay Wallace, "Normativity, Commitment, and Instrumental Reason," 3 *Philosophers' Imprint* 1 (2001) and Kieran Setiya, "Cognitivism about Instrumental Reason" *Ethics* (forthcoming). For Bratman's rejoinder, see his "Intention, Belief, Practical, Theoretical" in Jens Timmerman, John Skorupski, and Simon Robertson, eds., Spheres of Reason (forthcoming) and "Intention, Belief and Instrumental Rationality," in David Sobel and Steven Wall, eds., Reasons for Action (forthcoming).

⁴ Similarly, it follows that when a planner recognizes that her plan applies, she should not deliberate about the merits of the case at hand. The plan is supposed to settle the matter of whether he should act in a certain way and, thus, to deliberate before execution undermines the fundamental purpose of the plan. The planner rationally executes the plan because he adopted it and it applies and for no other reason.

⁵ See, e.g., David Hume, A Treatise on Human Nature, Part II, Section II (1739).

⁶ See, e.g., Ludwig von Mises, "Economic Calculation in the Socialist Commonwealth," in Friedrich von Hayek, ed., Collectivist Economic Planning (1935); Friedrich von Hayek, Individualism and Economic Order, chs. 7-9 (1948).

⁷ See, e.g., Joseph Raz, The Morality of Freedom, ; Jules Coleman, Risks and Wrongs, 62-64.

⁸ How would we know that Dudley and Stephens did not intend to paint the house? Suppose that halfway through his painting the fresh coat, Stephens announces that he quits. Dudley replies that it doesn't matter to him – after all, he will get his money regardless. If we assume that Dudley is rational, sincere and hasn't changed his mind once he accepted the job, we can infer from this exchange that Dudley never formed a plural intention. This is so because to intend that they paint the house entails a rational commitment on the part of Dudley to the joint activity of their painting the house. This commitment, if present in a rational participant, must express itself in some form of action designed to result in their painting of the house. Dudley might pick up the brush and paint the house himself, or notify me of Stephens's departure, or try to convince Stephens that he shouldn't leave. The fact that Dudley does nothing indicates that he is not so committed. He has the singular intention to do as I say and, hence, to scrape the paint off the house, but not the plural intention that *they* paint the house. Yet, despite the fact that Dudley is not so committed to the joint activity, it will be true that they share a plan if Stephens changes his mind and returns to the job.

⁹ Because Dudley and Stephens do not intend to act together, they are not subject to the same rationality constraints as Henry and I are. Dudley must fill in all his subplans related to scraping, ensure that they are consistent with all his other subplans and not reconsider prior decisions absent a good reason. Dudley must also make sure that he does not totally interfere with Stephens' painting. But beyond that, Dudley need not worry about Stephens' subplans. Suppose Stephens does not know how to paint. Dudley is under no rational obligation to help Stephens paint the house; after all, Dudley gets paid regardless of whether Stephens paints. Of course, since I know that Dudley is not rationally obligated to help Stephens, I should give Stephens very detailed instructions on how to paint the house. Unlike them, it is rationally incumbent on *me* to ensure that the house painters know what they are doing.

¹⁰ The analogy here is to individual agency: just as individual action is individual behavior explained by an individual plan, shared action is group behavior explained by a shared plan.

¹¹ For a fascinating discussion of these colossal failures of social planning, see James C. Scott, Thinking Like a State: How Certain Schemes to Improve the Human Condition Have Failed (1998).

¹² Friedrich Hayek, Road to Serfdom, 37

¹³ *Id.* at 36-37.

¹⁴ The reason that the stipulation is prefaced with "For the purposes of island land-use policies" is that the plan does not require the group to consider the family in question to be owner for, say, tithing purposes. As is the case in many jurisdictions, someone can be deemed the owner (to have title) for the purpose of determining who bears the risk of loss, but not for whether they owe property or sales tax.

¹⁵ Moreover, since we designed our system of private property and market exchange to increase the amount of food available for the group to eat, our policies are *sub-plans* of the overall shared plan to boost the agricultural output of the island. They organize the behavior of the group so that the goal of the shared plan can be achieved. It turns out, then, that market activity – an activity that has traditionally been thought to be the paradigm of individualism – on our island is a shared activity. As long as we seek to maximize our own wealth subject to the cost internalizations mandated by our shared plan, we will collectively maximize food production.

¹⁶ Robert Ellickson's Order Without Law is the *locus classicus* for how communities are often able to order their affairs without resorting to law. The discussion in the text is much indebted to Ellickson's discussion.

¹⁷ On the importance that offices play in routinizing authority, see Max Weber, Economy and Society, ed. G. Roth and C. Wittich, Chs. 11 and 14 (1968).

¹⁸ The term "circumstances of legality" is a reference, of course, to David Hume's famous doctrine of the "circumstances of justice." Hume argued that the justice is a virtue only in situations of moderate conflict. He identified the circumstances of justice as ones of (1) modest scarcity, (2) limited altruism, (3) rough equality and (4) moderate social interdependence. See, e.g., David Hume, An Enquiry Concerning the Principles of Morals, Sec. 3, Pt. I. Other philosophers have added further conditions. Rawls, for example, claims that justice is appropriate also where people have divergent conceptions of the good life. John Rawls, A Theory of Justice, 109–12. See also Hart, The Concept of Law, 193–200. The Planning Theory does not take the circumstances of justice to be sufficient for rendering law valuable. For the circumstances of justice can obtain for a certain community, but the problems of justice posed may be resolvable through non-legal forms of social ordering. Jeremy Waldron speaks of the "circumstances of politics," which he identifies as obtaining whenever a group of people must act together over time but persistently disagree about the principles of justice that will regulate their joint activities. See, e.g., Jeremy Waldron, The Dignity of Legislation, 153-54 (1999). On the Planning Theory, the circumstances of legality include the circumstances of

politics, but are not exhausted by it. Waldron's account unfortunately neglects the importance that complexity and arbitrariness play in rendering legal institutions valuable.

¹⁹ To say that the function of the law is to resolve those moral problems that cannot be solved as efficiently by custom, tradition, persuasion, consensus and promise is not to suggest, of course, that the law *never* relies on these other mechanisms. When simpler methods of organizing behavior work, it would be irrational to abandon or overturn them in favor of accomplishing the very same ends through more sophisticated methods. Thus, when customs or contracts solve moral problems, the law will typically let them stand. And the master plan may itself be filled out through customary means or private agreements as well. In other cases, the law will interfere with private orderings in a modest way, by supporting them when they emerge, such as when it enforces certain types of customs and contracts, and by empowering others to form when they would not otherwise, such as when it authorizes persons to engage in various types of property and contractual relationships. In these cases, the task of the law is to use minimally invasive forms of social planning in order to promote the solution of moral problems through private ordering.

²⁰ Bratman, Faces of Intention, 5.

²¹ See, e.g., Thomas Hobbes, Leviathan, Chs. 13-15 (1668).

²² See, e.g., John Locke, Second Treatise of Government, Chs. 2-3, 8-9 (1690).

²³ See, e.g., David Hume, A Treatise on Human Nature, Part II, Section VII (1739).

²⁴ Hans Kelsen, General Theory of Law & State, 21 (1949) ("The law is an organization of force.")

²⁵ Notice that while the authorization to plan for others merely depends on the mental states of legal officials, the ability to plan depends on the mental states of ordinary citizens as well. Legal authorities can motivate obedience in the normal course only if members of the community are sufficiently virtuous, submissive or cowardly. Trying to plan for a "stiff-necked people" may simply not be possible and, when this is so, legal authority will not be possible either.

²⁶ Since a legal system exists only when it is generally efficacious, i.e., members of the community normally obey the law, X will have legal authority over Y in S only when S is generally efficacious. The legal point of view, therefore, ascribes moral legitimacy only to those who are authorized by a master plan of a generally system.

²⁷ As mentioned in the previous note, X will have legal authority over Y in S only when S is generally efficacious. Hence, it will not be sufficient for a master plan to authorize someone to plan in order for that person to have legal authority. The master plan must be the plan of a generally efficacious planning system.

²⁸ Plato, The Republic, BK II (trans. Jowett)