

DRAFT

PLEASE DO NOT CITE OR DISTRIBUTE¹

June 2000

Recipes and Songs
Towards a Theory of Production

Beth Preston

Department of Philosophy
and
Artificial Intelligence Center
University of Georgia
Athens, GA 30602-1627
(706)542-2819
FAX (706)542-2839
epreston@ai.uga.edu

¹Permission to distribute only to students in Prof. Rapaport's CSE 410/510//PHI 498 Philosophy of Computer Science course, Spring 2004, granted by Beth Preston via email, 16 Oct 2003.

Listen! Wasn't that the bell? Damn! the day and the dance begin and we don't know the schedule! We have to improvise—all the world improvises its day. Let us proceed today as all the world does!

Friedrich Nietzsche, *The Gay Science*, §22

1 Introduction

There is an area of philosophical research called action theory which traces its descent back to Aristotle. Now Aristotle emphatically distinguished action (*praxis*) from production of artifacts (*poiesis*).² But there is no corresponding area of research called production theory. Indeed, action theorists routinely assimilate production to action, as evidenced, first, by the absence of discussion in the action theory literature of Aristotle's action-production distinction; and second, by the use of examples which, although mostly actions in Aristotle's sense, occasionally (and without comment) include productions. For instance, Myles Brand lists building a bridge as a paradigm case of action right alongside raising your arm and buying a loaf of bread (1970: 3).³

The most important reason for this assimilation is that the basis on which Aristotle distinguished production from action does not have much valence for contemporary analytic philosophers. Aristotle says that production has an end outside itself, whereas action does not (*NE* 1140^b5 – 10). In other words, craft (*techne*) aims at creating a freestanding product, whereas practical wisdom (*phronesis*) aims at doing well. Thus the end of action, excellence, is internal to the action itself. On the other hand, Aristotle describes the rational capacities involved in generating action and production in much the same way—both involve deliberating about how to attain our ends with regard to things we can change (*NE* 1140^a – 1140^b10). In short, Aristotle distinguishes production from action not because of any difference in the process involved, but solely because of a difference in the status of the result of that process. But action theory concerns itself almost exclusively with the explanation of action, i.e., with questions about the nature of the process by which actions are generated.⁴ So Aristotle's motivation for distinguishing action from production is not operative in this contemporary context, and the consequence is the assimilation of production to action.

I do not hold a brief for Aristotle's rigorous distinction between production and action—indeed, I would be prepared to argue against it. But for the purposes of this paper I will

²See for example *Nicomachean Ethics* 1140^a1 – 10. References to Aristotle will henceforth be incorporated in the text using *NE* as the abbreviation for *Nicomachean Ethics* and *M* for the *Metaphysics*.

³In the Continental tradition, on the contrary, heroic attempts have been made to counteract the threat of assimilation and reestablish the action-production distinction on a firmer basis. Such a project lies at the heart of Jürgen Habermas's longstanding critique of Marx's production paradigm, and his own theory of communicative action. See Grumley (1992) for an overview. Similarly, it underwrites Hannah Arendt's (1958) analysis of the *vita activa*, in which she argues for a rigorous three-fold distinction between labor, work (fabrication), and action. What I have to say has implications for this project too, but I will not be able to present them here.

⁴See the introduction to Mele (1997) for an overview.

consider production as a distinct phenomenon subject to investigation in its own right. I have two reasons for this. First, I am pursuing a philosophical account of artifacts and other use objects,⁵ and this requires an account of the production of artifacts in any case. Second, on the assumption that production really is a subspecies of action, my account of production suggests a critique of certain aspects of action theory.

In the first part of this paper I outline an account of production descending from Aristotle *via* Karl Marx and Randall Dipert. I contend that in spite of some surface variation, the underlying account has remained substantially the same. I characterize this standard account in terms of a model I call ‘centralized control.’ In the second part of the paper I explain why the centralized control model is not an adequate account of production, and I propose and defend an alternative model I call ‘collaborative improvisation.’ Finally, I briefly sketch the proposed critique of action theory by showing that its basic account of intentional action is committed to the centralized control model, and that the criticisms I bring against this model in the context of production generalize in interesting ways to the wider context of action.

2 The Standard Account of Production

A central topic in action theory concerns the explanation of intentional action, i.e., the question of how intentional states of human beings are involved in the generation of actions. To the extent that there has been a separate theory of production, it has revolved around this same question. In action theory, it used to be thought that intentions to act are reducible to beliefs and desires. More recently it has been argued that to have an intention is to have something over and above a desire and some relevant beliefs. First, to have a desire and to be committed to fulfilling that desire are two different things. Second, representing an action crucially involves planning how to carry it out, rather than merely entertaining beliefs about it. Consequently, having an intention to act is now commonly characterized as having a commitment to executing a mental plan.⁶ Production theory has followed a similar trajectory in characterizing the genesis of artifacts.

I will be focusing on the purely representational aspect of the producer’s intentional state (the plan, in other words), and leaving the the motivational aspect (the commitment) aside. As commonly construed, there are two possible *loci* of interaction between the artifact and such representational states of the producer: the antecedent design phase and the subsequent construction phase. As we shall see, the standard account of production has focused almost exclusively on an antecedent design phase as the crucial aspect of production. The construction phase has been viewed as secondary at best. The standard account of production also routinely characterizes production in terms of the activity of a single individual, and this individualistic emphasis has become increasingly explicit over time. These two themes—an emphasis on an antecedent design phase, and on individual producers—will be explicated

⁵See Preston 1998a, 1998b, and forthcoming.

⁶See Mele (1997: 16-20) for an overview and references to the current literature.

by examining the accounts of production in the work of Aristotle, Karl Marx, and Randall Dipert.

2.1 Aristotle

For Aristotle, production consists in impressing a form on matter. We produce neither the form nor the matter, but only the union of the two (*M* 1033^a24 – 1034^b5). Doing this requires thinking about how to bring the desired form into union with suitable matter.

Hence a craft (*techne*) is the same as a state (*hexis*) involving true reason concerned with production (*poietike*). Every craft is concerned with coming to be; and the exercise of the craft is the study of how something that admits of being and not being comes to be, something whose origin is in the producer and not in the product. (*NE* 1140^a8 – 15; Irwin’s translation)

Aristotle describes with some precision the pattern of thinking involved.

Now the healthy is generated when a man thinks as follows: since health is so-and-so, if the subject is to be healthy it must have such-and-such, let us say uniformity, and if uniformity, then warmth; and he always thinks in this manner until he arrives at something final which he himself can produce. Then the motion from this instant onward, which here is a motion towards health, is called “production”. Thus, it turns out that in a sense health is generated from health, and a house from a house (that is, the material house from the house without matter), for the medical art and the building art are the forms, respectively, of health and of the house. . . . Of the generations and motions just considered, one of them is called “thinking” (*noesis*) and the other “production” (*poiesis*); thinking occurs from the principle or the form, production from the end of thinking and thereafter. (*M* 1032^b15 – 20; Apostle’s translation)

The first thing to notice here is that the representational component of the intention includes both a specification of the thing to be produced (the form in the mind of the producer) and a specification of the steps by which this form may be realized in matter. So at the end of the thinking process you have in your mind a mental design for the product, complete with step-by-step instructions for constructing it. Second, this mental design is finished prior to the production proper, the actual making. So for Aristotle there is a clearly demarcated antecedent design phase and subsequent construction phase. Moreover, the construction phase is an unintelligent execution process—it simply realizes the design by executing the embedded instructions. So the real interest of the production process as a whole (craft) lies in the intentions to produce laid out beforehand in the mind of the producer during the design phase. This separation of the design phase from the construction phase, and concomitant emphasis on design, is the first common theme we shall see repeated historically in production theory.

One slight qualification must be noted here. Aristotle says craft is a state, or habit (*hexis*) involving (*meta*) true reason. Lack of craft (*atechnia*) is a habit involving false reason (*NE* 1140^a20 – 25). This makes craft analogous to virtuous action, which is a habit of choosing in accordance with a mean defined by true reason (*NE* 1106^b36 – 1107^a2). So production requires not only thinking, but also habits which might not be fully articulable. Presumably these would come into play during the construction phase. For example, in order to execute instructions for building a house, the builder has to rely on a whole congeries of motor habits involved in using hammers, handling boards, and so on. This shows that Aristotle's view is not simplistically rationalistic, and it makes the construction phase independently interesting insofar as you might want to study the nature and genesis of such habits. However, it is still the design phase which is the crucial one for Aristotle, since the character and quality of the thinking involved in it makes the difference between having craft and not having it, i.e., between being able to produce and not being able to produce.

The second common theme, the individualistic understanding of production, does not come to the fore in Aristotle, since he does not explicitly discuss production by individuals *versus* production by groups. He does routinely speak of the producer as a single individual, though—typically, a skilled artisan working alone, like an architect or a physician. Thus Aristotle's account of production is at best implicitly individualistic.

2.2 Marx

For Marx, productive labor is *the* fundamentally human form of activity. It takes on specific forms under historically local conditions, but it has a core structure which is universal.

We presuppose labour in a form in which it is an exclusively human characteristic. A spider conducts operations which resemble those of the weaver, and a bee would put many a human architect to shame by the construction of its honeycomb cells. But what distinguishes the worst architect from the best of bees is that the architect builds the cell in his mind before he constructs it in wax. At the end of every labour process, a result emerges which had already been conceived by the worker at the beginning, hence already existed ideally. Man not only effects a change of form in the materials of nature; he also realizes [verwirklicht] his own purpose in those materials. (Marx, 1883/1976: 283-4)

Like Aristotle, Marx distinguishes an antecedent design phase from a subsequent construction phase, characterizes the latter as the execution of a design fully laid out in advance in the mind, and sees in this design not only a representation of the thing to be made but also a specification of how to make it. But there is one salient difference from Aristotle here. Rather than being just a realization of an independently existing form, the artifact is expressive of the producer's own needs and purposes. This marks Marx's account as a modern rather than an ancient conception of the role of production in human life. Specifically, it is an interested, economic conception rather than a disinterested, aesthetic one. For Aristotle, the producer is merely a calculator, someone who takes pre-given elements (the form and the matter) and

figures out how to realize the one in the other. For Marx, when you make something you do not merely change the form of the object you work on; you also endow it with a use-value corresponding to, and expressive of, your own historically conditioned needs and purposes. So Marx's laborer is creative in a way Aristotle's artisan is not.

Marx, like Aristotle, routinely speaks of the producer as a single individual. But in Marx's case there is some reason for thinking this is more than a *façon de parler*. There is a difference between conceiving of production as inherently social—which Marx does—and conceiving of it as inherently collaborative, which he does not. For Marx, production is inherently social, first, because the individuals who engage in it are the individuals they are only in virtue of the social conditions in which they were raised; and second, because the mode of production in which individuals engage is informed by historically local social conditions. But you are fully social in this sense even when *not* collaborating.

The individual *is* the *social being*. The expression of his life—even if it does not appear immediately in the form of a *communal* expression carried out together with others—is therefore an expression and assertion of *social life*. (Marx, 1932/1967: 306)

But what did Marx think about collaboration *vis-à-vis* individual production? Collaboration is necessary for, and typical of, large scale capitalist production, as opposed to the preceding era of peasant agriculture and independent artisanry. But capitalist collaboration is an ambivalent phenomenon—it increases productivity, but simultaneously deepens exploitation of the workers.

When the worker co-operates in a planned way with others, he strips off the fetters of his individuality, and develops the capabilities of his species. (Marx, 1883/1976: 447)

Moreover, the co-operation of wage-labourers is entirely brought about by the capital that employs them. . . . Hence the interconnection between their various labours confronts them. . . as a plan drawn up by the capitalist, and, in practice, as his authority, as the powerful will of a being outside them, who subjects their activity to his purpose. (Marx, 1883/1976: 449-50)

But what would happen to collaboration in a communist society? This comes out most clearly in Marx's early discussion of what he calls free human production.

(1) In my *production* I would have objectified my *individuality* and its *particularity*, and in the course of the activity I would have enjoyed an individual *life*; in viewing the object I would have experienced the individual joy of knowing my personality as an *objective, sensuously perceptible, and indubitable* power. (2) In your satisfaction and your use of my product I would have had the *direct* and conscious satisfaction that my work satisfied a *human* need, that it objectified *human* nature, and that it created an object appropriate to the need of another *human* being.

What happens so far as I am concerned would also apply to you. (Marx, 1932/1967: 281)

Arguably, if a product is to be the expression of an individual's particularity, then the production process, or at least the design phase, must be an individual activity rather than a collaborative one. To the extent that I collaborate with you in the design of a product, that product is not an expression of *my* individuality, but an expression of yours; or worse yet, a compromise which does not express the individuality of either one of us. Moreover, the point at which others enter the picture here is as *consumers* of my product, not as co-producers. Producer and consumer have a cooperative relationship only in the minimal sense of reciprocal relations of production and consumption; not as collaborators in production itself. Thus Marx sets up individual production as the *ideal* form to be achieved under true communism, and strengthens the implicit individualism of Aristotle's account by explicitly endorsing the paradigm of the individual artisan as the realization of this ideal.⁷

2.3 Dipert

The traditional account is crystallized beautifully in Randall Dipert's pathbreaking work in what he calls artifact theory. Dipert's leading idea is that artifacts must be understood in action-theoretic terms, i.e., in terms of the intentions of their creators.

A correct description of an artifact as an artifact describes the artifact in the way that its creator conceived of it—at least as much as is now possible. (Dipert, 1993: 15-16)

Dipert understands these intentions of the creator as the outcome of a “deliberative history” during which the creator contemplates the overall function of the artifact as her end and the possible means for achieving this end, and then forms a complex set of intentions in the form of a construction plan.

Of the many elements of the deliberative history, perhaps most important is the means-ends hierarchy, or plan, according to which the artifactual features were imposed on the object. (Dipert, 1993: 54)

My intentions, in making a screwdriver, to have another person come to believe it is a tool for turning screws, or actually come to use it easily for turning screws, are high-level intentions. My intention that the handle be out of stable plastic and that it have a certain shape, that the blade be metal, and so on, are middle-level intentions. They are conceived as means to my high-level intentions as ends. Finally, my intentions that I must use a lathe to create this shape, a drill press to place the hole in the handle for the blade, and so on, are low-level intentions. (Dipert, 1993: 151)

⁷Jon Elster calls this valorization of individual creativity ‘ethical individualism’ and says that Marx maintained a lifelong commitment to it. See Elster, 1985, especially sections 1.1 and 2.2.7.

Like Aristotle, Dipert thinks that this plan, which consists of fully conscious intentions, is supported in execution by a network of not fully conscious “half-intentional” habits. For example, my choice of wood as opposed to plastic for the handle of the screwdriver would be a fully conscious intention, but my choice of a particular piece of wood and my manipulation of the lathe in shaping it would be half-intentional habits. But the important point is that Dipert, like both Aristotle and Marx before him, gives us a theory of production in which a detailed design is first constructed in the mind, and then the artifact is constructed in reality in accordance with this design. The advances Dipert makes are, first, to spell out the genesis of this design in terms of a complex deliberative history having the overall structure of full-fledged means-end analysis; and second, to characterize the result of this deliberation in terms of hierarchies of intentions organized into a full-fledged plan to be executed in the actual construction of the artifact.⁸

Unlike Aristotle and Marx, Dipert is explicit about the individualistic tendency of his theory.

... [W]hat I call “artifacts” are necessarily conceived in terms of an individual agent, deliberative history, act of creation, and so on. This account of artifacts, and thus of art works, is highly biased toward “individualistic” accounts of agency, thought, and action. I think it is useful to have such an account of idealized individual agency, even if one ultimately ends up rejecting that there are many such examples in actual human behavior. I happen to think the idea is also necessary for making sense of many of our institutions and thoughts about ourselves and that the theory is, to some extent, of some people, and on some occasions, true. (Dipert, 1993: 194-5)

Dipert goes so far as to say that the deliberative history even of objects *known* to be collaboratively created should be reconstructed as if they were the work of a single individual. One motivation for this is to give an account of the unity of the artifact in terms of the unity of the producer’s intentions as directed towards an end.⁹ Thus for Dipert, as for Marx, the paradigm of the independent individual artisan functions as a sort of ideal against which actual production conditions can be measured or their products understood.

⁸Dipert (personal communication) objects to my characterizing this as unintelligent execution on the grounds that his theory of performance, which is an integral part of his overall theory of artifacts, clearly calls for intelligent execution. If performances are artifacts, the performer is the constructor of the work, but usually not its designer. So here design phase and construction phase are strictly separate. Dipert’s theory is that the performer is an agent whose intention is to carry out the intentions of the author of the work. But typically the author’s intentions are not fully known, and/or there are questions about how to fulfill them. So in practice the performer is forced into a creative role, fleshing out the author’s known intentions with some of her own. Dipert is right—this is not unintelligent execution. However, the *ideal* here is still unintelligent execution. If the author’s intentions *were* completely known the performer *should* be simply an unintelligent executor on Dipert’s view. This follows directly from his claim that a performer’s intention is to execute another agent’s intentions. Creative or intelligent execution is a failure in practice to achieve this, not in principle a positive characteristic of performance.

⁹See Dipert (1986: 406 and 1993: 32-7) on the notion of a “virtual” individual agent.

3 Centralized Control

I was central

I had control

R.E.M., “Country Feedback,” *Out of Time*

The standard account of production distilled from the previous sections can be characterized in terms of a model I will call ‘centralized control.’ The common sense of ‘control’ is that of a directive or determining force. But etymologically, ‘control’ is an accounting term from the French ‘contre-roller,’ referring to the practice of producing a copy of a ledger for purposes of account verification. The etymological and common senses are unified by the idea that in *faithful* copying the features of the original direct and determine all the relevant features of the copy. This is altogether similar to the way the standard account of production understands the relationship between mental design and artifact. The design specifies all the relevant features of the artifact, and in addition specifies the construction plan by which they can be realized. The actual construction of the artifact then is (ideally) a process which simply executes the plan and thus copies the artifact into reality, as it were. Or, to put the point another way, the standard account regards the mental design as controlling production in roughly the same way a program controls the operations and output of a computer.

Centralization is the idea that the controlling design is (typically) or should be (ideally) located in the mind of a single individual. This idea was implicit already in Aristotle, and has become progressively more explicit in recent accounts of production. It embodies the common, but often unreflective, assumption that the paradigm case of production is the case of a single individual who first designs and then constructs an artifact. The other side of this assumption is the idea that collaboration in production is at best a secondary phenomenon, and at worst an outright pathological one.

4 The Vicissitudes of Control

Control is problematic because the “faithful copy” ideal requires that all relevant features of the artifact be prefigured in the design, and that the construction process faithfully realize these features in the finished product. In other words, the construction plan ideally should be an *algorithm* (effective procedure) for realizing the prespecified features. I will use the example of cooking with recipes to show that this ideal model is not descriptively adequate to the phenomena typically observed in the production of everyday artifacts. I will then argue that rather than merely showing that the “faithful copy” ideal is rarely attained in practice, the descriptive inadequacy shows that control is *not* the purpose of design in the first place.

Food artifacts are produced daily by ordinary people, very often by following recipes which may be conveniently regarded as externalized versions of the originating cook’s men-

tal design.¹⁰ At first glance, recipes do seem to conform to the control model, since they normally include some sort of description of the dish and instructions for combining the listed ingredients in order to achieve the result described. But a closer look reveals an important divergence from the expectations of the control model—recipes routinely leave many details open.

Some of these open details involve features of the product. For example, a recipe may suggest either sour cream or yoghurt as a thinner for cucumber soup, or a cake recipe may suggest a number of possible frostings. Many recipes also list ingredients as optional altogether—chopped nuts in cookie recipes, for example. And recipes frequently specify some ingredients generically. A recipe might call for a cup of shortening, for instance, and then you have to decide between butter, margarine, vegetable oil, and so on.

Other open details involve the construction instructions. Consider the following cookie recipe.

Rolled Pecan Cookies

7 oz. butter (scant cup) 2 cups flour
 4 tablespoons powdered sugar 1 tablespoon ice water
 2 cups pecans (small pieces) 1/8 teaspoon salt
 1 teaspoon vanilla

Cream butter and sugar, add the rest. Roll with palms of hands into finger lengths. Bake 45 minutes (325F.). Roll in powdered sugar while warm, or shake in bag with 1/2 cup powdered sugar.¹¹

Notice, first, that the order in which to mix ingredients is not completely specified—you are on your own after creaming the butter and sugar. Even where recipes do specify this order more completely, it is largely conventional. For instance, it really does not matter whether you first sift the dry ingredients, then mix the wet ingredients, or *vice versa*. The only essential thing is that the dry and the wet ingredients be mixed separately before combining them. The instructions thus constitute at best a partial order on the steps in the construction process. Second, some crucial steps are not specified. Even a novice cook will know to put these cookies on a baking sheet before baking them, but the recipe does not tell you to do this. And more experienced cooks will realize there is an open question about whether to grease the sheet or not, and will probably realize this is unnecessary because these cookies have so much butter in them they could not stick to anything if they tried. Finally, the last instruction explicitly requires the cook to decide between two options for coating the cookies with powdered sugar.

¹⁰Recipes are typical of a large class of such externalized intentions to produce, such as sewing patterns, blueprints, instructions in “how to” books, outlines, and so on.

¹¹*The Settlement Cookbook*, 28th edition, revised and enlarged, Milwaukee: The Settlement Cookbook Company, 1947

So recipes diverge from the control ideal by not specifying some important artifact features and construction steps. In some cases the cook is expected to supply automatically the requisite features or steps on the basis of habits she has acquired (e.g., putting the cookies on a baking sheet). Here we see the importance of Aristotle's insistence on the necessity of having the right habits in addition to doing the right thinking. In other, more interesting cases the cook is implicitly expected to make a decision based on background knowledge she has (e.g., greasing the baking sheet). This is important because it shows that some of the thinking necessary for producing the artifact takes place in the construction phase—i.e., it is not all carried out beforehand in the design phase and then simply executed in the construction phase. Most interestingly, recipes explicitly prompt cooks to make decisions which will affect what features the artifact will have, or how those features will be achieved (e.g., optional ingredients, or alternative construction methods). Here some of the thinking and decision making is explicitly relegated to the construction phase.

This brings us to the question of how people actually do use recipes. The control model expectation is that cooks faithfully *follow* recipes to the maximum extent possible. But what more typically happens is that cooks use recipes as a basis for *improvisation*.¹² Improvisation is normally a response to local conditions. Sometimes these are difficulties encountered in the construction process. For example, when you do not have, or cannot get locally, an ingredient called for by the recipe, you can often substitute something else—cocoa and butter for baking chocolate, for instance. On the other hand, sometimes these conditions involve resources available locally which you can exploit. A cook with a walnut tree in his backyard might substitute walnuts for pecans in the recipe above, for instance. A third type of condition involves the special needs or desires of the cook and/or her clientele. For example, in the cookie recipe above a cholesterol conscious cook might use margarine instead of butter.

Cooks sometimes arrive at a stable customization of a recipe after a period of trying a range of variations. For example, when making the cookies above, because my oven runs a little hot I tried out a number of different baking times and oven temperatures, and finally settled on baking at the specified temperature for forty minutes, turning the baking sheet at the twenty minute mark for more even browning. On the other hand, recipes represent a permanent possibility of doing something different from what is specified, depending on who is coming to dinner, how calorie or cholesterol conscious you are feeling, what you have on hand, and so on. The important point here is that in addition to all the details left open by the recipe itself, the regular practice of cooks is to change around even the details that *are* specified to suit their own situations and purposes.

¹²My thinking on this issue owes a lot to Agre and Chapman (1990) and Suchman (1987). I also want to stress here that I am talking about *experienced* cooks. Novice cooks do tend to follow recipes rather slavishly. But novice cooks by definition do not yet know how to cook—in particular, they do not yet have the fund of habits and background knowledge the experienced cook has accumulated. Consequently, there is no particular reason to think that experienced cooks are merely doing better what novice cooks do, and some good reasons to think they are doing something *different*, namely, using recipes as resources for improvisation rather than as controlling devices. These reasons are detailed in what follows. See Dreyfus et al. (1986) for more on the dangers of trying to understand skilled behaviors in terms of what novices do.

So recipes diverge from the control ideal in two ways. They explicitly or implicitly require the construction phase cook to do some of the design work, and in addition they are routinely used as a basis for further improvisational alterations to, or extensions of, the design rather than faithfully followed. This means the constructor is typically *not* an unintelligent executor. Thus, contrary to the picture offered by the control ideal, the intentions of the constructor are crucial for understanding the production of an artifact; first, because they are required for filling in gaps in the express intentions of the designer, and second, because in practice they are often improvisationally substituted for the express intentions of the designer.¹³

But what accounts for the descriptive inadequacy of the control ideal? The main culprit is an underlying assumption that the world is stable over time and homogenous across agents. Consequently, the designer is viewed as approaching omniscience and the constructor omnipotence. The control ideal of advance specification and unintelligent execution really makes sense only in light of this assumption, which was close to explicit in the early artificial intelligence planning literature, for instance.

The agent is given a goal, it computes a plan for achieving it, and then, at least in principle, it executes that plan. The environment is quiescent; the agent is the only force acting on it. So nothing of significance happens while the agent is forming its plan. And nothing happens while the agent is executing that plan, except what the agent itself causes to happen. (Pollack, 1992: 45)

But our actual world is neither quiescent nor homogenous. So even knowledgeable, experienced agents are not able to predict future or local conditions, either for themselves or for other agents, with any great reliability. In such an environment it does not make sense to work up elaborate advance plans to be faithfully executed. First, if you are wrong in your predictions the plan cannot be executed and you are back to square one. So elaborate designs militate against flexible accommodation of problematic contingencies. Second, if you are unable to foresee local conditions, you cannot plan to take advantage of opportunities which might serve to advance your interests in ways different from those initially envisaged. So such designs militate against serendipity as well. In short, the control ideal aims to minimize the exercise of creative improvisation, but our environment is one in which we are actually very well served by improvising.¹⁴

¹³You might object that this is only true of cases where the constructor and the designer are different individuals. But take the process of producing a paper by first making an outline. The outline does not contain all the relevant details of the final paper. So a topic heading might be ‘Aristotle’s theory of artifact production,’ with nothing more about the content of that theory specified. Indeed the writer may at this point not yet have worked out his interpretation of Aristotle. Similarly, the outline is only a partial order on the steps of the construction process. In principle, you can write the sections you have specified in any order. And finally, in practice outlines are no more sacrosanct than recipes. By the time you are half way through the paper you have ordinarily (in my experience, anyway) deviated from the outline at any number of points, changing the topics to be covered or rearranging the order, as the actual writing progresses. So even where designer and constructor are the same individual, we have a design which requires further design work to be done during the construction phase, and which is used as a framework for improvisation rather than a controlling device.

¹⁴Vinod Goel (1995) makes a similar point. In producing designs which are themselves artifacts (e.g.,

So the control ideal is not just descriptively inadequate—it is a false ideal. Since we are agents with finite knowledge in a world which is neither stable nor homogenous, what we should ideally aim to be is good improvisors, not good controllers. So rather than concentrating on our capacity to come up with elaborate designs and plans, we should concentrate on our capacities to improvise solutions to unforeseen problems and exploit unlooked for opportunities. And rather than understanding improvisation as merely a patch for incomplete or unexecutable plans, we should understand designs and plans as resources for activity which is inherently improvisatory. Moreover, we should be aware that designs are only one such resource—some others are seeking advice or information, trial-and-error exploration of possibilities, and use of tools and other artifacts in ways they were not designed to be used. Finally, taking improvisation as the ideal rather than control requires revision of the traditional understanding of the relation between design and construction. Construction is not unintelligent execution, not even ideally. It is an inherently creative phase of production, and must be examined on its own terms by any adequate theory of production.

4.1 A Deflationary Objection Deflated

Devotees of the control model may respond: ‘We agree with you about the phenomena, but improvisation does not represent an alternative model. Improvisation is just scruffy planning. Precisely *because* planners are not omniscient and executors not omnipotent, plans must be revised frequently due to unforeseen contingencies arising in the execution phase. And the original goal sometimes must be revised in light of these contingencies as well. So longterm planning and execution is broken up in practice into small increments of replanning in light of the current situation. And often enough an agent simply proceeds from the start in small increments, as when you cook dinner by repeatedly looking to see what is in the refrigerator and figuring out what you can do with each item as you go along. But each small increment retains the basic control structure—visualizing a way the world might be and what you would have to do to make it that way, and then realizing what you have visualized. So improvisation is just reality adjusted control.’¹⁵

This objection is likely to prove tenacious, because it appeals to the inherent future directedness of action, which rests on the agent’s ability to visualize the world as other than it is at the moment. The very etymology of the term ‘improvisation’ suggests that here there is no future directed visualization, and if that were true then improvisation *would* (blueprints) there is a prolonged sketching phase in which alternative possibilities are generated, explored, and refined. Cognitive virtue here resides in producing sketches which are coarse-grained, ambiguous, subject to multiple interpretation, and so on, because these characteristics keep the design open-ended and easily transformable, which is what you want at this stage. His point—that cognitive virtue does not always reside in precision, elaborate detail, univocal designation, and so on, and in particular that it does not do so in design contexts—dovetails nicely with my claim that improvisational activity is more virtuous than planning in production contexts.

¹⁵This objection might well be voiced by Michael Bratman (1987, 1990, 1999), for instance, who, to his credit, has repeatedly emphasized the necessity for plans to take account of the cognitive limitations of agents, and the resulting need for partial plans and frequent replanning.

be a different model of action, albeit a wildly implausible one. But of course it is not true. Improvisation always has a direction and a destination, however vaguely delineated and indefinitely revisable these may be. It is continuously responsive to changing local conditions, but this responsiveness is intelligent and even creative, not blindly reactive. To exploit an unforeseen opportunity, for instance, you first have to see it *as* an opportunity, and this comes down to visualizing *how* to turn it to account. But this visualization, which is clearly central to planning as well, seems to bring improvisation under the aegis of the control model.

Agre and Chapman distinguish improvisation from planning in the first instance by pointing out that in improvised activity “each moment’s action results. . . from a fresh reasoning-through of that moment’s situation” (1990: 21). In consequence, they say, the improvising agent does not control its world, but rather interacts with, or participates in it.¹⁶ This is a good direction, since it at least supplies an alternative metaphor. But it will not defuse the deflationary objection since it does not address the concern that the underlying structure of the activity is the same in both cases, and that the only difference is one of degree in the size of the activity units so structured. To do that we first have to make a distinction between *mechanisms* and *strategies*.¹⁷ A mechanism is an ability to perform some relatively basic (i.e., domain independent) and unitary function. Examples of mechanisms are the ability to visualize the world as otherwise than it is; to tell when an operation has succeeded or when it has failed; to recognize cause and effect relationships; to recognize relevance relationships, and so on. A strategy is a higher level, relatively domain dependent, pattern of activity which selects among available mechanisms and organizes their use relative to each other in a distinctive way in order to carry out the life tasks of the agent. Examples of strategies are trial-and-error learning; learning by apprenticeship; linguistic communication; non-linguistic communication (gesture and body “language”), and so on. The difference between mechanisms and strategies is not just a matter of scale. Although a mechanism may be dependent on the possession of more basic mechanisms, it is applied by the agent as an inflexible unit—you either have the whole ability or you do not. A strategy depends on a repertoire of mechanisms and sub-strategies; but the agent applies this repertoire flexibly, so the pattern of activity characteristic of the strategy varies from domain to domain, agent to agent, and occasion to occasion.

Control and improvisation are different *strategies*. This means that although they use many of the same mechanisms, they embed this use in quite different overall patterns of activity. The deflationary objection quite properly recognizes that both control and improvisation use the mechanism of visualization, but it fails to acknowledge the differences in the overall activity pattern. Most importantly in this respect, for the control model visualization is *the* central mechanism and the use of other mechanisms and substrategies is subordinated to it. For the improvisation model, visualization is not central, but is interleaved with a

¹⁶Agre and Chapman (1990), pp. 20-24, *et passim*.

¹⁷Agre and Chapman make a similar distinction between “cognitive machinery” and “*dynamics* or regularly occurring patterns of activity” (1990: 21). But since they do not articulate the deflationary objection in the first place, they do not apply this distinction to defuse it.

number of other, equally preferred, mechanisms and sub-strategies, such as being constantly on the *qui vive* for unforeseen opportunities (which relies on relevance detection after the fact); and trying out a number of options to see if any of them work, and if so, which one works best (which relies on trial and error experimentation rather than visualization). To put the point another way, the characteristic pattern of the control model is to maximize the use of visualization to construct plans which are as long-range and detailed as possible, and to resort to other mechanisms and sub-strategies only when faced with internal limitations (lack of omniscience) or external difficulties (lack of omnipotence in execution). The characteristic pattern of the improvisation strategy is to minimize the use of visualization in order to maximize the use of other mechanisms as a matter of preference, not as a matter of need. So improvisation is not just reality adjusted control—it is a different way of managing the resources of the agent and the affordances of the world in order to get things done.

5 The Vicissitudes of Centralization

Centralization assumes individual production is the paradigm case, but in fact artifacts are often designed and/or constructed collaboratively. Although cooking is a good example—from homes to restaurants, it is quite typically collaborative—accounts of this activity are not readily available. So here I will adopt songwriting as the example domain, since accounts of this are available in interviews with songwriters.¹⁸

One obvious way production is collaborative involves collaboration among constructors. This is often necessary because it is physically impossible, or impossibly inefficient, for one individual to construct the artifact. In general, any artifact which is large (e.g., buildings) or complex (e.g., computers) or involves simultaneous performance (e.g., music) calls for collaborative construction. I will not dwell on this aspect of collaboration because it is the least disputed. Centralization is primarily the thesis that the *design* phase of production is paradigmatically individualistic.

But design is often collaborative as well. One way this occurs is when the designer and the constructor are different individuals. As we have discovered, the constructor is typically an intelligent executor with ongoing responsibility for decisions affecting the final form of the artifact. So there is a full-fledged collaborative relationship between constructor and original designer in the sense that understanding the contributions of both is necessary for understanding how and why the artifact got to be the way it is. This holds even if the designer and constructor do not work face-to-face, or communicate with each other, or even know each other personally. In music this kind of collaboration is common because the performer is often not the original composer. The contribution of the performer, roughly speaking, is the interpretation of the piece; but interpretation covers a wide range of phenomena. At one end you have traditions like the recent Western classical one where the performer is expected to

¹⁸I have relied here on both published interviews, and on interviews with some local songwriters I conducted myself. When not otherwise attributed, the quotations from songwriters in this section come from this latter set of interviews.

render a note-for-note version of an elaborate score which may constrain innovation even with regard to dynamics and tempo. In the middle you have a vast and varied range of traditions where written scores are less elaborate or non-existent, and the performer is expected to fill in many of the details, e.g., the earlier Western classical tradition, where the performer was responsible for ornamentation, cadenzas, and so on; and the jazz and Asian gamelan traditions, where the performer is expected to improvise.

At the other end of this range are cases where the line between performing an existing song and writing a new one begins to blur. A simple example is writing new lyrics to an existing tune—the Christmas carol “What Child Is This,” for instance, which is sung to the tune of “Greensleeves.” More interesting cases involve the interpretive appropriation of a song in such a way that it is substantially redesigned. Led Zeppelin did this brilliantly with Memphis Minnie’s “When the Levee Breaks,” for instance, by changing the tempo and the key, abandoning most of the original verses, rewriting some lines of the verses they kept, changing the melody substantially, and writing a bridge verse with a different melody and lyrics. Signalling both the continuity and the collaboration, they credited the song to Memphis Minnie and all four members of the band. A more recent example is R.E.M.’s “Hope.” Although the lyrics are completely new and the melody substantially different, this song remained sufficiently close to its origin in Leonard Cohen’s “Suzanne,” that it is credited to Cohen along with the members of the band.¹⁹ In these last cases the line between performing and creating is virtually erased. I stress this phenomenon of creative appropriation, because I think it represents one of the major strategies by which new artifacts are created. Indeed, Pete Seeger defines folk music in terms of it.

[Folk music is] a process by which ordinary people take over old songs and make them their own. They don’t just listen to it. They sing it. . . . And they change it. (Zollo, 1997: 4-5)

Similarly, cooks create new dishes by creatively appropriating existing recipes, and dress-makers and carpenters create new apparel and furniture by creatively appropriating existing patterns and designs. This strategy is more heavily relied on in some artifact traditions than in others, and more frequently used by some individuals than by others—there are half a dozen examples from Led Zeppelin, whereas “Hope” is unique in R.E.M.’s *œuvre*, for instance. But its importance for present purposes is that it is inherently collaborative in the sense defined above—the intentions and decisions of several individuals must be taken into account in understanding how and why the artifact got to be the way it is.

On the other hand, design is often straightforwardly collaborative *ab initio*. This is understood with regard to sophisticated, mass-produced artifacts like musical instruments or cars, where companies employ teams of designers to devise new models. But it is also true for many things we think of as typically individually designed, like songs, which are

¹⁹Memphis Minnie’s original version of “When the Levee Breaks” can be found on *Kansas Joe and Memphis Minnie, 1929-1934* (Document, 1991). Led Zeppelin’s version is on their untitled fourth album (Atlantic, 1971). “Suzanne” is on Cohen’s *The Songs of Leonard Cohen* (Columbia, 1968). “Hope” is on R.E.M.’s *Up* (Warner Brothers, 1998).

often written by teams of various sorts. R.E.M. routinely credits all songs to all the band members, for instance, and when asked whether the band's actual practice is to write songs together, bassist Mike Mills replied in the affirmative.

We put them [songs] together that way. Everybody sits at home and diddles around. Sometimes you'll come up with little ideas and sometimes you'll come up with a huge part of a song. And then you'll take that in to everyone else and piece it together until you get a song. Other times, things just come out of, literally, just the four of us sitting around making noise. All of a sudden it will reemerge into a song. (Zollo, 1997: 631)

Here we also glimpse the structure of R.E.M.'s collaborative process. First, there are two different strategies for generating initial song elements, which Mills termed the "show each other" method and the "chaos" method in a later interview. The "show each other" method, which is the predominant one, involves band members bringing in song elements they have come up with on their own. The "chaos" method is a strategy the band resorts to when the "show each other" method fails because no one has anything to show. As Mills describes it:

By "noise"...mostly what I meant was the three of us instrumentalists would just start playing something without listening to the other.... But if any one of us were to hear something the other one did that we felt excited by, then we might begin to follow them in some way.... And...once the drummer starts following somebody you have to follow that. The chaos is out the window and now you're in a little bit more of a formed thing.

Here we also learn that the initial face-to-face session usually involves only the band's instrumentalists. Once one of the abovementioned methods has generated a song element which interests all of them, they work on it together, revising and extending it. The result is then relayed to singer Michael Stipe, either by playing it for him in person or giving him a tape. If it interests him he supplies a melody and lyrics. Sometimes this is a matter of continuing the revise-and-extend process, using the instrumental material as the direct basis and inspiration for the melodic material. But Stipe also has a fund of melodies and lyrics he has come up with on his own, which he can often fit to the instrumental material provided by the other band members. The nascent song then undergoes further revision and refinement by the band in face-to-face sessions before (and sometimes during) recording. Variations on this basic pattern can and do occur—Michael Stipe sometimes participates in the initial session with the instrumentalists, for instance; or one of the instrumentalists may write the lyrics. But in any case, the overall structure of the process is cyclical. The "show each other" method, where individual band members bring in contributions they have generated on their own, is reinvoked at several levels; and it is interleaved with face-to-face methods where some or all of the band members work together to generate initial song elements (the "chaos" method) or to develop song elements they have in hand (revise-and-extend).²⁰

²⁰I have pieced together this picture of R.E.M.'s songwriting practice from several sources, including an interview with Mike Mills. See especially Zollo (1997: 629-41) and Gray (1997: 59-96).

Two points emerge from this brief sketch of an actual collaborative design process. First, the relationship between collaboration and improvisation is not entirely contingent. Improvisation is the best—and sometimes the only workable—strategy where events and conditions are not predictable. But in collaborative design, the contributions of other collaborators cannot be reliably anticipated in any detail. As R.E.M.’s songwriting process illustrates, the pattern is to wait for someone else to present something, and then respond to it; or else present something yourself and wait for responses from the others. And the relationship is not entirely contingent in the other direction either. Since the improvisation strategy, unlike the control strategy, relies on continuous interaction with the environment, and since other people are a salient aspect of the environment, improvisational activity naturally tends to generate collaborative relationships. In this sense collaborative interactions are like conversations. You may plan to *have* a conversation—perhaps even a conversation on a specific topic—but you cannot plan the conversation itself. And it is pretty difficult to have a real conversation with yourself—the difference between a monologue and a dialogue is a *structural* difference, not simply a numerical one.

The second point is that the contributions of individual collaborators do not combine in an additive or linear fashion. Rather than fitting together like the parts of a jigsaw puzzle, they tend to revise and override each other in complex ways. Members of R.E.M. mention being surprised regularly by the changing form and fate of their own contributions, as well as by the ultimate character of the song as a whole. For instance, an initiating contribution by Mike Mills, which he described as a “quiet little ditty” played on acoustic guitar, ended up as a rousing chorus in a song described by Michael Stipe as “stomp rock” (Zollo, 1997: 632). This shows how little amenable collaborative activity often is to any attempt to bring it under the centralization thesis by means of an analysis which would centralize control in one of the collaborators. The individual members of R.E.M. do not seem much inclined to insist on either the integrity of their own contributions, or on any privileged status for whatever ideas and intentions they may originally have had about how to develop these contributions. In short, the process is markedly decentralized and unhierarchical from the start.

Given the ubiquity and complexity of collaborative phenomena in production, then, what accounts for the persistent individualist slant of traditional theories? First, there is often a commitment to individualism for methodological reasons. Jon Elster has argued that Marx was committed, at least intermittently, to explaining social action reductively in terms of the beliefs and actions of the individuals making up the social group.²¹ And Dipert (personal communication) has indicated that he sees group or collaborative production as “phenomenologically, even metaphysically, parasitic on the individualistic case.” There are almost as many methodological individualisms as there are methodological individualists, but the common denominator is the view that individual activity is epistemologically fundamental in the sense that understanding small scale collaborative activity, as well as large scale social institutions and movements, depends upon first understanding individual activity, and not *vice versa*.²²

²¹Elster, 1985: 5-8.

²²The individual may be regarded as ontologically fundamental as well, but methodological individualism

This view has been repeatedly challenged. Even Marx, in other moods, was disposed to challenge it. As previously described, he thought of the individual as always already social, a view which bids fair to reverse the epistemological priorities. As Hans Joas points out, the predominant trend in sociology, American pragmatism, and most strains of Continental philosophy has been to consolidate this insight into the ‘primary sociality’ of the individual, by working out detailed theories of socialization (e.g., Mead, Scheler), and of the ongoing role of social institutions and rituals in the maintenance, dissolution, and restabilization of ego boundaries (e.g., Nietzsche, Durkheim).²³ If there is a bastion of individualism anywhere, it is in analytic philosophy. But even there challenges to methodological individualism have been raised in the philosophy of social science. For example, Harold Kincaid (1986) has carefully sorted out the different versions of methodological individualism and argued that only the very weakest of them has any real plausibility. Kincaid’s argument is complex, but he comes back repeatedly to the point that it is virtually impossible to talk about individuals in the first instance without employing concepts involving social roles or institutions. So here too there are echoes of Marx’s view that the individual is always already social.

As previously noted, sociality and collaboration are distinguishable phenomena. Sociality is a matter of institutions and practices in the abstract, whereas collaboration involves concrete interactions with others in the context of specific tasks. For example, sociality in music means being conversant (in practice, not necessarily in theory) with the standard scales, song forms, performance practices, and so on, of a musical tradition. Collaboration means writing a song with someone else, performing a song someone else has written, performing as a member of an ensemble, and so on. Although it would bolster the case against methodological individualism, critiques hinging on the primary sociality of the individual fail to note the centrality of concrete collaborative interactions in both the socialization process and the ongoing activity of the socialized individual.

Socialization is not just an initiation into social institutions and practices, but a developmental process during which skills are learned and abilities acquired. This is particularly obvious in the case of artifacts, since the ability to produce them usually requires training of a fairly explicit sort. This training often takes an apprenticeship form, which standardly involves the collaborative production of artifacts. For example, people usually learn how to cook from an experienced cook who shows them how to do subparts of the cooking task until they first become fully competent collaborators and then eventually start cooking on their own. Although it is easy to imagine someone learning how to write songs by apprenticing with a more experienced colleague, songwriters in fact report other types of collaborative interaction as being developmentally important. For example, they commonly report writing their first songs by using other people’s songs as “blueprints,” as Elvis Costello puts it.

I was using yesterday’s records as blueprints. . . . I wanted to take some of the ready-made clichés that Goffin and King or Smokey Robinson would come up with and come up with my own photo-negative versions of them. Almost every

focuses on the epistemological question.

²³Joas, 1996: 184-195.

song on my first album was an opposite—a diseased version—of another kind of song. (Flanagan, 1986: 196-7. See Gray (1997: 60) for a similar assessment of R.E.M.'s earliest, and mostly never recorded, songs.)

Although on the face of it the individual is working and learning on their own here, the process is nevertheless collaborative to the extent that it involves the creative appropriation phenomenon, which I described above as a distinct and important species of collaboration.

In addition, songwriters frequently report writing their first songs in direct collaboration with other, equally novice, songwriters, as R.E.M. guitarist Peter Buck did when asked about his first songs.

Michael and I wrote a few before the band started. . . . We'd write in tandem: me and Mike or me and Michael or Bill and Mike or Bill and Michael. And show it to the other guys. And after two or three months we started writing all together. All of us can write songs on our own. But having the four of us all do it has really made the difference. (Zollo, 1997: 638)

Jason Slatton, singer and guitarist for The Lures, reports an early reliance on the creative appropriation strategy, but describes later direct collaborations as the turning point in the development of his ability to write songs on his own.

They [the first songs] were on my own, but I didn't feel as strongly about them, and it was only until I had the confidence to sit with somebody else and show them ideas that it gave me the confidence to then go back on my own and work on it. When I was young I was more or less just trying to ape people that I listened to and liked, and now I'm . . . trying to find my voice. . . . And so writing with Randall and Russ gave me the confidence to work on that.

These examples show that collaborative production is integral to the developmental process rather than being something that occurs only after individuals have acquired the skills to produce on their own. So from a developmental point of view, collaborative production is an epistemologically fundamental phenomenon in the sense that understanding how individuals come to be able to produce artifacts on their own usually depends on understanding how they produce artifacts in collaboration arrangements of various sorts, rather than *vice versa*.

Collaboration is epistemologically prior with regard to ongoing productive activity as well. Verifying and expanding a hypothesis first advanced by Wolfgang Köhler, Peter Reynolds (1993) analyzed cross-cultural and cross-species videotapes to show that it is the social organization of tool use and tool construction which is the most striking difference between humans and apes. Apes invariably construct tools individually whereas humans usually construct them collaboratively in what Reynolds calls the 'face-to-face task group.' Moreover, humans typically use this collaborative mode of construction even in cases where the artifact could easily be constructed by an individual.²⁴

²⁴Reynolds uses an videotaped example involving simple stone tools to support this claim, but it is also supported by Peter Buck's remark quoted above that all the members of R.E.M. can write songs on their own but prefer to do it collaboratively.

The essence of human technical activity is the anticipation of the action of the other person and performance of an action complementary to it. . . . I call this process *heterotechnic* cooperation (‘different crafts’) to emphasize the complementarity of social roles. Heterotechnic cooperation may be contrasted with *symmetric* cooperation, in which all the participants do the same thing at the same time. . . . as when both men stir the spinifex together or a pack of wolves runs down a caribou. Thus human technology is not just ‘tool use,’ and not just ‘cooperative’ tool use, but tool use combined with a social organization for heterotechnic cooperation. (Reynolds, 1993: 412)

This social organization involves task specialization, assembly of separately constructed components, and symbolically mediated communication in addition to complementarity of social roles. Accordingly, Reynolds hypothesizes that the other most frequently cited difference between human and non-human animal tool construction—that human artifacts are typically composed of several separate components fastened together rather than being all of a piece like the chimpanzee’s termite dipping stick—is a co-evolutionary development, emerging in synergy with the social organization of the face-to-face task group. This organization has rudimentary analogues in such reciprocal activities as grooming; but, significantly, these reciprocal behaviors are not applied in technical activity contexts among non-human animals. So collaboration is not merely a central phenomenon in ongoing tool use and tool construction, but a distinctively human one. This suggests that phylogenetically, collaboration is epistemologically more fundamental than individual technical activity.

Methodological individualism thus faces serious challenges on multiple fronts. But it is frequently accompanied by another assumption—that individual production is fundamental in the sense that it is the *ideal* form of productive activity. As noted above, both Marx and Dipert hold such a view, although for different reasons.²⁵

On Marx’s view, an artifact is in the first instance an objectification or expression of the individual artificer’s needs, desires, hopes, beliefs, and so on. Thus it builds in an idealization of individual production, since collaboration would interfere with that expression, as previously noted.²⁶ But this expressive theory of the artifact is conjoined in Marx with an unwavering commitment to the free development of the individual through autonomously chosen creative activity. Famously, Marx thought this could be achieved only by doing away with the socially regulated division of labor.

[T]he division of labor offers us the first example for the fact that man’s own act becomes an alien power opposed to him and enslaving him instead of being controlled by him. . . . For as soon as labor is distributed, each person has a particular, exclusive area of activity which is imposed on him and from which he cannot escape. He is a hunter, a fisherman, a herdsman, or a critical critic, and he must remain so if he does not want to lose his means of livelihood. In communist

²⁵Dipert’s position has considerable interest, but for reasons of space I will omit discussion of it here.

²⁶It should also be noted that the expressive theory sorts well with the control ideal, since to express something is often understood as copying into an external form something already internally articulated.

society, however, where nobody has an exclusive area of activity and each can train himself in any branch he wishes, society regulates the general production, making it possible for me to do one thing today and another tomorrow, to hunt in the morning, fish in the afternoon, breed cattle in the evening, criticize after dinner, just as I like, without ever becoming a hunter, a fisherman, a herdsman, or a critic. (Marx, 1932/1967: 424-25)

But consolidated divisions of labor appear at the level of small scale collaborative production too. This is clear in R.E.M.'s songwriting process, for example, where singer Michael Stipe routinely takes responsibility for the lyrics and melody, and the instrumentalists take responsibility for the riffs and chord changes and so on. In addition to this overall division of labor, each of the instrumentalists assumes some special responsibility for writing the part for the instrument he usually plays in performance. And indeed, this is what we should expect if Peter Reynolds is right that what is distinctive about human collaborative activity is the complementarity of social roles and its concomitant task specialization. At this scale the division of labor may be somewhat more fluid in terms of the latitude an individual may have to assume different roles on different occasions or in different collaborative ensembles. But Marx's point is quite general. Any socially organized distribution of tasks, even if it is small scale, local, or transient, compromises the freedom of the individual because she becomes identified with a given role, and cannot arbitrarily switch from one role to another just as she likes. In short, there is a power outside us directing our activity in any collaborative situation, and it is not the opposing power of another individual, but the institutional power inherent in the distinctive social structure of collaborative interaction. Thus the idealization of individual production is grounded from another, complementary direction, since any collaborative arrangement compromises not only the free expression of my individuality but my freedom of activity as well.

For Marx, then, centralization as an ideal makes sense on the assumption that as producers our fundamental desires are for untrammelled self expression and absolute autonomy in action. So just as the control ideal rested on assumptions about the epistemic structure of the agent, the centralization ideal, in Marx's case, rests on assumptions about motivational structure. But if his assumptions are ungrounded, the centralization ideal might be a false ideal too—an ideal which does not make sense in light of the actual motivational structure of the agent. One consideration which already points in that direction is the ubiquity of collaboration in contexts where it is not necessary. Under Marx's assumptions, this is anomalous—individuals should be motivated to avoid collaboration whenever possible. But in songwriting, for instance, collaboration is very common,²⁷ and most of the songwriters involved can and have written songs on their own. Collaborative writing relationships are actively sought out and maintained, and some of them are of legendary longevity—Mick Jagger and Keith Richards of the Rolling Stones have been writing songs together for some

²⁷Obviously I do not have hard statistics here, but some evidence is provided by a recent National Public Radio series on the 100 most important American musical works of the 20th century. The selections are listed on their website (www.npr.org/programs/specials/vote/list100.html), and cover a wide range of genres. Roughly half the selections represent collaborations of some sort.

thirty years, for instance. Other, less monogamous, songwriters go from one collaborative relationship to another over the course of their careers, or engage in several collaborations simultaneously. For example, Russ Hallauer, guitarist for The Lures, is now in his third band, and third collaborative songwriting relationship; while his bandmate, Jason Slatton, writes regularly not only with Hallauer and the other Lures, but also with Randall Bramblett of The Randall Bramblett Band, and he has been known to write opportunistically with other people as well. So what motivates this widespread collaborative activity?

First, songwriters report that collaboration is psychologically rewarding. As the quote above from Jason Slatton indicates, this is not only a matter of being currently motivated by other people's interest in your musical ideas, but can result in permanent psychological gains, such as increased confidence in your own abilities. Second, and more importantly, collaborators are regarded as a resource for musical ideas. For example, they may have some expertise you do not possess, like a facility with lyrics, or a deeper knowledge of musical theory. Or it may be simply that their different musical experience and preferences enable them to generate ideas you would never have come up with yourself, as David Crosby (of The Byrds and Crosby, Stills, and Nash fame) indicates.

I've been trying more and more to write with other people. . . . [I]f this is the width of your palette. . . , when you're working with someone else, it's about twice as wide. (Zollo, 1997: 377)

In addition, collaborators have an important evaluative function, as Russ Hallauer points out.

[I]f you bring it up a couple times and it doesn't take. . . that's your judgement of a song, really, is if you can't get your band members interested in it, you know, how are you going to get anyone else interested in it? So I think you kind of trust your band members to judge you that way.

And finally, collaborators often combine the evaluation function with the musical resource function by recognizing the value of a musical idea you have, in some literal sense, produced but have not recognized as interesting yourself, as in the following story related by Mike Mills.

[T]he song "Me in Honey" . . . , the bass riff that is the song—it's the entire song—I was just sitting there, I started playing that riff. . . and I wasn't even thinking about anything, I was just sitting there doing it because I was killing time between songs, and Michael goes 'Keep going!'. . . . So it's good to have him there, because sometimes he'll take things over that would never become songs otherwise.

Clearly these aspects of collaboration are regarded as advantageous by collaborators. But this is precisely what is anomalous under Marx's assumptions—why is it that having other people evaluating your musical ideas and incorporating their own ideas into what would otherwise be "your" song not regarded as compromising individual self expression and abrogating individual freedom of activity? Here it is important to note what songwriters say about their overriding incentives. Rolling Stones guitarist Keith Richards put it this way.

It's great to watch other people's methods of working and to work with them. . . . Everybody needs a way of thinking about doing whatever it takes. As long as at the end of the day you come out with the song or the record, it don't really matter. What matters is the end product. (Flanagan, 1986: 184)

And Mike Mills adds this.

Well, you want it to be good. . . . Bill and I already had songs that we'd written from back in Macon, Georgia, that he and I had been goofing around with. So we showed them [Michael Stipe and Peter Buck] these songs, and we really liked what they did with them. . . . We thought it was great.

Moreover, on Mills' view there really is no important distinction to be had between "his" songs and R.E.M.'s songs.

I don't like to divert myself or dilute myself by writing my own songs, because to me my songs are R.E.M. songs. It's the same thing.

These remarks show that autonomy and individual self expression are not overriding concerns for songwriters. Rather they are focused on the end product itself, and on its quality. But in that case, the aspects of collaboration enumerated above appear as unadulterated advantages, since they conduce to those ends. Moreover, the division of labor which so bothered Marx appears as a positive aspect of collaborative activity for the same reason. There may be frustrations attendant upon being identified with a specific social role, but these are not frustrations of the overriding concern. On the contrary, division of labor is often the best way, and sometimes the only way, to generate a high quality end product. So like control, centralization is a false ideal. It is unresponsive to the actual incentives which motivate people in production contexts. And given those incentives, being a good collaborator looks like a much more reasonable ideal to adopt.

Like control, centralization is descriptively inadequate—collaborative phenomena are ubiquitous in production contexts. The attempt to salvage the centralization thesis as a methodological priority founders on standard objections to methodological individualism which appeal to the primary sociality of the individual; but it is also subject to a parallel critique which appeals to the developmental and phylogenetic centrality of collaborative phenomena. Marx's attempt to salvage it as an ideal runs into difficulties as well, because it is unresponsive to the actual motivational structure of producers.

6 Conclusion: Applications to Action Theory

It should come as no surprise that contemporary action theory is committed to the centralized control model, since, like the traditional production theory we have been discussing, it traces its ancestry back to Aristotle. With regard to the control aspect, action theory has most recently committed itself wholeheartedly to the view that as agents, we are fundamentally planners. This view derives from a question about what an intention is. An earlier theory,

now rarely defended, held intentions to be belief-desire complexes. But merely entertaining beliefs does not explain how I go about structuring the course of my action in terms of them. This can be explained, however, if we regard intentions as having plans, rather than merely beliefs, as their representational component. This indigenous development in philosophical action theory has converged happily with artificial intelligence planning theory, which developed independently as the reigning paradigm in robotics.²⁸

The position I have taken in this paper is that activity is not planfully structured, but improvisationally structured. This has two implications for action theory. First, it suggests that recourse to planning is at least not the only way, and arguably not the right way, to respond to the question of how beliefs that we hold structure actions we undertake. This is important, because other than the belief-desire view which it supersedes, there has been to my knowledge no challenge to the planning view of agency in recent action theory. Second, it suggests that the representational component of intentions still has not been correctly characterized, so there is a revised theory of intention in the offing here as well. This is particularly important because of its connection to questions about the nature of rationality. Characterizing intentions as including a plan component makes both intentions and intentional action rational in a very straightforward and traditional sense which will not be available to the improvisation theorist.

Action theory has traditionally been implicitly committed to the centralization aspect of the model as well. The best evidence for this is the simple fact that only in the last ten years or so has there been any substantial discussion of social or group action.²⁹ Some of the leading participants in this discussion are explicitly committed to some version of methodological individualism.³⁰ My position in this paper clearly has implications for this issue. But more importantly, it has implications for several general tendencies characterizing this discussion now that it has begun.

First, because a framework for understanding individual action was already firmly in place before group action ever came under consideration, a natural tendency is to approach group action as a matter of extending the individual action framework to cover multiple agents. One specific result here is that group action tends to be conceived in terms of multi-agent planning.³¹ But one of the claims I have advanced here is that the relationship between collaboration and improvisation is not entirely contingent. If this is right, conceiving group action in terms of shared plans is exactly the wrong thing to do—what you really need is a theory of how agents improvise in concert. Second, the analysis of group action has revolved largely around the question of what it is for a group to have an intention. It is admitted on all hands that whatever this amounts to, it cannot be exactly the same as

²⁸Articulations of the planning view of intentions and its rationale are now ubiquitous. Its leading proponent has perhaps been Michael Bratman (1987, 1990 and 1999); but see also Brand (1986) and Mele (1992, and 1997: 16-20). See Cohen, et al. (1990), for the connection to AI.

²⁹Margaret Gilbert (1989) is widely regarded as having initiated this discussion. See also Bratman (1999), Searle (1990), and Tuomela (1995). It is worth noting that one important recent anthology contains no articles at all on this topic, which the editor, Alfred Mele, refers to explicitly as “less traditional” (1997: 26).

³⁰See Bratman (1992), and Tuomela (1995), especially chapter 9, for instance.

³¹See especially Bratman (1997), chapters 5-8; and the essays in Cohen, et al. (1990).

what it is for an individual to have an intention. Nevertheless, the starting point of the analysis is the individual case, and the tendency is to model the collective case as closely as possible on the individual one. For example, David Velleman (1997) insists that we have to have an account of how intentions can literally be shared, rather than simply converging. The intuition here, I think, is that if individuals get things done in virtue of having a single-minded intention, then the question of how groups get things done must have some analogous answer. But if the direction of the analysis were to be reversed in light of my suggestion that collaborative activity is in some important respects more fundamental than individual activity, this intuition would lose its force. Specifically, I suspect that although the question of how intentions can be shared would still come up at some point, it would be a subsidiary question; and the leading questions would instead be about how the unshared intentions of individuals interact with each other in the often very complex process of getting things done. In conclusion, then, application to action theory of the improvisatory collaboration model of production I have presented here would have important effects not only on the answers which are given, but on the questions which are asked in this important area of philosophical investigation.

Acknowledgements

A primitive version of this paper was read at the 1996 Fall Philosophy Colloquium at Indiana University and in the Lunch-in-Theory Series of the The Center for Humanities and Arts at the University of Georgia. A somewhat more evolved version was read at the 1999 American Philosophical Association Central Division Meeting in New Orleans. I thank these audiences for helpful comments and questions; and especially Brian Smith and David Vessey, my commentators in Indiana and New Orleans, respectively. This paper has also been improved by comments from Philip Agre, David Chapman, Randolph Clarke, Randall Dipert, Edward Halper, Katarzyna Paprzycka, and Clark Wolf. Very special thanks to Russ Hallauer, Mike Mills, and Jason Slatton for patiently and reflectively answering a lot of questions about songwriting; to Cecilia Herles for acting as liaison to The Lures; to Bertis Downs, R.E.M.'s astonishing personal manager, for acting as liaison to R.E.M.; and to Kevin O'Neil and Meredith Williams at R.E.M./Athens for courtesy and efficiency beyond the call of duty.

References

- Agre, Philip E. and Chapman, David. 1990. What Are Plans for? *Robotics and Autonomous Systems* 6: 17-34.
- Arendt, Hannah. 1958. *The Human Condition*. Chicago & London: The University of Chicago Press.
- Aristotle. 1979. *Aristotle's Metaphysics*, Hippocrates G. Apostle, trans., Grinnell, IA: The Peripatetic Press.

- Aristotle. 1985. *Nicomachean Ethics*, Terence Irwin, trans., Indianapolis: Hackett Publishing Company.
- Brand, Myles. 1986. Intentional Actions and Plans. In *Midwest Studies in Philosophy*, Volume X, Peter A. French, Theodore E. Uehling, Jr., and Howard K. Wettstein, eds., Minneapolis: University of Minnesota Press, 213-30.
- Brand, Myles, ed. 1970. *The Nature of Human Action*. Glenview, IL: Scott, Foresman and Company.
- Bratman, Michael E. 1999. *Faces of Intention: Selected Essays on Intention and Agency*. Cambridge, UK: Cambridge University Press.
- Bratman, Michael E. 1992. Shared Cooperative Activity. *The Philosophical Review* **101**, No. 2: 327-41.
- Bratman, Michael E. 1990. What Is Intention? In *Intentions in Communication*, Philip R. Cohen, Jerry Morgan, and Martha E. Pollack, eds., Cambridge, MA and London: The MIT Press, 15-31.
- Bratman, Michael E. 1987. *Intention, Plans, and Practical Reason*. Cambridge, MA & London: Harvard University Press.
- Cohen, Philip R., Jerry Morgan, and Martha E. Pollack, eds. 1990. *Intentions in Communication*. Cambridge, MA and London: The MIT Press.
- Dipert, Randall. 1993. *Artifacts, Art Works, and Agency*. Philadelphia: Temple University Press.
- Dipert, Randall. 1986. Art, Artifacts, and Regarded Intentions. *American Philosophical Quarterly* **23** (4): 401-8.
- Dreyfus, Hubert L. and Stuart E. Dreyfus with Tom Athanasiou. 1986. *Mind Over Machine: The Power of Human Intuition and Expertise in the Era of the Computer*. Oxford, UK: Blackwell.
- Elster, Jon. 1985. *Making Sense of Marx*. Cambridge UK and Paris: Cambridge University Press and Maison des Sciences de l'Homme.
- Flanagan, Bill. 1986. *Written in my Soul: Rock's Great Songwriters Talk About Creating Their Music*. Chicago and New York: Contemporary Books, Inc.
- Gilbert, Margaret. 1989. *On Social Facts*. London and New York: Routledge.
- Goel, Vinod. 1995. *Sketches of Thought*. Cambridge, MA and London: The MIT Press.
- Gray, Marcus. 1997. *It Crawled From the South: An R.E.M. Companion*, revised edition. New York: Da Capo Press.
- Grumley, John Edward. 1992. Two Views of the Paradigm of Production. *Praxis International* **12**: 181-204.
- Habermas, Jürgen. 1981/1984. *The Theory of Communicative Action*, Thomas McCarthy, trans., Boston: Beacon Press.

- Joas, Hans. 1996. *The Creativity of Action*. Jeremy Gaines and Paul Keast, trans., Chicago: The University of Chicago Press.
- Kincaid, Harold. 1986. Reduction, Explanation, and Individualism. *Philosophy of Science* **53**: 492-513.
- Marx, Karl. 1883/1976. *Capital: A Critique of Political Economy*, Volume I, Ben Fowkes, trans., London: Penguin Books in association with New Left Review.
- Marx, Karl. 1932/1967. *Writings of the Young Marx on Philosophy and Society*, Lloyd Easton and Kurt Guddat, trans., Garden City, NY: Doubleday & Company, Inc.
- Marx, Karl. 1953/1973. *Grundrisse*, Martin Nicolaus, trans., New York: Vintage Books.
- Mele, Alfred R., ed. 1997. *The Philosophy of Action*. Oxford and New York: Oxford University Press.
- Mele, Alfred R. 1992. *Springs of Action: Understanding Intentional Behavior*. New York & Oxford: Oxford University Press.
- Pollack, Martha E. 1992. The Uses of Plans. *Artificial Intelligence* **57**: 43-68.
- Preston, Beth. The Functions of Things: A Philosophical Perspective on Material Culture. Forthcoming in July 2000 in *Matter, Materiality and Modern Culture*, Paul Graves-Brown, ed., London: Routledge.
- Preston, Beth. 1998a. Why is a Wing Like a Spoon? A Pluralist Theory of Function. *The Journal of Philosophy* **95** (5): 215-254.
- Preston, Beth. 1998b. Cognition and Tool Use. *Mind & Language* **13** (4): 513-547.
- Reynolds, Peter C. 1993. The Complementation Theory of Language and Tool Use. In *Tools, Language, and Cognition in Human Evolution*, Kathleen R. Gibson and Tim Ingold, eds., Cambridge, UK and New York: Cambridge University Press.
- Searle, John R. 1990. Collective Intentions and Actions. In *Intentions in Communication*, Philip R. Cohen, Jerry Morgan, and Martha E. Pollack, eds., Cambridge, MA and London: The MIT Press, 401-15.
- Suchman, Lucy. 1987. *Plans and Situated Action*. Cambridge, UK: Cambridge University Press.
- Tuomela, Raimo. 1995. *The Importance of Us*. Stanford, CA: Stanford University Press.
- Velleman, J. David. 1997. How To Share An Intention. *Philosophy and Phenomenological Research*, **57** (1): 29-50.
- Zollo, Paul. 1997. *Songwriters on Songwriting*. New York: Da Capo Press.