

Double Heuristics and Collective Knowledge: the Case of Expertise

Stephen Turner. 2012. Double Heuristics and Collective Knowledge: the Case of Expertise
 The definitive version of this paper has been published in *Studies in Emergent Order* 5: 82-103, all rights reserved. <http://studiesinemergentorder.org/current-issue/sieo5-64/>

Stephen Turner
 Professor Stephen P. Turner
 Department of Philosophy

Expert knowledge, heuristics, Michael Polanyi, social epistemology, democracy, consilience
 There is a large literature on social epistemology, some of which is concerned with expert knowledge. Formal representations of the aggregation of decisions, estimates, and the like play a larger role in these discussions. Yet these discussions are neither sufficiently social nor epistemic. The assumptions minimize the role of knowledge, and often assume independence between observers. This paper presents a more naturalistic approach, which appeals to a model of epistemic gain from others, as mutual consilience—a genuinely social notion of epistemology. Using the example of Michael Polanyi’s account of science as an illustration, it introduces the notion of double heuristics: that individuals, each with their own heuristics, each with cognitive biases and limitations, are aggregated by a decision procedure, like voting, and this second order procedure produces its own heuristic, with its own cognitive biases and limitations. An example might be the limited ability of democracies to assimilate expert knowledge.

There is a burgeoning literature on social epistemology. Some of it purports to illuminate the problem of expert knowledge. Much of this literature applies epistemological notions, such as reliabilism, to expert claims, which are interpreted in terms of notions familiar to epistemology, such as testimony. Another body of literature is concerned with the contrast between individual and collective rationality or collective knowledge, and is concerned with issues of emergence, specifically with the claim that collective knowledge processes are different from and arrive at different and better results than individual knowledge acquisition. Many of these are discussions of collective rationality, and use formal methods. To do so, they typically simplify the issues by assuming independent individual judges. Independence implies epistemic independence, namely, that people act on their own knowledge. Discussions of the related problem of expertise typically follow the same pattern: expertise is compared to testimony, which the *individual* judges as reliable. The classic prisoner’s dilemma is based explicitly on the mutual ignorance of the prisoners with respect to intentions. Both the social relations between the prisoners and the possibility of sharing knowledge are defined away. In this respect, these approaches follow standard economic theories, which assume information equality or otherwise assume the irrelevance of differences in quality of information between market participants in market transactions. Nor is there an easy alternative to these assumptions. Asymmetric information theorizing, for example, is technically difficult even in small scale transactions with limited dimensions of relevant information, such as theorizing the issues of agency in a used car purchase. Expanding these considerations and expanding considerations of variations between market participants makes calculating outcomes intractable.

My concern in what follows will be with cases of extreme “information asymmetry” in which members of the audience of the experts have knowledge that is different from the knowledge of experts. The knowledge is often relevant, and the decision by a member of the audience of the expert to accept or reject the expert’s claims is not, as the models imply, based simply on beliefs

about the reliability of the experts, but on the knowledge that the member of the audience already has, and has solid grounds for. In these cases, the better model for understanding how the member of the audience assesses the expert involves the content of the knowledge, not merely the evaluation of the expert. The member of the audience makes an epistemic judgment on the primary knowledge material, not merely on the credentials of the expert. My concern in this paper will be to provide a way of thinking about these epistemic judgments. But this discussion will be mostly illustrative.

My primary aim will be to suggest a way of thinking about the aggregation of these judgments and how this aggregation can be understood. In the course of this I will treat the problem of expert knowledge as a special case of knowledge aggregation. My suggestion will be that the application of specific decision procedures, such as voting, produces, at the collective level, an emergent form of knowledge acquisition with its own features. Nothing about this account, however, requires an appeal to super-individual entities or processes, collective intentionality, and so forth. My point, rather, will be that to understand these processes it is necessary to eliminate some familiar prejudices about knowledge acquisition and our dependence on others. To put it in a slogan, my point is that “collective epistemology” or social epistemology has failed to be either sufficiently social or sufficiently epistemological. My approach will be to bring both back in, without resorting to appeals to collective facts.

The Background

Social epistemologists have long been concerned with cases in which collective decisions, through such means as judges voting on verdicts, differ from individual judgments. Philip Pettit formulates two of the standard assumptions of these cases as follows:

- you are amongst many people who face a certain question;
- you are all equally intelligent, equally informed, and equally impartial. (2006a: 179)

In normal social situations, neither of Pettit’s assumptions holds, and in expert situations the exact opposite is assumed: that people who know something you don’t are more intelligent, and may have fewer biases. Moreover, they claim to understand the question better than you do, a second-order claim with unusual epistemic properties: among other things, it undermines the usual ways of judging the claims of others as one understands them.

Pettit interprets the epistemic issues which arise when one’s own opinion differs from the majority opinion as a normative problem: when should one defer to majority testimony and when should one not do so? To answer this, he adds to his standard assumptions about equal knowledge and the like

- you, however, differ from most others in the answer you give;
- you are aware that these things [i.e. the full set of prior assumptions about equal intelligence, knowledge, and impartiality] are true. (2006a: 179)

If we begin with the question of how we would ever know that the people you are among “are all equally intelligent, equally informed, and equally impartial,” we can see that these models are epistemically problematic. In this case, the problem of expert knowledge is excluded by definition: expert knowledge is precisely that case in which equality in intelligence, knowledge, and impartiality are denied by both the expert claiming expertise and the member of the expert’s audience assessing it. So any direct application of these assumptions to the notion of expert knowledge will fail. They could apply only if the problem of expert knowledge is reduced to the problem of assessing expert

“testimony,” so that the question of when to defer to an expert becomes a problem of when to defer to majority opinion about the expert. As noted earlier, this way of understanding assent to expert claims strips the knowledge claims themselves of content, making the knowledge of the content possessed by the members of the audience irrelevant.

Another approach to the problem of aggregation involves the suspicion of systematic bias in the production of expert knowledge claims. Miriam Solomon has in mind the idea that gender biases and the like distort theory choice. Although this approach was originally motivated by feminist considerations, it applies more generally. Solomon constructs this as a problem of epistemic diversity, and, rather than dealing with expert authority, deals with the problem of theory choice in the scientific community as a model for epistemic decision-making generally. She suggests that what is needed is a means of eliminating the effects of biases by balancing biases against one another and demanding empirical support for all options (2006a: 37-8; cf. 2006b). This differs slightly from Pettit’s approach, by assuming that despite being equally impartial, people have biases. But it also takes a valuable step in epistemologizing the problem of aggregation. “Theory choice by a community” is a collective procedure, although it is a theoretical construction of the observer rather than something that scientists collectively perform, as in voting. And the term “biases” does provide some, very minimal, epistemic content to the notion of epistemic diversity. But this is too small a step. The problem of bias is dwarfed by a much bigger problem of epistemic diversity: that we know different things and can understand different things.

Double Heuristics

The Pettit assumptions are simply false. The true, but difficult, situation is this. We know something already, about experts and what makes them acceptable, and also often about the content of their claims. Our knowledge is not “equal” to that of others, or the same as others. We have our own experiences and practical and sometimes theoretical knowledge that either fits or fails to fit with the expert claims. The (descriptive rather than normative) epistemological problem is to understand this kind of knowledge and to understand how we rely on our knowledge of others—the social aspect—and how we use our own knowledge to assimilate it to what others know.

The literature in social epistemology has been dominated by technical solutions. But if one adds actual epistemology to the social, by considering how we use the content of the knowledge of others as distinct from simply accepting on trust, these solutions become unstable. A less technical, but more usable way of conceptualizing the problem would be this: to think of our use of the knowledge of others as governed by more or less standard heuristics, which may go wrong in abnormal situations, and thus have biases. To discuss this problem, however, one needs some sort of model. The image of the individual knower I propose to work with is itself a simplification, as it must be. But it is a simplification that allows for a discussion, however limited itself, of the general problem of knowledge aggregation. The model is this: the individual is limited, operates with complexity reducing epistemic strategies, arrives at knowledge, and makes knowledge judgments. The individual knower, in short, uses heuristics, which, like all heuristics, work in some situations and not others. They have biases and blind spots. This is hardly an innovative idea, of course. It is enshrined in the literature on empirical models of rational choice (Tversky and Kahneman 1974, 1981).

The value of this starting point is this: it provides us with a model for thinking of emergent and “collective” forms of the same thing. We can think of decision procedures, such as democracies, and aggregation mechanisms without collective decisions but with “collective” outcomes, such as markets, as themselves being heuristics. We can think of procedures which function as if they were decision procedures, such as market decisions that put firms into bankruptcy, as heuristics as well.

The market itself is not a heuristic, nor is a rule like majority voting. But it is a procedure which makes selections. If a procedure is understood as made up of people operating with particular heuristics that include heuristics about the knowledge of others and how to assimilate it, plus some sort of analogue to a decision, these emergent processes themselves can be understood to have normal situations in which they “work” and others in which their “biases”—biases being the source of the efficiency gained by heuristics—lead to bad results.

This notion of a “double heuristic” then allows us to conceptualize the issues that arise with, for example, the (now commonplace) claim that liberal democracy needs to be abolished to save the human race in the face of global warming. We can ask what sort of alternative collective heuristic there is, such as the heuristic of uncritical acceptance of the assertions of scientific experts, and what the biases of this heuristic might be; which is to say, to ask what the normal and abnormal situation is for this heuristic when it is understood as a heuristic made up of the aggregation of the heuristics of people judging experts with the biases of these heuristics, and of experts themselves making decisions with their biases.

Nevertheless, the contrast between individual and collective results is an important one, and can be generalized beyond voting examples. If we think of individual and collective procedures of dealing with questions, one thing is immediately obvious—collective “decisions,” whether it is “the market decides,” voting, or counting up the guesses about beans at the county fair, all happen differently than individual decisions. One makes up one’s mind about a bean-count, and decides to submit the estimate. The collective act of adding up the estimates and taking a mean takes place on a schedule. No “minds” are made up. The market makes pricing decisions continuously—buyers make them one at a time, sellers look at their stock and respond by changing prices. The collective result is a theoretical construction from the actual prices that are charged transaction by transaction. Juries deliberate and vote, in accordance with protocols and admonitions. Jurors decide they are persuaded when they reach an epistemic threshold of doubt that they individually determine by self-examination, but the collective threshold is unanimity or some other rule.

Putting Epistemology Back In

The problem of experts in politics has epistemic content, but the content is highly problematic. Both the literature in what Alvin Goldman calls “classical social epistemology” and the literature of the Mertonian sociology of science have focused on the authority of science. As I have noted, epistemologists, not fond of the term “authority,” have construed the issue in terms of testimony. This allows authority to be interpreted in traditional philosophical terms, in this case in terms of reliability and therefore in terms of reliabilism, as an epistemology. Other social epistemologists have focused on cases in which collective knowledge is “better” than individual knowledge, or at least different. The model in these cases is guesses at the number of beans in a jar at a country fair: the mean is closer to the correct number than the individual guesses.

What is striking about these cases is their tractability to formal reasoning. One can put up a matrix of judges’ votes, for example, and show that the collective result of the votes differs from individual votes (List and Pettit 2002). What is also striking is their inattention to content. Guesses about beans have little epistemic content. Moreover, one’s knowledge of the guesses of others is irrelevant or assumed to be irrelevant.

But actual cases of judgments of expert opinions in political contexts are far richer, in a number of ways. The citizen has a variety of epistemic resources, including beliefs about the world, experiences, grounds for making judgments about the sources of claims, and personal knowledge to bring to the making of beliefs about the subject matter discussed by experts and of the experts themselves.

The classic discussion of this is Brian Wynne's paper "May the Sheep Safely Graze?" (1996) that considered the case of nuclear power experts making claims about the effect of radiation. The sheep owners to whom the expert discussion was addressed were skeptical of the claims, based on their knowledge of the actual grazing habits of the sheep. This is a case of two kinds of knowledge fitting together. But the fitting together involved content, and the product of the fitting together would alter what each side believed, rather than merely combining independent estimates to create a third without altering at least the epistemic weight of the beliefs of one side or the other.

Assuming content away, using the model of bias and similar devices, does not help much with these cases. Empirically, content-free judgments about expertise based on the pure kinds of assessments involved in testimony, in the extremely abnormal and purified sense of testimony in which the reliability of the witness is the only consideration, are nowhere to be found. When people on a jury assess real testimony, they do so on the basis of their prior knowledge and actual experience of the world, as lawyers know very well, which is why they are careful to select juries that are as ignorant as possible about the topics they are going to hear testimony on.

A Classic Model of Science as a Collective Heuristic

Philosophy of science provides some models for thinking about "fitting together," such as Michael Polanyi's picture of science as a big jigsaw puzzle into which we each fit our little pieces of knowledge. Polanyi did this more than once, and the differences between his accounts are revealing with respect to the phenomenon of the relation between individual and collective heuristics. In the essays collected in *The Logic of Liberty*, he describes the collective process in terms of the

adjustment of each scientist's activities to the results hitherto achieved by others. In adjusting himself to the others each scientist acts independently, yet by virtue of these several adjustments scientists keep extending together with maximum efficiency the achievements of science as a whole. At each step a scientist will select from the results obtained by others those elements which he can use best for his own task and will thus make the best contribution to science; opening thereby the field for other scientists to make their optimum contribution in their turn— and so on indefinitely. ([1951] 1980: 34-35)

This implies what I am calling a heuristic—an efficient method for getting "collective" knowledge results out of individual contributions to a process of aggregating knowledge.

The only way to get the job finished quickly would be to get as many helpers as could conveniently work at one and the same set and let them loose on it, each to follow his own initiative. Each helper would then watch the situation as it was affected by the progress made by all the others and would set himself new problems in accordance with the latest outline of the completed part of the puzzle. The tasks undertaken by each would closely dovetail into those performed by the others. And consequently the joint efforts of all would form a closely organized whole, even though each helper would follow entirely his own independent judgment. ([1951] 1980: 35)

A collective process is defined by its decision procedure, which in these early writings Polanyi described as a "twofold condition," consisting of rapid publicity plus acclamation, in which

. . . each suggested new step can be readily judged as to its correctness or otherwise, and that each new step is rapidly brought to the notice of all participants and taken into account by

them when they make their own next move. ([1951] 1980: 36)

Pure science, as distinct from applied science or technology, required this heuristic, rather than others.

Polanyi was arguing against planned science, which represented a distinct heuristic or set of possible heuristics. On the surface, planning seemed to be the perfect way to avoid waste in science and produce valuable results. A great deal of thinking, and a social movement of Left-wing scientists in Britain, promoted planning generally, and the planning of science specifically, in the 1930s, as part of a general enthusiasm for the idea of planning. Polanyi's argument against planning had to do with the problem of knowledge.

Put negatively, planning is simply impracticable, at least for the most important problems in science: No committee of scientists, however distinguished, could forecast the further progress of science except for the routine extension of the existing system . . . the points at which the existing system can be effectively amended reveal themselves only to the individual investigator. And even he can discover only through a lifelong concentration on one particular aspect of science a small number of practicable and worthwhile problems. ([1951] 1980: 89)

The argument against planned science then depends on an argument about the distribution of knowledge. Translated into our terms, the argument is this: knowledge in science is specialized, so a heuristic that depends on the knowledge of some small group or any collective decision-making process will lose the advantages that the specialist has in deciding how to pursue his or her own problems.

But science cannot avoid collective decision procedures. Money has to be doled out. In this respect it is necessary to construct another decision procedure. Here, notoriously, Polanyi and his sympathizers found the going more difficult. The process of doling out determines the content of the science that can be extended. So there is no escaping the consequences of the system of supporting science. The best that can be done is to have a system that retains the advantages of the heuristic described above.

The pursuit of science can be organized . . . in no other manner than by granting complete independence to all mature scientists. They will then distribute themselves over the whole field of possible discoveries, each applying his own special ability to the task that appears most profitable to him. Thus as many trails as possible will be covered, and science will penetrate most rapidly in every direction towards that kind of hidden knowledge which is unsuspected by all but its discoverer; the kind of knowledge on which the progress of science truly depends. The function of public authorities is not to plan research, but only to provide opportunities for its pursuit. All that they have to do is provide facilities for every good scientist to follow his own interests in science. ([1951] 1980: 89-90)

With this general heuristic in mind, one can turn to problems of institutional design. This is the model that the Howard Hughes Medical Institute applies by giving six-year appointments on a principle of scientific promise with no restrictions on what the scientist will choose to do. The American National Science Foundation does this by evaluating proposals on merit by specialized panels. In each case, of course, many choices need to be made, each of which involves biases, biases which diminish the odds of science penetrating in some of the directions where hidden knowledge can be revealed.

The epistemology involved here is still individualistic: the discoverer seeks hidden knowledge individually. Verification is collective, by acclamation. Both of these are of course caricatures of the actual practice of science. But there is already a hint in these early writings of the problems of fitting knowledge together, which Polanyi later makes more central, in the idea of extension and the admission that specialization of a very extreme kind is characteristic of scientists. What disqualifies these scientists from making the kinds of general judgments about how science should be planned is this very specialization, and this also disqualifies them for the role of acclaiming scientific achievements. In some of his early writings, Polanyi spoke of the decision-makers in science as a group analogous to Plato's Guardians. But the Guardians were, so to speak, possessors of the most general knowledge; scientists, in contrast, are specialists. These were conflicts that he later resolved, in his classic essay "The Republic of Science" (1962).

The resolution is of interest not only because of its explicit appeal of the concept of spontaneous coordination, but because of the new kind of knowledge relation he identified as the connecting link between the specialized worlds of science. The new emphasis is especially relevant to "social epistemology," because Polanyi makes an explicit contrast between activities which are "coordinated" and those that are not. The examples favored by Pettit and Solomon are uncoordinated—the judges rendering verdicts, people guessing locations of submarines or numbers of beans in a jar are independent and take no account of the knowledge that others have. In science, Polanyi says, the result of this would be as follows:

Each scientist would go on for a while developing problems derived from the information initially available to all. But these problems would soon be exhausted, and in the absence of further information about the results achieved by others, new problems of any value would cease to arise and scientific progress would come to a standstill. (1962: 54)

This is enough to show that science is a coordinated activity, "and it also reveals the principle of their coordination. This consists in the adjustment of the efforts of each to the hitherto achieved results of the others. We may call this a coordination by mutual adjustment of independent initiatives—of initiatives which are coordinated because each takes into account all the other initiatives operating within the same system" (1962: 54). The emphasis on coordination is new. And one soon sees why, when he considers the analogue to a decision procedure in science.

When he turns to the problem of explaining the way in which coordination works, he reasons in terms of what I have called here double heuristics. He introduces the image of the puzzle: "Imagine that we are given the pieces of a very large jig-saw puzzle, and suppose that for some reason it is important that our giant puzzle be put together in the shortest possible time" (1962: 55). He gives three examples of increasingly effective methods of doing this with a group of helpers. The first is independence, in which each person is given a few pieces to fit together: "Suppose we share out the pieces of the jig-saw puzzle equally among the helpers and let each of them work on his lot separately" (1962: 55). He remarks "it is easy to see that this method, which would be quite appropriate to a number of women shelling peas, would be totally ineffectual in this case, since few of the pieces allocated to one particular assistant would be found to fit together" (1962: 55). The second would be to supply each with copies of all the pieces, but still make them work independently: "We could do a little better by providing duplicates of all the pieces to each helper separately, and eventually somehow bring together their several results" (1962: 55). Polanyi's verdict is that "even by this method the team would not much surpass the performance of a single individual at his best" (1962: 55).

The best collective heuristic would be this, which Polanyi takes to be a model of the coordination heuristic for science itself.

The only way the assistants can effectively cooperate and surpass by far what any single one of them could do, is to let them work on putting the puzzle together in sight of the others, so that every time a piece of it is fitted in by one helper, all the others will immediately watch out for the next step that becomes possible in consequence. Under this system, each helper will act on his own initiative, by responding to the latest achievements of the others, and the completion of their joint task will be greatly accelerated. (1962: 55)

He goes on to note that “Such self-coordination of independent initiatives leads to a joint result which is unpremeditated by any of those who bring it about. Their coordination is guided as by an invisible hand towards the joint discovery of a hidden system of things. Since its end-result is unknown, this kind of cooperation can only advance stepwise, and the total performance will be the best possible if each consecutive step is decided upon by the person most competent to do so” (1962: 55). He expands the thought with this image: “We may imagine this condition to be fulfilled for the fitting together of a jig-saw puzzle if each helper watches out for any new opportunities arising along a particular section of the hitherto completed patch of the puzzle, and also keeps an eye on a particular lot of pieces, so as to fit them in wherever a chance presents itself” (1962: 55). The “competence” in question is epistemic and localized. The person knows more about this little patch. But in the later text Polanyi admits that this has implications for the acceptance and evaluation of science.

From the point of view of discovery, the argument is the same. The result is a heuristic that is more effective than the others, at least with respect to speed. “The effectiveness of a group of helpers will then exceed that of any isolated member, to the extent to which some member of the group will always discover a new chance for adding a piece to the puzzle more quickly than any one isolated person could have done by himself” (1962: 55). The term invisible hand invites comparison to markets, and Polanyi suggests that the comparison is apt, but the conclusion should not be that science is a special case of markets, but that “the coordinating functions of the market are but a special case of coordination by mutual adjustment” (1962: 66). This notion of mutual adjustment, however, is a new emphasis. The contrast to the market is this:

In the case of science, adjustment takes place by taking note of the published results of other scientists; while in the case of the market, mutual adjustment is mediated by a system of prices broadcasting current exchange relations, which make supply meet demand. But the system of prices ruling the market not only transmits information in the light of which economic agents can mutually adjust their actions; it also provides them with an incentive to exercise economy in terms of money. (1962: 56)

The motivations in science differ: “by contrast, the scientist responding directly to the intellectual situation created by the published results of other scientists is motivated by current professional standards” (1962: 56). Current professional standards, as we will see, play a special role. The choices about what lines of inquiry to follow that the scientist makes in the face of these standards, Polanyi admits, have an economic character. The scientist does not want to waste time on insoluble problems, or those which are too easy, or hypotheses that are implausible from the standpoint of present professional knowledge. But originality is prized by these professional standards.

The puzzle, from the point of view of the application of collective heuristics, is in the decision procedure, meaning in this case understanding how these standards are applied. As noted, Polanyi earlier seemed to rely on a kind of general acclamation. Now he recognizes the conflict between specialization and the idea that each scientist is omniscient to act as judge.

No single scientist has a sound understanding of more than a tiny fraction of the total domain of science. How can an aggregate of such specialists possibly form a joint opinion? How can they possibly exercise jointly the delicate function of imposing a current scientific view about the nature of things, and the current scientific valuation of proposed contributions, even while encouraging an originality which would modify this orthodoxy? (1962: 59)

The solution to this is of course to invoke a new collective heuristic, or what Polanyi calls an “organizational principle.”

In seeking the answer to this question we shall discover yet another organisational principle that is essential for the control of a multitude of independent scientific initiatives. This principle is based on the fact that, while scientists can admittedly exercise competent judgment only over a small part of science, they can usually judge an area adjoining their own special studies that is broad enough to include some fields on which other scientists have specialised. We thus have a considerable degree of overlapping between the areas over which a scientist can exercise a sound, critical judgment. And, of course, each scientist who is a member of a group of overlapping competences will also be a member of other groups of the same kind, so that the whole of science will be covered by chains and networks of overlapping neighbourhoods. Each link in these chains and networks will establish agreement between the valuations made by scientists overlooking the same overlapping fields, and so, from one overlapping neighbourhood to the other, agreement will be established on the valuation of scientific merit throughout all the domains of science. (1962: 59)

Crudely, there are scientists in adjacent areas of science who know enough to judge the work of the specialist, and this enforces consistency in the application of professional standards.

These relations of adjacency produce a network, which is the place in which we can interpret this as a collective heuristic. Polanyi puts this in his own terms:

This network is the seat of scientific opinion. Scientific opinion is an opinion not held by any single human mind, but one which, split into thousands of fragments, is held by a multitude of individuals, each of whom endorses the other's opinion at second hand, by relying on the consensual chains which link him to all the others through a sequence of overlapping neighborhoods. (1962: 59-60)

But scientific opinion, even when it is distributed in this way (and Polanyi has more to say about whose opinions count most) is still opinion, as Polanyi always insisted. These procedures, and the collective heuristics system they create through their operation, are not epistemic guarantors of truth.

In his earlier writings Polanyi discussed the corruption of scientific opinion. In “The Republic of Science” he concedes that the system of control by the application of professional standards can lead to bad results. But on a collective level, it is, in our terms, the best heuristic.

scientific opinion may, of course, sometimes be mistaken, and as a result unorthodox work of high originality and merit may be discouraged or altogether suppressed for a time. But

these risks have to be taken. Only the discipline imposed by an effective scientific opinion can prevent the adulteration of science by cranks and dabblers. In parts of the world where no sound and authoritative scientific opinion is established research stagnates for lack of stimulus, while unsound reputations grow up based on commonplace achievements or mere empty boasts. Politics and business play havoc with appointments and the granting of subsidies for research; journals are made unreadable by including much trash. (1962: 61)

However, “Though it is easy to find flaws in [the] operation [of the organizational principles of this system], they yet remain the only principles by which this vast domain of collective creativity can be effectively promoted and coordinated” (1962: 61).

Evaluating these claims is not my concern here. This is an illustration of the basic concept of double heuristics. But a few points need to be made. Polanyi changes his ideas about the nature of the relations between areas of science in “The Republic of Science” by emphasizing the way in which specialists in adjacent areas evaluated new findings. He tells us little about *their* heuristics for doing so, though clearly his general ideas about professional standards are of the kind that has local variations and applications. More can be said.

The considerations Polanyi applies to science were also applied by him to the economics of planning, and this opens a related domain of inquiry. Peter Boettke argued in *The Political Economy of Soviet Socialism* that the idea of central planning was applied in earnest only for the first decade after the Russian Revolution (1990). What emerged in its place was a collection of loosely related plans. The way the plans were collected and collated had little to do with the idea of central planning in the contemporary economics literature. The models of perfect knowledge that Oscar Lange (and even Frank Knight in response to Lange) had discussed when they examined the theoretical possibility of centralized socialist planning were completely unlike the actual process. Yet well into the 1940s, planning continued to be discussed in these theoretical terms, and the successes of the Soviet Union were taken to vindicate, in some sense, the theoretical possibility of a kind of virtual knowledge of demand. Polanyi himself pointed out that, although the planners got their knowledge in ways other than the open market, they were ways that were quite mundane. As David Prychitko notes,

Polanyi argued that, as opposed to the theoretical model, the Soviet economy has been composed of numerous, conflicting planning centers—a “polycentric” as opposed to “monocentric” order. Coordination, to the extent that it occurred at all, took place not at the center of the planning hierarchy, but at the lower levels, among the individual enterprise managers who used their own discretionary authority and engaged in black market exchanges. Though the quantity and quality of outputs chosen and produced at the enterprise level became aggregated into a so-called central plan, and indeed were later published as a unified, centrally issued plan established by the directives of GOSPLAN (the Soviet central planning bureau), in fact the coordination of economic activities took place at the enterprise level, at the bottom of the hierarchy. (Prychitko 2002: 154n8)

Managers engaged in black market operations knew what the values of goods were, whether there was demand, and so forth (cf. Roberts 2005). These were imperfect means, but nevertheless means subject to real world discipline in the form of facts about prices. The managers used this information to plan their own production, and the planners aggregated these plans and tinkered with them. This was a system with its own epistemic biases, resulting in part from the limited knowledge of the contributing players. It was far from the ideal rationalized central planning. Nevertheless, it was a system that aggregated knowledge of diverse kinds into a collective result.

The most famous alternative to this form of centralized pseudo-planning in the “Socialist” countries was the decentralized Yugoslavian system of worker self-managed productive units, which operated with limited market mechanisms for resources, and more or less open markets for the consumer goods. As Prychitko showed in *Marxism and Workers Self-management* (1991), these units, which purported to solve the problem of alienation by giving the workers control of what happened to the products they made, transferred the problems of planning and making market related decisions to the workers, or rather to workers’ councils, which were supposed to be democratic, participatory, and to produce a decentralized bottom up control of the economy that eliminated the waste of financial speculation and advertising.

Democratic participation is a knowledge aggregation system. And comparing the scheme of aggregation of knowledge of an ordinary hierarchical firm to a democratically self-managed one, as Prychitko does, points out some interesting differences in how knowledge is shared. The workers had a greater propensity and incentive to share information about what was happening in their part of the production process, for example. But the workers were not especially willing to undertake the knowledge related tasks of aggregating this knowledge, or taking responsibility for decisions, problems that would be solved in an actual case of democratic rule by rewarding winners in political competition. This points to a whole range of questions about how propensities and abilities to use and share knowledge differ among organizations, and thus to the epistemic problem of what sort of different kinds of learners different organizations are, and what are the biases, efficiencies, and blind spots in the different “collective” information processing and aggregating heuristics that result from the way in which the organizations operate.

Bilateral Asymmetric Consilience

Fitting the pieces together is a metaphor, as is the term ‘network’. Polanyi doesn’t inquire into the epistemology of fitting together. What is the relevance of having the pieces fit, or having new knowledge in adjacent areas of science? Certainly it has some bearing on our sense that we are on the right track, that the previous steps leading up to the new knowledge were the correct ones, and so forth. This model suggests consilience of induction as both a ground for belief in “our piece” and the solution provided by the puzzle as a whole, but also suggests a model of collective outcomes of epistemic contributions that go beyond individual knowers. The “adjustments” which Polanyi stressed are also adjustments in what people believe and the weight they give to beliefs.

Other examples can be used to reveal the problematic results of interpretations of experiments which ignore the “social” cues on which we ordinarily and necessarily rely in coming to beliefs. It is nevertheless awkward to think epistemically about what fitting together might mean because the traditions of epistemology are individualist. To help this along, let me give an example of a non-individual epistemic notion. Here the (descriptive rather than normative) epistemological problem is to understand this kind of knowledge and to understand how we rely on our knowledge of others—the social aspect—and how we use our own knowledge to assimilate it to what others know.

Suppose that the doctor supplies a diagnosis that is based on your self-reported symptoms, but that also predicts symptoms that you did not report because you did not think they were relevant, but can now recognize as part of the syndrome. The situation is one of asymmetric knowledge, but also of distributed knowledge: the patient knows something the doctor doesn’t as well as the reverse. They are different kinds of knowledge: the doctor supplies the means of fitting together without knowing in advance at least some fact that turns out to fit. This is consilience in the original Whewell sense of correctly predicting some novel fact that the theory was not constructed to account for (Whewell 1858: 88-90). In this case, the doctor was accounting for the symptoms that

were presented.

This is different from reliabilism, also “social” and “epistemic” in a sense independent of judging testimony, yet which still reflects acceptance of the asymmetric knowledge of others. The basic thought is that the fact of consilience itself adds to the epistemic weight of the facts considered independently, in contrast to the aggregation and voting cases. The added weight is not something done by judging the expertise of the source, though this is part of the story. It is social because you don’t get the epistemic payoff, namely consilience, without the social, in this case the acceptance as minimally weighty the fact and content of the beliefs of others, which are then combined with one’s own for the payoff.

In contrast, assessing testimony adds no epistemic weight, content, or predictive power to the original testimony—it is a subjective weighing of something else’s epistemic weight (or “probative force” —you can pick your favorite term). “Consilience” or what we might in this case call “Asymmetric Bilateral Consilience,” is more than merely consistency with the diagnosis, which is another, weaker sense, which might be the case where a jury rejects testimony based on their own knowledge of relevant facts that are not consistent with the testimony, which we could call “Asymmetric Bilateral Consistency.” But maybe we could do the world a favor and not call them anything fancy.

Like all heuristics, this one can go wrong, and typically “going wrong” means applying them under circumstances that don’t work for reasons the user does not know. The place that this heuristic most obviously can go wrong is when the novel fact predicted by the expert is not as independent of the facts known to the non-expert as one or the other might believe. This would occur, for example, when both expert and non-expert are describing facts in accordance with a common but unacknowledged ideology, as when the non-expert reads an ideologically selective newspaper report and “discovers” that this fits an expert “truth” that is generated in a hidden way by the same ideology. The hypothesis of a hidden variable producing the facts, like the case of assumptions of independence in statistics, is normally one that is beyond the limits of the heuristic itself. And this raises questions about the way in which the heuristics we employ in assessing and giving weight to other people’s opinions, for example in the case of the problem described by Pettit of when to defer to majority opinion, are themselves potentially compromised. The heuristics are limited by our failure or inability to assess whether these opinions are indeed the result of more or less independent judgments of others, or are the product of a consensus produced artificially by some other means. This fits with the many social psychological experiments on conformity of the 1950s, to be discussed shortly, which showed how readily people would accept false beliefs if a group of which they were a part affirmed them. If the subjects of the experiments had known that there was a conspiracy by the members of the group to affirm these beliefs, specifically, that they were not independently arrived at, they would have responded differently. In the case of the physician’s diagnosis, the same point holds: if one’s descriptions of one’s own symptoms are influenced by the same therapeutic ideology as the physician, the independence of the two acts of description, one motivated from and derived from the physician’s diagnosis, the other from the private experience of the patient, is an illusion.

Polanyi’s puzzle model of science depends on giving epistemic weight to the beliefs of others, and “fitting” in some way with their beliefs. What is fitting, in an epistemic sense? This strong kind of consilience is one example of fitting. Collective rationality, extended mind, etc. models locate the knowing in the collective knower. Polanyi’s model doesn’t do that: it relies on the notion of networks. But it also allows for, and indeed forces us to begin, thinking about the kinds of heuristics, both individual and (actually) collective rather than merely social, that are in fact employed, and how they produce the double heuristic pattern. Bilateral asymmetric consilience is a very strong source of epistemic weight. But it too makes some assumptions about independence: the

scientific workers are supposed to be specialists working on their own little patch of science and thus uninfluenced by what is going on in the parts of the network, or the puzzle, that they are not working on. And this is not the only form of fitting: there can be heuristics that work using weaker heuristics, such as deference to scientific authority as such and consciously fitting our observations and beliefs to whatever appears as the consensus. But of course these heuristics have their own weaknesses and biases.

The Cutlery Problem

As part of the training for American diplomats, they are shown a table of cutlery, with dozens of implements. Why? So they know how to use the right fork, and avoid a diplomatic gaffe. One could have expert knowledge of such things, but most of us, in the face of the problem of how to use a fork, use the simple heuristic “when in Rome, do as the Romans.” This is suggestive. Isn’t it normally right to accept what others believe, or to give it great weight? Isn’t it a significant problem for a believer in x that others reject x ? Doesn’t this produce a potentially large explanatory burden?

Heuristics work in normal situations. This one would work as well, unless one were copying the wrong person. And here we have heuristics as well: the Castilians who, according to the apparently false legend, started lisping because King Phillip lisped were following a heuristic, and successive generations followed them based on their own heuristic of talking like their betters. Does this make sense as a normative rule? Of course not, in the sense of an abstract approach to ethical truth. But this is misleading, as are a large number of psychology experiments which come to mind in these cases.

Here are a few examples. Vance Packard, in the 1950s, gave the example of an umbrella for sale in a department store. At a low price, the umbrella failed to sell. The price was doubled, and it sold briskly. What is going on here epistemically? The question of whether an umbrella is any good is not one that we are ordinarily able to determine by looking at it. The heuristic that says “you get what you pay for” would lead you to think a cheaply priced umbrella was no good; a higher priced one would be good. Since this is the only information we have in this case, the rule misleads us.

The Asch conformity experiments involved subjects who were placed with a group of confederates who gave different measurements of a line. Asch wondered about the circumstances under which the subject would capitulate to the majority. He found that some people did, others were confused, and others were resistant. The findings were that if the confederates were unanimous, people conformed; if there were a few dissenters, or even an inconsistent dissenter, the rate of conformity dropped drastically. The Milgram experiments seemed to show a lot of conformity. But they can’t even be run again because people would know what was going on. Subjects did question the experiment as it went on, but the experimenters were trained to fend the questions off.

The experiments are all about abnormal situations: settings are information poor, so that often the only added information is the beliefs of others; access to other opinions is manipulated; information costs are manipulated or high, so that only cheap information supports the outcome. So these are, from the point of view of normal heuristics, abnormal situations. They are useful only for revealing our normal heuristics. But they also show that one can create abnormal situations that allow these normal social heuristics to be used against people. The problem, as indicated in connection with bilateral asymmetric consilience, is that the heuristics themselves, by definition, do not detect the abnormality of the situation. And this is particularly important in relation to notions like consensus, which we know from these experiments to have powerful effects on people’s beliefs: assumptions about the independence of the parties to the consensus are almost certainly false, but the illusion of agreement is still powerful.

Expert Knowledge and Democracy

We can think of the problems with which we began, the problem of expertise in liberal democracy, in terms relative to normal and abnormal heuristics rather than “truth.” Simple models of democracy assume that people have interests, knowledge of the basic functions of government and information on how they are being carried out, and a capacity to assess the interests and motives of others. They operate to advance their own interests, and make common cause with those who can articulate interests that coincide with theirs, or are not in too great a conflict with theirs, or match their vision of a harmonious, decent society. But the heuristics they employ, according to various bodies of research, involve getting information from trusted sources, such as local influentials, rather than making these assessments on their own. This is a heuristic: trust those who you know to be well-informed, responsible, and with a stake in the same things you have a stake in.

“Influence,” however, is a crude term, which implies some sort of occult psychological force. Perhaps, in these cases, it should be understood epistemically, in the manner of the physician. If the influential says things that imply things that the hearer knows, it should help strengthen both of their beliefs. Even the weak epistemic support provided by the fact that the influential, who is similarly situated, has these beliefs is still epistemic rather than a matter of occult psychology. Nevertheless the reliance on influentials is a system with obvious biases. Some of these, under normal circumstances, are beneficial. It provides an obvious protection against such classic evils of democracy as demagoguery: if one is relying along with one’s friends on one’s local influentials, it is unlikely that waves of political enthusiasm for false prophets will overwhelm the system. At the same time, it is a heuristic that is relatively immune to totalitarian ideology: local influentials tend to think for themselves and not behave in a uniform manner. But it is also true that such a system is not especially receptive to assertions of expert authority that do not operate through trusted local influentials.

Of course, all of this is a greatly simplified model even of traditional democracies. Modern democracies are composed of people with memberships in a variety of groups, which operate in ways that differ, and have their own heuristics. Because of the sheer variety of heuristics found in different groups, the possible combinations of them in a collective procedure are also large. The problem of designing a decision procedure that produces good results given the individual heuristics of the participants is daunting. But posing the question in terms of double heuristics does allow us to give these questions some content. What if it is claimed that liberal democracy, because of its open discussion, which fails to adequately defer to scientific consensus, needs to be abolished or corrected by policing utterances about science in order to save the world by enacting proper policies on climate change? These are translatable into questions about the joint operation of individual and collective heuristics, and pose questions that might be solved by altering collective decision procedures to produce heuristics with different biases.

We can ask the same kinds of critical questions about the double heuristics involved in the production of collective expert opinion out of individual expert heuristics. Does scientific groupthink and grant-driven bandwagoning make science unreliable as a source of the kinds of facts that political bodies need to make? What if Ulrich Beck was right to complain that experts had a conservative epistemic bias which led them to be skeptical about the evidence for risks, and to systematically under-rate risks, and we have a system for collective decision-making that defers to experts? (Beck 1995). We magnify the error producing potential of the system in a specific direction. But if we have a system in which experts benefit by asserting risks, we have the opposite result.

In the end it will be clear that there is no such thing as a perfect heuristic, that each has blind

spots or biases. We can also see what the “normal” situations are in which the heuristics can be said to be the best, and ask whether the situation we are in is abnormal, and perhaps requires a differently designed decision procedure which implies a different collective heuristic. There is no general solution to the problem of whether the situation in which the heuristic is applied is normal. But that is not the point. We will at least have a vocabulary in which to ask these questions, and ask them about historical cases which are similar, as well as to think about problems of institutional design. This is a failing of much present discussion of expertise and liberal democracy, which is concerned instead with the question of whether expertise is genuine.

And it is a failing of social epistemology to ignore the epistemic dimension of “the social,” the fact that much of the content of our social relations with others involves epistemic weightings—indeed, it is hard to see anything in our social relations that does not involve changes in the weighting of our own beliefs on the basis of the actions and beliefs of others. Thinking in terms of double heuristics compels us to think about collective decision procedures in terms of the same problems of bias, selectivity, and so forth that characterize the individual knowledge related activities of which collective activity is composed.

References

- Beck, Ulrich. 1995. *Ecological Enlightenment: Essays on the Politics of the Risk Society*, trans. M. Ritter. Atlantic Highlands, NJ: Humanities Press International, Inc.
- Boettke, Peter. 2010. *The Political Economy of Soviet Socialism: The Formative Years, 1918-1928*. Norwell, MA: Kluwer.
- Jacobs, Struan. 2000. “Critical Review of Michael Polanyi and Spontaneous Order,” *International Social and Political Philosophy* 3(4): 49-67.
- Jacobs, Struan. 1997-98. “Michael Polanyi and Spontaneous Order 1941-1951,” *Tradition and Discovery* XXIV(2): 14-27.
- List, Christian and Philip Pettit. 2002. “Aggregating Sets of Judgments: An Impossibility Result,” *Economics and Philosophy* 18: 89-110.
- Pettit, Philip, 2006a, “When to Defer to Majority Testimony – and When Not,” *Analysis* 66(3): 179-87.
- Pettit, Philip, 2006b, “No Testimonial Route to Consensus,” *Episteme: A Journal of Social Epistemology* 3(3): 156-165.
- Polanyi, Michael. [1951] 1980. *Logic of Liberty: Reflections and Rejoinders*. Chicago: The University of Chicago Press.
- Polanyi, Michael. 1962. “The Republic of Science,” *Minerva* 1: 54–73.
- Prychitko, David. 2002. *Markets, Planning, and Democracy: Essays After the Collapse of Communism*. Northampton, MA: Edward Elgar.
- Roberts, Paul Craig. 2005. Polanyi the Economist. In Struan Jacobs and Richard Allen (eds.) *Emotion, Tradition, Reason: Essays on the Social, Economic and Political Thought of Michael Polanyi*. Aldershot, UK: Ashgate, pp. 127-32.
- Solomon, Miriam. 2006a. “Groupthink versus the Wisdom of Crowds: The Social Epistemology of Deliberation and Dissent,” *The Southern Journal of Philosophy* 44: 28-43.
- Solomon, Miriam. 2006b. “Norms of Epistemic Diversity,” *Episteme: A Journal of Social Epistemology* 3(1-2): 23-36.
- Tversky, Amos and Kahneman, Daniel. 1974. “Judgment under Uncertainty: Heuristics and Biases,” *Science* 185(27 September): 1124-31.
- Tversky, Amos and Kahneman, Daniel. 1981. “The Framing of Decisions and the Psychology of Choice,” *Science* 211(30 January): 453-58.

Whewell, William. 1858. *Novum Organon Renovatum: Being the Second Part of the Philosophy of the Inductive Sciences* 3rd edn. London: John W. Parker and Son.

Wynne, Brian. 1996. May the Sheep Safely Graze? A Reflexive View of the Expert–Lay Knowledge Divide. In S. Lash, B. Szerszynski and B. Wynne (eds.), *Risk, Environment, Modernity: Towards a New Ecology*. London: Sage, pp. 27–43.