



CHICAGO JOURNALS



The Independence Thesis: When Individual and Social Epistemology Diverge

Author(s): Conor Mayo-Wilson, Kevin J. S. Zollman, and David Danks

Reviewed work(s):

Source: *Philosophy of Science*, Vol. 78, No. 4 (October <sc>2011</sc>), pp. 653-677

Published by: [The University of Chicago Press](#) on behalf of the [Philosophy of Science Association](#)

Stable URL: <http://www.jstor.org/stable/10.1086/661777>

Accessed: 04/05/2012 14:20

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at

<http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The University of Chicago Press and Philosophy of Science Association are collaborating with JSTOR to digitize, preserve and extend access to *Philosophy of Science*.

<http://www.jstor.org>

The Independence Thesis: When Individual and Social Epistemology Diverge*

Conor Mayo-Wilson, Kevin J. S. Zollman, and David Danks^{†‡}

Several philosophers of science have argued that epistemically rational individuals might form epistemically irrational groups and that, conversely, rational groups might be composed of irrational individuals. We call the conjunction of these two claims the Independence Thesis, as they entail that methodological prescriptions for scientific communities and those for individual scientists are logically independent. We defend the inconsistency thesis by characterizing four criteria for epistemic rationality and then proving that, under said criteria, individuals will be judged rational when groups are not and vice versa. We then explain the implications of our results for descriptive history of science and normative epistemology.

Philosophers and social scientists have often argued (both implicitly and explicitly) that rational individuals can form irrational groups and that,

*Received June 2010; revised January 2011.

†To contact the authors, please write to Conor Mayo-Wilson: Department of Philosophy, Baker Hall 135, Carnegie Mellon University, Pittsburgh, PA 15213; e-mail: conormw@andrew.cmu.edu.

‡The authors would like to thank three anonymous referees and audiences at Logic and the Foundations of Game and Decision Theory 2010; Logic, Reasoning, and Rationality 2010; the London School of Economics; and the University of Tilburg for their helpful comments. Conor Mayo-Wilson and Kevin Zollman were supported by the National Science Foundation grant SES 1026586. David Danks was partially supported by a James S. McDonnell Foundation Scholar Award. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or the James S. McDonnell Foundation.

Philosophy of Science, 78 (October 2011) pp. 653–677. 0031-8248/2011/7804-0007\$10.00
Copyright 2011 by the Philosophy of Science Association. All rights reserved.

conversely, rational groups might be composed of irrational individuals.¹ We call the conjunction of these two claims the *Independence Thesis*, as together they imply that prescriptions for individual and group decision making, respectively, are logically independent of each other.

In the context of science, the Independence Thesis is the assertion that methodological prescriptions for scientific communities and those for individual scientists are logically independent,² and in recent years, this thesis has been defended in various forms. Individuals who desire credit more than truth may divide “cognitive labor” across competing research programs better than do truth-seeking scientists, thereby inadvertently improving the scientific community’s chances of discovering truths (Goldman 1992; Kitcher 1993; Strevens 2003). Individuals who refuse to abandon their favored theory, even in light of strong evidence against it, may help to make sure good theories are not prematurely abandoned by the broader scientific community (Feyerabend 1965, 1968; Popper 1975; Kuhn 1977; Hull 1988; Zollman 2010). Groups composed of a random assortment of problem solvers may outperform a group of the individually best problem solvers (Hong and Page 2001, 2004). Unreliable individuals might pool their information in such a way as to create reliable groups (Surowiecki 2004; Goodin 2006). Groups in which a significant amount of information is ignored might do better than groups in which information flows freely because an appropriate amount of diversity can be maintained in the absence of information (Zollman 2007, 2010). And so on. The underlying spirit of the Independence Thesis has motivated the creation of a new field—social epistemology—which focuses explicitly on the epistemic properties of groups rather than individuals (Goldman 1999).

Despite the recognition that individual and group epistemic rationality might differ, discussions of scientific methodology (and of epistemology generally) often depict the scientist as studying the world in isolation, rather than in a community of other researchers. Moreover, philosophers of science continue to draw methodological prescriptions for scientific practice in light of such idealized models. For example, in Bayesian models

1. Here, we are thinking of a host of economic phenomena. In game theory, the prisoner’s dilemma illustrates that rational agents might nonetheless act so as to necessitate an outcome that is Pareto dominated by some other outcome; i.e., the alternative outcome is strictly preferred by all agents. The free-rider problem and the tragedy of the commons are famous examples of such multiple-person prisoner’s dilemmas. In social choice theory, Arrow’s theorem asserts that it may be impossible to form a rational “group preference,” even when each agent in the group has a rational set of preferences. We thank an anonymous referee for suggesting these comparisons to our results.

2. A similar claim is called the *autonomy thesis* by Bishop (2005). As there are some differences between the two theses, we will use a different name.

of science, researchers are typically conceived as responding to evidence in isolation. Yet the real-world scientist is always part of some larger research community. Hence, philosophers who employ Bayesian models to draw prescriptions for working scientists often implicitly assume the falsehood of the Independence Thesis: they uncritically conclude that each member of a scientific community ought to adopt Bayesian methodology from the assumption that Bayesianism is rational for an isolated individual.³ Similar remarks can be made for formal learning theory, belief revision, ranking theory, and a host of other inductive methods recommended by philosophers of science and epistemologists.

If true, the Independence Thesis suggests that the methods prescribed by Bayesianism, belief-revision theories, formal-learning theory, and so on, might be correct for an isolated scientist but that this correctness need not extend to a community of scientists. Reliable, correct inference for a community of scientists might depend crucially on the internal organization of that community.

In this article, we (i) develop a model of communal scientific inquiry; (ii) describe several applications of the model to scientific practice, including an extended application to modeling theory choice in psychology; and (iii) prove several theorems that characterize when four different criteria for individual and group epistemic rationality converge and diverge. In addition to generalizing existing models of scientific communities,⁴ the model we develop and the theorems we prove contribute to understanding the relationship between individual and social epistemology in at least four novel ways.

First, our arguments make precise several possible formulations of the Independence Thesis and show that some formulations are true while others are not. In other words, there is no single Independence Thesis concerning the relationship between individual and social epistemology but rather a myriad of claims—some of which are true, and others not.

3. Three caveats are necessary. First, some regard Bayesianism only as a description of scientific practice. Our arguments address only those philosophers who draw prescriptive consequences from Bayesian models. Second, while some have attempted to analyze how Bayesians should respond to evidence from others (e.g., Bovens and Hartmann 2003), the focus of Bayesian philosophy remains the performance of an individual inquirer, not the performance of a group of inquirers. Finally, the propriety of groups of Bayesians may, in fact, depend on properties of individual Bayesians in isolation, but in light of the Independence Thesis, one is obliged to give an argument about why those properties “scale up.” One cannot simply presume that they will.

4. Our model of scientific practice generalizes some features of models already extant in the philosophical literature, but it also differs in several important respects (Kitcher 1990, 1993, 2002; Weisberg and Muldoon 2009). Space prevents a detailed comparison of these models.

In discussing the relationship between individual and social epistemology, therefore, one must precisely characterize the criteria by which individual or group epistemic quality is judged.

Second, although economists and other social scientists frequently argue that rational individuals might compose irrational groups, it is less clear that rational groups might consist predominantly of irrational members. Two of our theorems suggest this more radical conclusion: prescriptions for scientific communities might permit (or even require) every individual researcher to adopt a method that would be fundamentally irrational for a scientist learning in isolation.

Third, we consider different inductive and statistical methods that have been advocated in various scientific disciplines.⁵ Beyond simply asserting the existence of methods that are good from an individual perspective and bad from a group perspective (or vice versa), we point to particular methods that have been put forward as good for scientific inference. Thus, we simultaneously evaluate those methods along with generating more abstract claims about the relationship between group and individual judgments of rationality.

Finally, the formulations of the Independence Thesis that we defend have important descriptive and normative implications for history and philosophy of science. The descriptive implication is that history and sociology of science can proceed neither by focusing exclusively on individual researchers nor by focusing exclusively on aggregate properties of scientific communities. Rather, both levels must be considered in understanding how groups come to further scientific knowledge.⁶ For example, although Aristotle's scientific acumen was unparalleled during his time, one cannot uncritically conclude that his methods, if employed by many scientists for centuries, would prove fruitful for discovery. Conversely, knowing that Bourbaki discovered an immense number of fruitful theorems does not tell us, after all, about the rationality of any particular individual mathematician in the group. An accurate historical record of science ought to incorporate both detailed descriptions of the achievements of individual scientists and also a social history of the relevant scientific communities and institutions, including an analysis of how learning methods are shared and research results are communicated.

The normative implications of our arguments are also important. A

5. Epsilon greedy methods are described by Sutton and Barto (1998), Roth-Erev reinforcement learning is analyzed by Beggs (2005), and Bayesian methods are analyzed by Berry and Fristedt (1985).

6. In this way, our argument bolsters one of the central claims of Philip Kitcher (1993), wherein he argues against the move from the "irrationality" of individual scientists to the "irrationality" of science as a whole.

normative theory that evaluates scientific methodology by focusing on a single scientist will not necessarily “scale up” to a theory about groups. Similarly, a normative theory that only considers properties of groups should not necessarily require that each individual adopt a methodology that would necessarily succeed in isolation.

For these reasons, the Independence Thesis supports a separation between social epistemology and more traditional individualistic epistemology, at least with regard to certain standards of reliability. Moreover, these results make precise the intuitions that others have had about the context sensitivity of the connections between individual and group rationality. For example, we now have a clear picture of how dogmatic adherence to a particular theory can stymie science in one context and yet promote appropriate epistemic diversity in another (Feyerabend 1965, 1968; Popper 1975; Kuhn 1977; Hull 1988; Zollman 2010).

1. A Model of Scientific Inquiry. To distinguish between and make precise several formulations of the Independence Thesis, we introduce the following idealized model of scientific inquiry.⁷ We describe features of our model fairly informally; technical details can be found in Mayo-Wilson, Zollman, and Danks (2010). In our model, there is a finite collection of scientists, whom we will also call *learners*. Each scientist has access to the research being conducted by a set of his or her peers, whom we will call *neighbors*. A scientist’s neighborhood can be understood as her closest colleagues: those researchers with whom she communicates and exchanges information, papers, draft papers, and so on. We assume that the sharing of information is symmetric, in the sense that if scientist 1 knows about scientist 2’s research, then 2 knows about 1’s research.

We represent the relationships in a scientific community using undirected graphs like the one in figure 1. We will often refer to such a community as a *research network* or simply a *network*. For simplicity, we assume that the network structure is fixed for the duration of inquiry, although real-world scientific communities are of course dynamic and

7. In our model, each individual scientist is confronted with the same “bandit problem.” In other words, what we call a learning problem below is more frequently called a bandit problem in economics and psychology. See Berry and Fristedt (1985) for a survey of bandit problems. The difference between our model of scientific inquiry and standard bandit problems is that our model captures the social dimension of learning. We assume that, as they learn, scientists are permitted to share their findings with others. Thus, our model is identical to the first model of communal learning described in Bala and Goyal (2011). Importantly, many of our assumptions are either modified or dropped entirely by authors who have developed similar models, and we urge the reader to consult Bala and Goyal (2011) for a discussion of how such modifications might affect our results.



Figure 1. Research network (*left*) and neighborhood of that same network (*right*). Dark node = a scientist; gray nodes = neighborhood.

change over time. Finally, we assume that our research networks are *connected*: there is an “informational path” (i.e., a sequence of edges) connecting any two scientists. An informational path represents a chain of scientists who each communicate with the next in the chain. In this model, one learns about the research results of one’s neighbors only; individuals who are more distant in the network may influence each other but only by influencing the individuals between them.⁸ We restrict ourselves to connected networks because unconnected research networks do not adequately capture scientific communities but rather several separate communities. Examples of a connected and an unconnected network are depicted in figure 2.

At each stage of inquiry, each scientist may choose to perform one of finitely many *actions*, such as conducting an experiment, running a simulation, making numerical calculations, and so on. We assume that the set of actions is constant through inquiry, and each action results (probabilistically) in an *outcome*. An outcome might represent the recording of data or the observation of a novel phenomenon. Scientific outcomes can usually be regarded as better or worse than alternatives. For example, an experiment might simply fail without providing any meaningful results. An experiment might succeed but give an outcome that is only slightly useful for understanding the world. Or it might yield an important result. In order to capture the “epistemic utility” of different scientific outcomes, we will represent them as real numbers, with higher numbers representing better scientific outcomes.⁹

8. Our model does not allow the sharing of secondhand information, which is undoubtedly an idealization in this context. However, we do not think it is likely that a more realistic model that included such a possibility would render our conclusions false, and such a model would be significantly more complex to analyze.

9. For technical reasons, we assume outcomes are nonnegative and that the set of outcomes is countable. Moreover, although we assume (for presentation purposes) that the number of actions and scientists is finite, it is not necessary to do so. See Mayo-Wilson et al. (2010) for conditions under which the number of actions and agents might be considered countably infinite.



Figure 2. Connected network (*left*) and unconnected network (*right*).

There is a set of possible *states of the world*, and the probabilities of outcomes (after an action) can depend on the state of the world. This captures the idea that many outcomes (e.g., experimental observations) depend on the world being a certain way. For example, the outcome of the detection of the bending of light during an eclipse depends on the state of the world—specifically, whether general relativity is true or false. Of course, detection is only possible if one sets up particular measuring apparatuses and so forth; some action must be taken by the scientist. We assume the probability of some outcome given an action and world is constant through time. Such an assumption may not be realistic. For example, a *discovery* is an outcome that can only happen once, regardless of how many times the scientist takes the relevant action.

Throughout, we will consider a running example of theory choice in cognitive psychology. Specifically, consider the problem of determining the nature of human concepts. There is a range of theories of concepts, including exemplar-based theory (Medin and Schaffer 1978; Nosofsky 1984), prototype-based theory (Smith and Minda 1998; Minda and Smith 2001), causal-model theory (Rehder 2003a, 2003b), theory theory (Carey 1985; Gopnik and Meltzoff 1997), and others. The states of the world in this example correspond to the actual cognitive representations that people have. Actions in the model correspond to a scientist acting to verify a particular theory: conducting experiments, performing mathematical derivations, and so forth.¹⁰ Note that a scientist's action in this model is not some psychological act or state but rather is an observable action that leads to an observable outcome. (We return to this point below.) Outcomes in the model are the importance, quality, and nature of the resulting research products. Clearly, outcomes depend on both what cognitive representations are actually in people's heads (i.e., what our concepts really are) and also the type of investigation conducted. At the same time, outcomes are not a deterministic function of cognitive representations and

10. Because of the “first past the post” incentive system in science, an action cannot correspond to a specific experiment since (as noted earlier) the value of the outcome (i.e., the discovery) is not stable over time.

the type of investigation, as many other factors can influence the research products (e.g., experimental design, measurement error).

Returning to the abstract model of scientific inquiry, we say that a *learning problem* is a quadruple $\langle \Omega, A, O, p \rangle$, where Ω is a set of states of the world, A is a set of actions, O is a set of outcomes, and p is a measure specifying the probability of obtaining a particular outcome given an action and state of the world. In general, at any point of inquiry, every scientist has observed the finite sequence of actions and outcomes performed by herself, as well as the actions and outcomes of her neighbors. We call such a finite sequence of actions and their resulting outcomes a *history*. A *method* (also called a *strategy*) m for an individual scientist is a function that specifies, for any individual history, a probability distribution over possible actions for the next stage. In other words, a method specifies probabilities over the scientist's actions given what she knows about her own and her neighbors' past actions and outcomes. Of course, an agent may act deterministically, simply by placing unit probability on a single action $a \in A$. A *strategic network* is a pair $S = \langle G, M \rangle$ consisting of a network G and a sequence $M = \langle m_g \rangle_{g \in G}$ specifying the strategy employed by each learner, m_g , in the network.

To return to our running example, a particular psychologist directly knows about the research products of only some of the other psychologists working on the nature of concepts; those individuals will be her graphical neighbors. The experiments and theories considered and tested in the past by the psychologist and her neighbors are known to the psychologist, and she can use this information in her learning method or strategy to determine (probabilistically) which cognitive theory to pursue next. Importantly, we do not constrain the possible strategies in any significant way. The psychologist is permitted, for example, to use a strategy that says "do the same action (i.e., attempt to verify the same psychological theory) for all time" or to have a bias against changing theories (since such a change can incur significant intellectual and material costs).

Although we will continue to focus on this running example, it is important to note that the general model can be realized in many other ways. Here are two further examples of how the model can be interpreted:

Medical Research: The scientists are medical researchers or doctors who are testing various potential treatments for a disease. An action represents the application of a single treatment to a patient, and the outcome represents the degree to which the treatment succeeds for that patient (including the evaluation of side effects). The set of worlds represents the states in which different treatments are superior, all things considered. At each stage of inquiry, a scientist administers a treatment to a particular patient, observes the outcome, and then

chooses how to treat future patients given the success rates of the treatments she has administered in the past. Her choice of treatment, of course, is also informed by which treatments other doctors and researchers have successfully employed; those other researchers are the neighbors of the scientist in our model.

Scientific Modeling: When attempting to understand some phenomenon, there are often a variety of potential methods that could be illuminating. For instance, a biologist wanting to understand some odd animal behavior in the wild could turn to field observation, laboratory experiments, population genetic models, game theoretic models, and so on. Under this interpretation, each action represents an attempt by a scientist to apply a particular method to a given domain. The outcome represents the degree to which the scientist succeeds or fails, and the state of the world represents the state in which a particular method is more or less likely to provide genuine understanding of the problem. Our running example of cognitive modeling is, therefore, an instance of one of many interpretations of the model.

Because each outcome is assigned an epistemic utility, we can speak of a particular action as being optimal in a given state of the world. An action is optimal in a state if its expected utility in that state is at least as high as that of any other action. For example, it might be that exemplar-based theories are the best theories for understanding concepts, and so the action of attempting to verify an exemplar-based theory will be optimal.¹¹ Focusing on other theories might yield useful results occasionally, but on average they would (in this world) be worse.

Some learning problems are easier to solve than others; for example, if one action always dominates all others regardless of the state of the world, then there is relatively little for scientists to learn. This is rarely the case in science since the optimal actions typically depend on the state of the world. We are principally interested in such difficult problems. More precisely, recall that outcomes can be interpreted as utilities. Thus, for any state of the world ω , there is an expected value for each action that is constant throughout time. Hence, in any state of the world ω , there is some collection of optimal actions that maximize expected utility. Say a learning problem poses the problem of induction if no finite history reveals that a given action is optimal with certainty. In other words, the problem

11. A theory might be “best” because it is true or accurate or unified, etc. Our model is applicable, regardless of which of these standards one takes as appropriate for evaluating scientific theories.

of induction (in this context) is that, for any finite history, there are at least two states of the world in which such a history is possible, but the sets of optimal actions in those two states are disjoint. Say a learning problem is difficult if it poses the problem of induction, and for every state of the world and every action, the probability that the action yields no payoff on a given round lies strictly between zero and one. That is, no action is guaranteed to succeed or fail, and no history determines an optimal action with certainty. For the remainder of the article, we assume all learning problems pose the problem of induction and will note whether a problem is also assumed to be difficult.

2. Individual versus Group Rationality. Undoubtedly, one goal of scientific inquiry is to eventually find the “best” theories in a given domain. In our model, “eventually finding the best theories” corresponds to performing optimal actions with probability approaching one in the infinite limit—this property is generally known as *statistical consistency*. In this section, we investigate several versions of the Independence Thesis that assert that individuals employing statistically consistent methods might not form consistent groups and that consistent groups need not contain consistent individuals. We find that, depending on exactly how the notion of consistency is construed, individual and group epistemic quality may either coincide or diverge.

Two questions are immediate. Why do we focus on statistical consistency exclusively, rather than considering other methodological virtues? And why is consistency characterized in terms of performing optimal actions, rather than holding true beliefs about the state of the world (and, hence, of which actions are in fact optimal)?

In response to the first question, we argue that statistical consistency is often the closest approximation to reliability in empirical inquiry, and therefore, employing a consistent method is a necessary condition for attaining scientific justification and knowledge. In the context of science, reliabilism is the thesis that an inductive method confers a scientist with justification for believing a theory if and only if the method, in general, tends to promote true beliefs. Unfortunately, even in the simplest statistical problems, there are no inductive methods that are guaranteed (with high probability) to (i) promote true rather than “approximately true” beliefs and (ii) promote even approximately true beliefs given limited amounts of data.

Suppose, for example, you are given a coin, which may be fair or unfair, and you are required to determine the coin’s bias, that is, the frequency with which the coin lands heads. Suppose you flip the coin $n_1 + n_2$ times and observe n_1 heads. An intuitive method, called the “straight rule” by Reichenbach, is to conjecture that the probability that the coin will land

heads is $n_1/(n_1 + n_2)$. Suppose that, in fact, the coin is fair. With what probability, does the straight rule generate true beliefs? If you flipped the coin an odd number of times, it is impossible for $n_1/(n_1 + n_2)$ to be $1/2$, and so the probability that your belief is correct is exactly zero. But it is worse than that. As the number of flips increases, the probability that the straight rule conjectures exactly $1/2$ approaches zero. This argument does not show a failing with the straight rule: it applies equally to any method for determining the coin's bias.

Therefore, reliability, in even the simplest settings, cannot be construed as "tending to promote true beliefs," if "tending" is understood to mean "with high probability." Rather, the best our inductive methods can guarantee is that our estimates approach the truth with increasing probability as evidence accumulates. But this is just what it means for a method to be statistically consistent.

However, one still might be worried that we have construed statistical consistency in terms of performing optimal actions rather than possessing true beliefs. Our focus is thus different from the (arguably) standard one in philosophy of science: we consider convergence of action, not convergence of belief. We are not attempting to model the psychological states of scientists but rather focusing on changes in their patterns of activity.¹² This is why our running example from cognitive psychology focuses on the actions that a scientist can take to verify a particular theory, rather than the scientist accepting (or believing or entertaining or so on) that theory. Of course, to the extent that scientists are careful, thoughtful, honest, and so forth, their actions should track their beliefs in important ways. But our analysis makes no assumptions about those internal psychological states or traits, nor do we attempt to model them.

Moreover, we think that there are compelling reasons to focus on convergence in observable actions rather than convergence in a scientist's beliefs. Suppose only finitely many actions a_1, a_2, \dots, a_n are available and that a scientist cycles through each possible action in succession $a_1, a_2, \dots, a_n, a_1, a_2$, and so on. As inquiry progresses, such a scientist will acquire increasingly accurate estimates of the value of every action (by the law of large numbers). Thus, the scientist will also know which actions are optimal, even though she chooses suboptimal as often as she choose optimal ones. So the scientist would converge in belief without that convergence ever leading to appropriate actions.

Consider for a moment how strange a scientific community that follows such a strategy would appear. Each researcher would continue to test all potential competing hypotheses in a given domain, regardless of their

12. Those action patterns are presumably generated by psychological states, but we are agnostic about the precise nature of that connection.

apparent empirical support. Physicists would occasionally perform experiments that relied on assumptions from Aristotelian physics, astronomers would continue developing geocentric models, geologists would use theories that place the age of the earth at 8,000 years, and so on. Such a practice would continue indefinitely, despite the increasing evidence of the inadequacy of such theories. We believe that commitment to a scientific theory (whether belief, acceptance, or whatever) ought to be reflected in action and, thus, that good scientific practice requires a convergence in not only the beliefs of the scientists but also the experimentation.¹³

More radically, one might even question what role belief would have for individuals who continue to rely on increasingly implausible theories. What exactly does it mean for a scientist to believe in a theory that she uses as often as its competitors? We do not wish to delve too deeply into the issue of the relationship between belief and action, except to suggest that our would-be detractors would have to devise a notion of acceptance that radically divorces it from action.

With this background in mind, we now consider one formulation of the Independence Thesis: Do optimal methods for an isolated researcher prove beneficial if employed by a community of scientists who share their findings? More precisely, suppose we have a fixed learning problem and consider the isolated strategic network consisting of exactly one scientist. We say that the scientist's method is *convergent in isolation* (IC) if, regardless of the state of the world, her chance (as determined by her method) of performing an optimal action (relative to the actual state of the world) approaches one as inquiry progresses.

In our running example, our concept-focused psychologist would be called IC if she, when working alone, would converge to always testing (and presumably publicly advocating) the optimal theory of concepts. For example, suppose she selects a focal theory at time n by the following method: the theory that has yielded the best outcomes in the past is tested with probability $n/(n+1)$, and a random theory is selected with probability $1/(n+1)$. The psychologist's method is an example of what are called *decreasing epsilon greedy* methods, and this particular decreasing epsilon greedy strategy is IC; the psychologist will provably converge toward testing the optimal theory with probability one. Interestingly, the method of always testing the best-up-to-now theory is not IC: testing the best-up-to-now theory can lead to abandoning potentially better theories early on, simply because of random chance successes (or failures) by suboptimal (or optimal) theories (Zollman 2007, 2010). The previous

13. Moreover, sometimes ethical constraints preclude the alternation strategy. A medical researcher cannot ethically continue to experiment with antiquated treatments indefinitely.

method is IC, precisely because our imaginary psychologist attempts to verify other theories sufficiently often to ensure that the best theory is not prematurely abandoned, while still acting to verify the best-up-to-now theory with increasing probability.

On first glance, convergence in isolation seems to be a minimal/necessary requirement for being a “good” method for learning. However, IC methods may have undesirable consequences when employed by a group of researchers. To see why, say that a collection of methods M is *convergent in an isolated group* (GIC) if whenever every scientist in a research network uses a method in M and every method in M is used by at least one scientist, then each scientist’s chance of performing an optimal action approaches one as inquiry progresses. That is, a GIC collection of methods is one in which every method finds at least one optimal action (in the limit) whenever the other methods in the collection are present in the network. Surprisingly, as the following two theorems state, there is no necessary relationship between IC and GIC learners.

Theorem 1. In any difficult learning problem, there exist IC methods m such that the (singleton) collection $M = \{m\}$ is not GIC.

One example is a variation on the decreasing epsilon greedy strategy described above.¹⁴

Theorem 2. In any learning problem that poses the problem of induction, there exist GIC collections of methods M such that no $m \in M$ is IC.

One example is a group of “preferred action” learners (described below).

The decreasing epsilon greedy strategy is an IC method that can lead to disastrous results when employed in groups. The psychologist’s method is IC because her probability of experimentation tapers to zero at the “right” speed. If that probability tapers too quickly (e.g., if the probability of randomly choosing a theory were $1/n^3$), then the method may fail to experiment frequently enough to find the best theory. In particular, if the psychologist’s rate of experimentation depends on the number of neighbors she has (e.g., a random sample rate of $1/n^k$, where k is the number of neighbors), then she will converge to testing an optimal theory in isolation but will not necessarily converge to an optimal theory in a group. Essentially, the potential problem here is one of groupthink: by herself, each psychologist patiently figures out the right theory; in a group, how-

14. Proofs sketches of all theorems are provided in the appendix. Full details are available in Mayo-Wilson et al. (2010).

ever, she might converge on a suboptimal theory too quickly and get stuck.¹⁵

As an example of the second theorem, consider the following collection of methods: for each action a , the method m_a chooses a with probability $1/n$ (where n is the stage of inquiry) and otherwise does a best-up-to-now action. One can think of a as the “preferred action” of the method m_a , although m_a is capable of learning. For our psychologist, a is the “favorite theory” of concepts. Any particular method m_a is not IC since it can easily be trapped in a suboptimal action (e.g., if the first instance of a is successful). The core problem with the m_a methods is that they do not experiment: they either do their favorite action or a best-up-to-now one. However, the set $M = \{m_a\}_{a \in A}$ is GIC, as when all such methods are employed in the network, at least one scientist favors an optimal action a^* and makes sure it gets a “fair hearing.” The neighbors of this scientist gradually learn that a^* is optimal, and so each neighbor plays either a^* or some other optimal action with greater frequency as inquiry progresses. Then the neighbors of the neighbors learn that a^* is optimal, and so on, so that knowledge of at least one optimal action propagates through the entire research network. In our running example, the community as a whole can learn, as long as each theory of concepts has at least one dedicated proponent, as the proponents of any given theory can learn from others’ research. Moreover, this arguably is an accurate description of much of scientific practice in cognitive psychology: individual scientists have a preferred theory that is their principal focus (and that they hope is correct), but they can eventually test and verify different theories if presented with compelling evidence in favor of them.

This latter example—and the associated theorem—supports a deep insight about scientific communities made by Popper and others: a diverse set of approaches to a scientific problem, combined with a healthy dose of rigidity in refusing to alter one’s approach unless the evidence to do so is strong, can benefit a scientific community even when such rigidity would prove counterproductive to a scientist working in isolation. The above argument also bolsters observations about the benefits of the “di-

15. This particular method might seem strange on first reading. Why would a rational individual adjust the degree to which she explores the scientific landscape on the basis of the number of other people she listens to? However, one can equivalently define this strategy as dependent, not on the number of neighbors but on the amount of evidence. Suppose a scientist adopts a random sample rate of $1/n^{x/y}$, where x is the total number of experiments the scientist observes and y is the total number of experiments she has performed. Our scientist is merely conditioning her experimentation rate on the basis of the amount of acquired evidence—not an intrinsically unreasonable thing to do. However, in our model, this method is mathematically equivalent to adopting an experimentation rate of $1/n^k$.

vision of cognitive labor” that have been advocated by Kitcher, Strevens, Weisberg, and Muldoon, among others.

The definitions of IC and GIC focus on learners who have some measure of control over their research network. Although there are numerous examples and case studies of scientific communities trying to maintain control over their membership, most actual scientific communities are not so “insular.” We must also consider cases in which a learner or collection of learners is embedded in a larger community that could potentially contain individuals who have different epistemic goals, standards, or methods. Say a method is *universally convergent* (UC) if in any state of the world, the researcher chooses optimal actions with probability approaching one as inquiry progresses, regardless of the network in which the researcher is embedded. Similarly, say a set of methods M is *group universally convergent* (GUC) if for all networks such that every method in M is employed at least once and everyone employing methods in M is connected via “informational paths” consisting only of scientists also employing methods from M , then each researcher employing some method in M chooses optimal actions with probability approaching one as inquiry progresses.

UC and GUC methods are those that are, in a specific sense, resistant to the other individuals in the network. A UC method converges to an optimal action, for example, even when it is surrounded by methods that produce suboptimal actions for all eternity. Notice that, in all learning problems, every UC method is IC, and every GUC collection is GIC. Why? Isolation is simply a special type of network structure, and since UC methods and GUC collections converge regardless of network structure, they must converge for these special cases. As one might suspect, UC and GUC are strictly stronger epistemic standards than IC or GIC.

Theorem 3. In all learning problems that pose the problem of induction and in which payoffs are bounded from below and above by positive real numbers, there exist IC methods that are not UC. There also exist GIC collections that are not GUC.

Perhaps the most commonly studied strategies witnessing the above theorem are called *reinforcement learning* (RL) strategies. Reinforcement learners begin with an initial, positive, real-valued weight for each action. On the first stage of inquiry, the agent chooses an action in proportion to the weights. For example, if there are two actions a_1 and a_2 with weights 3 and 5, respectively, then the agent chooses action a_1 with probability $3/(3 + 5)$ and a_2 with probability $5/(3 + 5)$. At subsequent stages, the agent then adds the observed outcome for all the actions taken in his neighborhood to the respective weights for the different actions. RL has




Figure 3.

been used as a descriptive model of learning opponents' behavior in games and various cognitive learning processes.¹⁶ It might also be recommended as a normative model for induction, as it is consistent when employed in isolation (Beggs 2005). RL strategies are IC but not UC, and groups of such learners can be GIC but not GUC.

Interestingly, however, the disconnect between group and individual epistemic norms found in isolation does not carry over in the same way to universal convergence, as shown by the following two theorems. In particular, there is a connection, but only in one direction.

Theorem 4. In all learning problems, every collection containing only UC methods is GUC.

Theorem 5. In any learning problem that poses the problem of induction, there exist GUC collections M such that every $m \in M$ is not UC.

One example is a set of “preferred action” methods (described above), which is a GUC collection of not-UC (not-IC) methods.

The relationships between IC, GIC, UC, and GUC can thus be neatly summarized by figure 3. The above five theorems provide qualified support for certain formulations of the Independence Thesis, as they show that, when epistemic quality is understood in terms of statistical consistency, there are circumstances under which individual and group rationality diverge.

Of course, one might object that our arguments rely on one particular formal model of inquiry that makes several unrealistic idealizations: network structure is constant, informational access is symmetric, and so on. One might continue that we have considered only stringent criteria of success and that there are any number of other standards by which one might evaluate epistemic performance of individuals and groups. Such an

16. This type of RL is different, though, from that studied in the animal behavior literature (e.g., RL models of classical conditioning).

objector might conclude that the above theorems say little about the relationship between individual and social epistemology in the “real world.”

First, note that a similar divergence between individual and group epistemic performance (measured in terms of statistical consistency) must emerge in any model of inquiry that is as complex as our model (or more so). For any model of inquiry capable of representing the complexities captured by our model, we would be able to define all of the methods considered above (perhaps in different terms), define appropriate notions of consistency, and so on. Thus, each of the above theorems would have an appropriate “translation” or analog in any model of inquiry that is sufficiently complex. In this sense, the simplicity of our model is a virtue, not a vice.

Moreover, we agree that our model is not the single “correct” representation of scientific communities. In fact, we would argue that no formal model of inquiry is superior to all others in every respect. Clearly, there are other criteria of epistemic quality that ought to be investigated, other models of particular scientific learning in communities, and so forth. However, significant scientific learning problems can be accurately captured and modeled in this framework, and the general moral of the framework should hold true in a range of settings, namely, that successful science depends on using methods that are sensitive to evidence and can learn, while avoiding community groupthink.

3. Conclusion. In this article, we have investigated several versions of the Independence Thesis that depend on different underlying notions of individual rationality (IC and UC) and group rationality (GIC and GUC). We have found some qualified support for the Independence Thesis; IC is independent of GIC. However, considering a stronger notion of rationality, we found that UC is not independent of GUC, although the entailment goes in one direction only.

Our formulations of these different notions of rationality illustrate important distinctions that have been previously overlooked. For example, the different properties of IC and UC (GIC and GUC, respectively) methods illustrate that, when considering individual (or group) rationality, one must consider how robust a method (or set of methods) is to the presence of others employing different methods. Subtle differences in definitions of individual and group rationality, therefore, can color how one judges the Independence Thesis.

In addition to illustrating the importance of this distinction, we have provided some formal support to those philosophers who have pushed the Independence Thesis (primarily on the basis of historical evidence). It is our hope that this proof will demonstrate the thesis, in a mathe-

matically rigorous way, and thus illustrate the limitations of discussions of scientific method that consider only the properties of inductive methods when employed in isolation, as is common in many formal models of scientific inquiry (e.g., Bayesianism, belief revision, formal learning theory). Finally, our arguments provide additional support to those who argue that the empirical study of science must include both the individual and the group properties of scientific groups if the prospects for scientific progress are to be understood properly.

: Appendix

Proofs of Theorems. The proofs of the theorems employ several lemmas, which are stated in the next section.

Theorem 1. In any difficult learning problem, there exist IC methods m such that the (singleton) collection $M = \{m\}$ is not GIC.

Proof. Define m to be the following method. At stage n , an individual employing m plays the action with the highest estimated value with probability $(n^k - 1)/n^k$, where k is the number of the individual's neighbors. If there are several actions that have the highest estimated value, then m splits the probability $(n^k - 1)/n^k$ evenly among all such actions. The method m plays every other action with equal shares of the remaining probability of $1/n^k$. We first show that m is IC and then show that m is not GIC.

To show that m is IC, consider the isolated network with one learner g employing m , and pick any state of the world ω . Regardless of history, the method m , by definition, assigns at least probability $1/(|A| \times n)$ to each action (where $|A|$ is the number of actions) on the n th stage of inquiry, as g has exactly one neighbor (herself). Thus, the probability that g plays action a on the n th stage of inquiry, regardless of what g has done previously, is at least $1/(|A| \times n)$. It is easy to check that $\sum_{n=1}^{\infty} 1/(|A| \times n) = \infty$, and so by lemma 1, it follows that g plays every action infinitely often. By lemma 2, it follows that, for any action a , the individual g 's estimate of the value of the action a approaches the true value. Now, by definition, the method m plays actions with highest estimated value with probability $(n - 1)/n$ at stage n , and this ratio approaches one as n approaches infinity. By lemma 3, the probability that m plays a truly optimal action approaches one as n approaches infinity. As ω was chosen arbitrarily, it follows that m is IC.

Next, we show that m is not GIC. To do so, consider the network consisting of exactly two learners g_1 and g_2 , each of whom employs the strategy m , and each of whom is the other's neighbor. We show that g_1 fails to play optimal actions with probability approaching one

in ω , and by symmetry, the same proof works for g_2 . In fact, we show that, with some positive probability, g_1 plays suboptimal actions at every stage after h .

To do so, fix a state of the world ω . Because the learning problem is difficult, there is some history h such that (i) every truly optimal action $a \in A_\omega$ yields zero payoff along h , (ii) some suboptimal action yields positive payoff along h , and (iii) the history h has nonzero probability $p_\omega^S(h) > 0$ in ω . Let j be the length of the history h . For any natural number n , let $A_{g,n}$ be the set of actions that have the highest estimated value at stage n . By definition of the method m , the learner g plays actions from $A \setminus A_{g,n}$ with probability no greater than $1/n^2$, as g_1 has two neighbors, namely, g_2 and herself. By choice of h , therefore, m assigns the set of optimal actions probability no greater than $1/(j+1)^2$ at the stage $j+1$ if h occurs. Hence, the conditional probability that g_1 plays an optimal action at stage $j+1$ given h is no greater than $1/(j+1)^2$. Because g_1 and g_2 choose their actions independently of one another, the probability that neither chooses an optimal action at stage $j+1$ given h is no greater than $[1/(j+1)^2]^2$.

Similarly, if h occurs and no optimal action is played at stage $j+1$, then m assigns probability no greater than $1/(j+2)^2$ to optimal actions at stage $j+2$. So the conditional probability that g_1 plays an optimal action at stage $j+2$ given h and that no optimal action is played at stage $j+1$ is no greater than $1/(j+2)^2$, and since g_1 and g_2 choose their actions independently of one another, the probability that neither chooses an optimal action is $[1/(j+2)^2]^2$. And so on. So the probability that h occurs and that suboptimal actions are played from every stage onward by both g_1 and g_2 is given by

$$p_\omega^S(h) \times \prod_{n=j}^{\infty} \left(1 - \frac{1}{n^2}\right)^2,$$

and this quantity can be shown to be strictly positive by basic calculus. So m is not GIC. QED

Theorem 2. In any learning problem that poses the problem of induction, there exist GIC collections of methods M such that no $m \in M$ is IC.

Proof. In fact, we show a stronger result. We show that there are collections M of methods such that M is GUC, but no member of M is IC. In the body of the article, we described the strategy m_a with the following behavior. If a is the best-up-to-now action, then it is

played with probability one. If a is not a best-up-to-now action, then it is played with probability $1/n$, and the remaining probability of $(n - 1)/n$ is divided evenly among the best-up-to-now actions. We claim that, in difficult learning problems, the set $\{m_a\}_{a \in A}$ is GUC, but no member of M is IC.

It is easy to show that no method m_a is IC. Suppose there is one individual employing the strategy m_a in isolation. First notice that, on the first stage of inquiry, a is one of the best-up-to-now actions, as no action has resulted in a payoff. So the individual plays a . Because outcomes are, by assumption, nonnegative, the action a remains a best-up-to-now action on the second stage of inquiry, and so the individual plays a again with probability one. By induction, it follows that the individual plays a on every stage of inquiry, regardless of the outcomes she obtains. Since the learning problem poses the problem of induction, there is some state of the world in which a is not an optimal action. Hence, in that state of the world, the isolated individual employing m_a not only fails to converge to an optimal action, but, moreover, with probability one, she plays a sub-optimal action for eternity. So m_a is not IC.

Next we show that the set of methods $M = \{m_a\}_{a \in A}$ is GUC. Let S be any strategic network in which the methods $M = \{m_a\}_{a \in A}$ are all employed at least once and such that, for any two individuals g_1 and g_2 employing methods from M , there is some sequence of individuals i_1, i_2, \dots, i_n employing methods from M such that i_1 is g_1 's neighbor, i_2 is i_1 's neighbor, and so on, until i_n , who is g_2 's neighbor. Fix any state of the world ω , and let g be some individual in the network employing a method from M . We want to show that, with probability approaching one, g plays an optimal action in ω .

As the set of actions is finite, there is at least one action a that is optimal in ω , and by definition of M , there is at least one learner g_a employing m_a in the network S . By lemma 1, with probability one, the learner g_a plays the action a infinitely often. Hence, by lemma 2, g_a 's estimate of the value of the action a approaches the true value of the action a . By lemma 3 and the definition of the strategy m_a , it follows that g_a plays optimal actions with probability approaching one; in fact, she plays a with probability approaching one. By lemma 4, this means that g_a plays optimal actions infinitely often.

Now consider neighbors of g_a who employ methods from M . Since neighbors of g_a observe the action a being played infinitely often, their estimates of the value of a will likewise approach the true value (again by lemma 1). By lemma 3 and the definition of the strategies in M , it follows that neighbors of g_a also play optimal actions with probability approaching one in ω . By lemma 4, it follows that neigh-

bors of g_a who employ methods from M will, with probability one, play optimal actions infinitely often. Hence, neighbors of neighbors of g_a will have an estimate of at least one optimal action that approaches the true value of that action. And so on.

That is, one can repeat the argument any finite number of times to prove that, for example, neighbors of neighbors of neighbors of g_a who employ strategies in M play optimal actions with probability approaching one. In this way, the optimal behavior of the individual g_a propagates through the subnetwork of agents employing methods from M . Now, by assumption, there is some sequence of individuals i_1, i_2, \dots, i_n employing methods from M such that i_1 is g_a 's neighbor, i_2 is i_1 's neighbor, and so on, until i_n , who is g 's neighbor. So g must converge to playing optimal actions in the limit as desired. QED

Theorem 3. In all learning problems that pose the problem of induction and in which payoffs are bounded from below and above by positive real numbers, there exist IC methods that are not UC. There also exist GIC collections that are not GUC.

Proof. When payoffs are bounded from above and below by positive real numbers, then a slight modification of the proof of theorem 1 in Beggs (2005) yields that any set of RL methods is GIC. See Mayo-Wilson et al. (2010) for details.

So we limit ourselves to showing that no finite set of RL methods is GUC in a learning problem that poses the problem of induction and in which payoffs are bounded from below and above by k_2 and k_1 , respectively, where k_1 and k_2 are positive real numbers. Let M be a finite sequence of RL methods. It suffices to find (i) a strategic network $S = \langle G, N \rangle$ with a connected M subnetwork $S' = \langle G', M \rangle$, (ii) an individual, and (iii) a state of the world such that $\lim_{n \rightarrow \infty} p_\omega^S(h^A(n, g) \in A_\omega) \neq 1$, where $p_\omega^S(h^A(n, g) = a)$ is the probability that g plays action a on the n th stage of inquiry in state of the world ω .

To construct S , first take a sequence of learners of the same cardinality as M and place them in a singly connected row, so that the first is the neighbor to the second, the second is a neighbor to the first and third, the third is a neighbor to the second and fourth, and so on. Assign the first learner on the line to play the first strategy in M , the second to play the second, and so on. Denote the resulting strategic network by $S' = \langle G', M \rangle$; notice that S' is a connected M network.

Next, we augment S' to form a larger network S as follows. Find the least natural number $n \in \mathbb{N}$ such that $nk_1 > 3k_2$. Add n agents

to G' and add an edge from each of the n new agents to each old agent $g \in G'$. Call the resulting network G . Pick some action $a \in A$, and assign each new agent the strategy m_a , which plays the action a deterministically. Call the resulting strategic network S ; notice that S contains S' as a connected M subnetwork.

Let ω be a state of the world in which $a \notin A_\omega$ (such an a exists because the learning problem poses the problem of induction by assumption). By construction, regardless of history, g has at least n neighbors each playing the action a at any stage. By assumption, payoffs are bounded below by k_1 , and so it follows that the sum of the payoffs to the agents playing a in g 's neighborhood is at least nk_1 at every stage. In contrast, g has at most three neighbors playing any other action $a' \in A$. Since payoffs are bounded above by k_2 , the sum of payoffs to agents playing actions other than a in g 's neighborhood is at most $3k_2 < nk_1$. It follows that, in the limit, one-half is strictly less than the ratio of (i) the total utility accumulated by agents playing a in g 's neighborhood to (ii) the total utility accumulated by playing all actions. As g is a reinforcement learner, g , therefore, plays action $a^* \notin A_\omega$ with probability greater than one-half in the limit, and so g plays a suboptimal action with nonzero probability in the limit. QED

Theorem 4. In all learning problems, every collection containing only UC methods is GUC.

Proof. Since UC methods converge regardless of the network in which they are placed, they must each also converge when informationally connected with other UC methods. QED

Theorem 5. In any learning problem that poses the problem of induction, there exist GUC collections M such that every $m \in M$ is not UC.

Proof. This follows immediately from the proof of theorem 2, as the methods M constructed there are GUC, but none is IC (let alone UC). QED

Lemmas. We state the following lemmas without proof. The first is essentially an immediate consequence of the (second) Borel-Cantelli lemma, and the second is an immediate consequence of the strong law of large numbers. However, some tedious (but straightforward) measure-theoretic constructions are necessary to show that the events described below have

well-defined probabilities. These technical details are omitted for brevity, but for the interested reader, they are available in Mayo-Wilson et al. (2010).

Lemma 1. Fix a state of the world ω , a strategic network S , and a learner g in S . Let $p_\omega^S(h^A(n, g) = a)$ be the probability that g plays action a on the n th stage of inquiry in state of the world ω . Let E_n be the event that g does not play an action a at any stage before n . Suppose that, for all natural numbers k , $\sum_{n=k}^\infty p_\omega^S(h^A(n, g) = a | E_n) = \infty$. Then, with probability one in ω , the individual g plays the action a infinitely often.

Lemma 2. Fix a state of the world ω , a strategic network S , and a learner g in S . Suppose that, with probability one in ω , an action a is played infinitely often in individual g 's neighborhood. Then, with probability one, g 's estimate of the value of a at stage n approaches (as n approaches infinity) the true value of a .

Lemma 3. Fix a strategic network S , an individual g in that network, and a state of the world ω . Let m be the method employed by g . Let $p_n(m)$ represent the probability that m assigns to actions with highest estimated value at stage n , and suppose that $p_n(m)$ approaches one as n approaches infinity. Further, suppose that, with probability one, for every action a , the individual g 's estimate of the value of a at stage n approaches the true value of a in ω . Then the probability $p_\omega^S(h^A(n, g) \in A_\omega)$ that g plays a truly optimal action approaches one as n approaches infinity.

Lemma 4. Fix a state of the world ω , a strategic network S , and a learner g in S . Let $p_\omega^S(h^A(n, g) = a)$ be the probability that g plays action a on the n th stage of inquiry in state of the world ω . Suppose that $\lim_{n \rightarrow \infty} p_\omega^S(h^A(n, g) = a) = 1$. Then, with probability one in ω , g plays the action a infinitely often.

REFERENCES

- Bala, Venkatesh, and Sanjeev Goyal. 2011. "Learning in Networks." In *Handbook of Mathematical Economics*, ed. J. Benhabib, A. Bisin, and M. O. Jackson. Amsterdam: North-Holland.
- Beggs, Alan. 2005. "On the Convergence of Reinforcement Learning." *Journal of Economic Theory* 122:1–36.
- Berry, Donald A., and Bert Fristedt. 1985. *Bandit Problems: Sequential Allocation of Experiments*. London: Chapman & Hall.
- Bishop, Michael A. 2005. "The Autonomy of Social Epistemology." *Episteme* 2:65–78.

- Bovens, Luc, and Stephan Hartmann. 2003. *Bayesian Epistemology*. Oxford: Oxford University Press.
- Carey, Susan. 1985. *Conceptual Change in Childhood*. Cambridge, MA: MIT Press.
- Feyerabend, Paul. 1965. "Problems of Empiricism." In *Beyond the Edge of Certainty: Essays in Contemporary Science and Philosophy*, ed. Robert G. Colodny, 145–260. Englewood Cliffs, NJ: Prentice-Hall.
- . 1968. "How to Be a Good Empiricist: A Plea for Tolerance in Matters Epistemological." In *The Philosophy of Science: Oxford Readings in Philosophy*, ed. Peter Nidditch, 12–39. Oxford: Oxford University Press.
- Goldman, Alvin. 1992. *Liasons: Philosophy Meets the Cognitive Sciences*. Cambridge, MA: MIT Press.
- . 1999. *Knowledge in a Social World*. Oxford: Clarendon.
- Goodin, Robert E. 2006. "The Epistemic Benefits of Multiple Biased Observers." *Episteme* 3 (2): 166–74.
- Gopnik, Alison, and Andrew Meltzoff. 1997. *Words, Thoughts, and Theories*. Cambridge, MA: MIT Press.
- Hong, Lu, and Scott Page. 2001. "Problem Solving by Heterogeneous Agents." *Journal of Economic Theory* 97 (1): 123–63.
- . 2004. "Groups of Diverse Problem Solvers Can Outperform Groups of High-Ability Problem Solvers." *Proceedings of the National Academy of Sciences* 101 (46): 16385–89.
- Hull, David. 1988. *Science as a Process*. Chicago: University of Chicago Press.
- Kitcher, Philip. 1990. "The Division of Cognitive Labor." *Journal of Philosophy* 87 (1): 5–22.
- . 1993. *The Advancement of Science*. New York: Oxford University Press.
- . 2002. "Social Psychology and the Theory of Science." In *The Cognitive Basis of Science*, ed. Stephen Stich and Michael Siegal. Cambridge: Cambridge University Press.
- Kuhn, Thomas S. 1977. "Collective Belief and Scientific Change." In *The Essential Tension*, 320–39. Chicago: University of Chicago Press.
- Mayo-Wilson, Conor A., Kevin J. Zollman, and David Danks. 2010. "Wisdom of the Crowds vs. Groupthink: Learning in Groups and in Isolation." Technical Report 188, Department of Philosophy, Carnegie Mellon University.
- Medin, Douglas L., and Marguerite M. Schaffer. 1978. "Context Theory of Classification Learning." *Psychological Review* 85:207–38.
- Minda, John P., and J. David Smith. 2001. "Prototypes in Category Learning: The Effects of Category Size, Category Structure, and Stimulus Complexity." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 27:775–99.
- Nosofsky, Robert M. 1984. "Choice, Similarity, and the Context Theory of Classification." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 10:104–14.
- Popper, Karl. 1975. "The Rationality of Scientific Revolutions." In *Problems of Scientific Revolution: Progress and Obstacles to Progress*, ed. R. Harre. Oxford: Clarendon.
- Rehder, Bob. 2003a. "Categorization as Causal Reasoning." *Cognitive Science* 27:709–48.
- . 2003b. "A Causal-Model Theory of Conceptual Representation and Categorization." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 29:1141–59.
- Smith, J. David, and John P. Minda. 1998. "Prototypes in the Mist: The Early Epochs of Category Learning." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 24:1411–36.
- Strevens, Michael. 2003. "The Role of the Priority Rule in Science." *Journal of Philosophy* 100 (2): 55–79.
- Surowiecki, James. 2004. *The Wisdom of Crowds: Why the Many Are Smarter than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*. New York: Doubleday.
- Sutton, Richard S., and Andrew G. Barto. 1998. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.
- Weisberg, Michael, and Ryan Muldoon. 2009. "Epistemic Landscapes and the Division of Cognitive Labor." *Philosophy of Science* 76 (2): 225–52.

- Zollman, Kevin J. 2007. "The Communication Structure of Epistemic Communities." *Philosophy of Science* 74 (5): 574–87.
- . 2010. "The Epistemic Benefit of Transient Diversity." *Erkenntnis* 2 (1): 17–35.