

Naturalized metaphilosophy

David R. Morrow · Chris Alen Sula

Received: 6 January 2008 / Accepted: 30 November 2008
© Springer Science+Business Media B.V. 2009

Abstract Traditional representations of philosophy have tended to prize the role of reason in the discipline. These accounts focus exclusively on ideas and arguments as animating forces in the field. But anecdotal evidence and more rigorous sociological studies suggest there is more going on in philosophy. In this article, we present two hypotheses about social factors in the field: that social factors influence the development of philosophy, and that status and reputation—and thus social influence—will tend to be awarded to philosophers who offer rationally compelling arguments for their views. In order to test these hypotheses, we need a more comprehensive grasp on the field than traditional representations afford. In particular, we need more substantial data about various social connections between philosophers. This investigation belongs to a naturalized metaphilosophy, an empirical study of the discipline itself, and it offers prospects for a fuller and more reliable understanding of philosophy.

Keywords Philosophy · Sociology · Networks · Reason · Naturalism

In addressing questions of how to represent philosophy, we should begin—in good philosophical fashion—by first asking what philosophy *is*. Historical answers have often prized the role of reason in addressing questions of existence, knowledge, and value. Thus, Plato recommends contemplation of the forms, Descartes' meditator attempts to deduce self-evident truths from first principles, Spinoza discusses an activity carried out *sub specie aeternitatis* by a rational thinker who has removed all traces of subjectivity and individuality from his thought, and Kant describes philosophy as

D. R. Morrow · C. A. Sula (✉)
The Graduate Center, The City University of New York, 365 Fifth Avenue,
New York, NY 10016, USA
e-mail: csula@gc.cuny.edu

D. R. Morrow
e-mail: morrow@uchicago.edu

the system of principles of pure reason. These metaphilosophies have tended to prize the role of reason in giving an account of the field. But for any such account, a curious puzzle arises: nearly all philosophers are influenced by other philosophers around them, and influence is not necessarily rational, warranted, or wise.

Cohen (2000) provides a vivid illustration of the persistently *social* nature of philosophy in discussing the analytic/synthetic distinction in the twentieth century:

people of my generation who studied philosophy at Harvard rather than at Oxford for the most part reject the analytic/synthetic distinction. And I can't believe that this is an accident. That is, I can't believe that Harvard just happened to be a place where both its leading thinker rejected that distinction and its graduate students, for independent reasons—merely, for example, in the independent light of reason itself—also came to reject it. And vice versa, of course, for Oxford. I believe, rather, that in each case students were especially impressed by reasons respectively for and against believing in the distinction, because in each case the reasons came with the added persuasiveness of personal presentation, personal relationship, and so forth. (p. 18)

Anecdotes like Cohen's are common in the published record of philosophy, especially in more informal settings such as book introductions, autobiographies, interviews, and correspondence. These confessions suggest that philosophy does not proceed *purely* through "rational" reflection; the social dimensions of the field may be as important—perhaps even more important, in some cases—than the exercise of our "rational" faculties.

In order to understand how much, if at all, these social factors influence philosophy, we need a more comprehensive grasp on the field than traditional representations afford. Where traditional representations have focused almost exclusively on ideas and arguments, alternative representations might also consider what *causes* ideas and positions to emerge, *why some* gain philosophical currency over time and *others* lose credit, and *how* the commitments, questions, and boundaries of the field shift over time. Where traditional representations focus on *what* has been talked about in past and present philosophy, these alternatives might ask *why* or *how* those views have achieved or lost dominance. To answer these questions, we need more evidence about the practice of philosophy in the past and present. We need to develop and test hypotheses about the interaction between social and rational factors in the development of the field. This investigation belongs to *naturalized* metaphilosophy, an empirical study of the discipline itself.

1 Two models of representation

Questions of how to represent philosophy depend, in part, on what the content and limits of the field are taken to be. From a rationalistic perspective, representing the discipline is a fairly straightforward task. Good philosophy is a success story, a triumph of human capacity in ascertaining the rational order of the universe, including ideal states of social, political, and moral affairs. In turn, good representations of philosophy should contain ideas and arguments that have survived the test of

rational scrutiny. They should capture the cumulative knowledge we have obtained, and the more truths a representation contains, the better the representation. One could even argue there is some single *best* representation that reflects the most truths ascertained and endorses the least falsehoods. This austere representation could still mention “mistakes” in the history of philosophy *if* their errors or oversights would prove instructive or fruitful for further discoveries. But failures and lesser theories would be footnotes to a larger history of successful claims supported by decisive reasons.

Alas. We don’t all agree—and rarely have—on which claims are recommended by right reason. If anything, present and past philosophy looks more like a history of *disagreement*, rather than convergence on certain claims. Maybe there is no determinate rational order guiding our work, or maybe we simply lack enough access to that order to sort the true claims from the false ones in any widely accepted manner. In either case, more recent philosophers have tended to regard their overall work as the search for good *questions*, rather than answers. As Russell (1912/1997) says in his popular introduction to the field,

Philosophy is to be studied, not for the sake of any definite answers to its questions since no definite answers can, as a rule, be known to be true, but rather for the sake of the questions themselves; because these questions enlarge our conception of what is possible, enrich our intellectual imagination and diminish the dogmatic assurance which closes the mind against speculation... (p. 161)

Against this pluralistic backdrop of speculation, the best representations of philosophy (and we had better assume there can be more than one) should take a value-neutral approach that does not bias one (now-admittedly) contentious position over any other. Good representations should simply catalog the important ideas of philosophy in the broadest and most charitable sense.

This line of thought probably undergirds much of the encyclopedic efforts made by philosophers in the past few decades. Since Macmillan’s 1967 *Encyclopedia of Philosophy*, there has been a movement toward collectively authored volumes that treat individual topics in an allegedly balanced way. Contrary positions are presented with equal regard—or at least given their day in court for the benefit of novices, scholars of other (sub)disciplines, and experts alike. These audiences expect that in consulting the catalog, they are receiving an accurate, thorough, and undistorted representation of philosophy. But is such an expectation really fulfilled, and, to ask the prior question, is such an expectation even warranted?

If rationalistic representations suffer empirical defeat at the hands of disagreement, encyclopedic models face a parallel problem from the limits and biases of human cognition. Decisions about the content of encyclopedias and the exposition of their topics are the choices of scholars, who, incidentally, have their own views to defend on the very topics they attempt to present in an unbiased fashion. Individuals *may* overcome bias in some cases, and decisions *may* improve if they are made by collectives rather than individuals. Still, as we discuss in Sect. 2, there are strong psychological tendencies that bias one towards arguments and evidence that support one’s own position and against positions and objections that do not. It is not unreasonable to think

that these tendencies compromise the neutrality advertised by many encyclopedic representations of the field.

Even if encyclopedists are successful in overcoming psychological biases, their model imposes an overly restrictive view of what information belongs in representations of philosophy. Encyclopedias present a *time-slice* view of the field: the ideas and arguments that are currently considered important for understanding the topic of any given entry. People show up incidentally as the vehicles of these ideas and arguments, but do not play any substantial role in explanation. As an example of this point, consider two answers to the question, Why did Kant disagree with Hume's account of morality? One explanation might be that Hume and Kant both began with the same basic premises, but Hume made some mistake in drawing a sentimentalist conclusion from them, which Kant corrected (or vice versa, as you like). If so, there is nothing more to their disagreement than a simple logical error on the part of at least one author. Another answer might be that Kant (or Hume) supplied or assumed additional premises to derive his conclusions. This, too, could be explained purely in terms of the rational manipulation of ideas.

But surely there is more to be said about the field. Why, for example, did specific premises or specific inferences seem sound to one philosopher, but not another? How did specific ideas earn their place in the canon, while others never attained such status? Among those that did find their way into the canon, what explains the ebb and flow of their credibility? These questions go beyond ideas and arguments to consider changes within the field as a whole. As such, they require *diachronic* representations of philosophy for their answers. In these representations, people will play a more substantial role. Consider why, for example, Confucianism, rather than Mohism, came to dominate Chinese thought. During Mozi's lifetime, "his fame was as great as that of Confucius, and his teaching no less influential" (Fung 1948, p. 49). Chinese philosophy might have resembled Western philosophy much more closely than it does had the Mohists come to dominate the Confucians, rather than vice versa. Conversely, Western philosophy might have resembled Chinese philosophy more than it does if different schools of thought achieved cultural dominance in ancient Greece. In the Chinese case, it was not for obvious want of philosophical abilities that the Mohists are now found only in history books and the Confucians still play an important role in Chinese culture. To understand the rise of Confucianism—and likewise, perhaps, the rise of Platonism and Aristotelianism—we need to place the development of these philosophical schools in their social contexts. We suggest that the dominance and survival of some philosophical ideas at the cultural level is due, not only to their rational significance, but instead to factors such as influence, dissemination, conflict, and loyalty. With this observation in mind, we turn to the influence of social factors on the field and to the alternative representation of philosophy they suggest.

2 Social influence in philosophy

If we abandon a fetishism of ideas, what might we include in our wider representations of philosophy? Let us begin with the simple observation that nearly all philosophers' views are influenced by the views of those around them, either through *actual*

interaction with one's teachers, students, colleagues, and critics, or merely *imagined* dialogue with past figures and foes. In either case, one's own philosophical ideas do not come purely from consulting reason, but instead depend on processes of training and enculturation in the field.

Collins (1998) develops this point in *The Sociology of Philosophies: A Global Theory of Intellectual Change* to describe the personal ties—often conflictual—between major figures, which drive the field of philosophy. Collins distinguishes four senses of 'school', only some of which are relevant to understanding the development of philosophy. First, there may be "schools of thought," or individuals who have similar modes of thought, but need not be personally connected—present-day Platonists have presumably not had contact with Plato. These taxonomic classifications can be rather vague and need to be studied through more tangible structures. Second, intellectual historians often use the word 'school' to demonstrate transmission of intellectual influence. As Collins points out, this "is not an explanation, in the sense of showing why, out of all the persons who could read or hear of an idea, a certain few become important..." (p. 65). Such an explanation depends on a rigorous study of the final two senses of 'school'. One of these involves chains of personal relationships between teachers and pupils (vertical ties) and personal contacts between contemporaries (horizontal ties). Another consists of organizations, such as the Platonic Academy, the Aristotelian Lyceum, the Vienna Circle, or perhaps the Harvard Philosophy Department, where teaching and discussion take place and authority may be passed down through explicit succession.

In these latter senses of the term, schools are engines of inquiry that steer the course of philosophy as a discipline. They are places (physical and social abstractions) where philosophers have developed their positions under the influence of others. As Cohen points out, two ideas may be equally (or nearly equally) compelling, and one's decision about which to endorse may be made on nonrational grounds including the presentation of the idea or one's personal relationship with its proponent. While it is hard to pinpoint the exact causes of a philosopher's views and their relative contributions to his or her beliefs, we can observe rather striking trends of association between specific people, places, and ideas. Collins' study of social networks across 2,000 years of Eastern and Western philosophy reveals that "*The most notable philosophers are not organizational isolates but members of chains of teachers and students who are themselves known philosophers, and/or circles of significant contemporary intellectuals*" (p. 65, emphasis Collins).

As an illustration of the master–pupil tie, consider the importance of teachers in one's own philosophical training. Most beginning students in the field do not have well-worked-out philosophical positions, but only develop them under the tutelage of more experienced figures, who have arrived at their own views through the very same process. One may, on occasion, seek out a particular teacher because of one's antecedent philosophical interests. But where did those interests arise? Mostly likely, they have developed under the influence of earlier teachers and discussion with one's peers, or perhaps from a chance encounter with one text rather than another. It is more plausible to explain philosophical development in terms of social influence than pure reason for the same reason Cohen suggests: two sets of equally rational inquirers may disagree about a certain claim (e.g., the analytic/synthetic distinction), where

the primary difference between them is their members' closer interaction with one proponent in the debate, rather than the other.

We need not take these sociological claims to suggest some spooky social causation in order to understand how social ties influence philosophical practice. There are ordinary, well-documented psychological mechanisms by which individual philosophers' social environments can affect their judgments. As Collins himself points out, these claims about the influence of groups can be brought down to earth in terms of familiar social interactions: "The history of philosophy is to a considerable extent the history of groups. Nothing abstract is meant here—nothing but groups of friends, discussion partners, close-knit circles that often have the characteristics of social movements" (p. 3). Although more empirical study is needed to determine the precise effects of various mechanisms on philosophical practice in particular, we hypothesize that two broad psychological mechanisms operate in tandem to influence the course of philosophy.

The first mechanism is what social psychologists call "uniformity pressure," which is a form of social pressure that induces members of a group to seek uniformity of opinion within the group (Kruglanski et al. 2006, pp. 88–89). Psychologists have observed uniformity pressure at work in enduring social groups, such as political parties and families, as well as ad hoc groups formed in experimental contexts. Such pressure makes itself felt through interaction with other group members, who are motivated by peer disagreement to eliminate discrepancies among group members in certain kinds of beliefs. Group members might seek to eliminate these discomfiting discrepancies through either a "change other" strategy, in which they try to bring others' beliefs in line with their own, or a "change self" strategy, in which they change their own beliefs in order to match those of their fellows (Festinger 1950). The "change self" strategy might play an especially important role in shaping the views of newer students in an academic department who find themselves unable to match the argumentative abilities of higher-status faculty and more advanced students, all of whom have themselves been shaped by uniformity pressure already. The end result is a strong convergence, within social groups, on a set of shared beliefs.

Although philosophers may be better than others at overcoming such biases, philosophical claims are precisely the kinds of claims for which uniformity pressure is strongest. Festinger cites the empirical testability of the beliefs in question as a major determinant of the strength of uniformity pressure. To use his example, an individual who forms a judgment that a particular object is fragile by smashing it with a hammer is unlikely to be swayed by another group member's insistence that the object is unbreakable; but to the extent that the individual's claims cannot be easily tested, as counterfactuals about complex social events cannot, that individual is more likely to depend on others' agreement in deciding whether to hold fast to his or her opinion. Many philosophical claims, of course, fall nearer to the untestable pole of that continuum than not, and so we should expect even stronger uniformity pressure among philosophers than among some other disciplines. This may offset whatever cognitive advantages philosophers gain from their insistence on careful argumentation.

This is not to say that philosophers will not tolerate disagreement in their ranks. As in any discipline, there are always prominent points of contention in philosophy. The point is rather that a known social psychological phenomenon predicts that

philosophers who form stable social groups, such as within academic departments, will converge in some of their philosophical opinions. This convergence is enough to begin a process by which social interaction can shape the course of philosophy.

Once social influence begins to affect individuals' philosophical views, a second psychological tendency known as confirmation bias will often entrench that influence more deeply. Following Raymond Nickerson, we take "confirmation bias" to subsume several more specific psychological tendencies that lead individuals to seek and believe information that is consistent with their existing beliefs and to ignore, disbelieve, or to be more critical of information that is inconsistent with their existing beliefs (Nickerson 1998). Confirmation bias is so pervasive and so potent that Nickerson suggests, "If one were to attempt to identify a single problematic aspect of human reasoning that deserves attention above all others, the confirmation bias would have to be among the candidates for consideration" (p. 175). What makes confirmation bias so insidious, however, is that it operates unconsciously rather than through the deliberate distortion of evidence, and, according to some psychologists (Ditto et al. 1998; Fischer et al. 2005), it can sometimes arise from reasonable methods for allocating cognitive resources.

We hypothesize that four mechanisms of confirmation bias, out of the many identified by Nickerson, are particularly relevant to philosophy. First, the "primacy effect" consists in individuals' tendency to give more weight to information acquired earlier in an inquiry than information acquired later. Thus, formative philosophical experiences, which may be saturated by social influence, are likely to carry more weight than arguments encountered later, tilting the balance of reasons concerning contentious issues in one direction rather than another. Second, "motivated skepticism" involves a tendency to be more critical of information or arguments that challenge one's beliefs than of information and arguments that confirm them. For instance, a patient is far more likely to ask for a second opinion or to entertain alternative hypotheses when a doctor informs her that she has cancer than when a doctor gives her a clean bill of health (Ditto and Lopez 1992). This tendency probably causes philosophers to devote more time and energy to rebutting arguments against their position than thinking critically about arguments for their position. Third, the "preferential treatment of evidence" involves giving more weight to belief-consistent evidence than belief-inconsistent evidence. This sometimes occurs just because people are more likely to notice or recall evidence that supports their beliefs than evidence that does not. It sometimes occurs because people are likely to interpret evidence in ways that support, rather than challenge, their beliefs, even when other individuals interpret the same evidence in the opposite way. Thus, philosophers are likely to give more weight to arguments and information that supports their views than to other arguments or information, perhaps simply because these are the arguments that come most frequently to mind. Fourth, "selective exposure" is the tendency to seek out information that confirms one's beliefs rather than information that challenges it. This tendency may lead philosophers to devote more energy to finding or concocting arguments that support their own position and undermine others' than they do seeking out objections to their own position.

Of course, academia is structured to counteract confirmation bias, and philosophers in particular may be well suited to resisting it. None of this is meant to suggest that philosophers never seek out objections or that they fail to take objections seriously.

Confronting objections is a central part of philosophical practice. But it is an empirical question whether professional philosophers avoid this ubiquitous psychological bias altogether. The point of invoking confirmation bias is only to show that there is a widespread and powerful psychological tendency that may cause philosophers to persevere in the beliefs that uniformity pressure may cause them to adopt within their social groups.

Taken together, these psychological processes could explain how social factors influence individuals' philosophical judgments. These applications of psychological theories to philosophical practice would, of course, need to be tested, and we should expect that the actual processes at work are more complicated. Nevertheless, these conjectures lay grounds for the hypothesis that social factors influence the development of philosophy. Testing these empirical claims is doing *naturalized* metaphilosophy—studying the discipline itself in an empirical fashion. In Sect. 5, we will propose more sophisticated forms of data collection and analysis to assist this effort. Before considering such methods, it will be useful to reflect on the view of the discipline implied by naturalized metaphilosophy. In particular, we will consider the role that reason plays in the practice of philosophy and the implications of naturalized metaphilosophy for the idea of progress in philosophy.

3 The role of reason

Some may fear that naturalized metaphilosophy will deny reason any role in the practice of philosophy. If we can explain why philosophers adopt the philosophical positions they do by appealing to psychological and sociological factors, it may seem that they are not basing their positions on rational argument. This suggestion clashes with philosophers' subjective understanding of what they are doing. Most philosophers take themselves to hold the philosophical positions that they do on the basis of rational argumentation. But if each individual philosopher adopts theories on the basis of reason alone, then sociological factors play no significant role in the field.

We take as a provisional assumption that reason *does* play an important role in philosophy. "Reason," as we understand it here, is a psychological capacity to engage in reasoning, where "reasoning" is "the cognitive activity of drawing inferences from given information" (Grafman and Goel 2002). Some types of inferences are more likely to yield true beliefs (given true premises) than others. Good reasoning involves using those, and only those, types of inferences. To say that "reason" or "reasoning" plays an important role in philosophy, then, is to say that philosophers are significantly more likely to accept a claim if (they believe that) it is supported by good reasoning based on true premises.

We do not intend to exclude from naturalized metaphilosophy philosophers who are prone to more inflated notions of reason. The challenge for such philosophers is to craft a definition of reason suitable for empirical testing, without which it is impossible to determine empirically whether "reason" does play an important role in the field.

One task of naturalized metaphilosophy in the immediate future, then, is to explain the connection between the social hypothesis put forward in the previous section and the hypothesis that reason, as defined above, plays an important role in determining

which views philosophers adopt. In advance of empirical inquiry, there are reasons to think that some kinds of social influence coincide with philosophical reasoning. As Collins puts it, “intellectuals are oriented toward what they believe is the truth. They do not want to undermine their own truths, even though it is socially useful to have flawed truths which will keep their names alive in subsequent generations of creative workers” (p. 33). Philosophers constantly challenge one another to defend their claims through argumentation. Philosophers who hold fast to claims that cannot be rationally defended presumably invite the ridicule of their peers and the exclusion of their views from the philosophical mainstream. Contrapositively, it is reasonable to hypothesize that most philosophers who *do* rise to positions of social prominence within the field are those who excel at offering compelling arguments for their claims. To the extent that these individuals shape the social influence at work in the field, that influence will reinforce the effects of argumentation on others’ philosophical views, rather than run counter to it. Both of these processes—the marginalization of those who cannot rationally justify their claims and the exaltation of those who can—require that reason plays a prominent role in shaping the discipline. If empirical study shows these processes to be as important as they initially appear, then this study will help clarify the ways in which reason and social influence interact in philosophical practice.

Naturalized metaphilosophy may be able to do more than just clarify this interaction; it may help us harness social factors for the good of the discipline. We can appreciate this point by way of a parallel discussion in the history and philosophy of science. As critics have pointed out, various social factors affect *which* hypotheses get proposed, *how* they are tested, and *whether* results received recognition. These factors include demographics (The Biology and Gender Study Group 1989; Hays-Gilpin and Whitley 1998; Hrdy 1981; Keller 1985; Lloyd 1993; Longino and Doell 1983; Spanier 1995); science funding, administration, and policy (Hull 1988; Kitcher 1990, 1993; Solomon 2001; Strevens 2003); institutional structures of reward (Keller 1983; Latour and Woolgar 1979/1986; Waring 1990); moral or political views; etc. Understanding the effects of these factors is a precondition for mitigating undesired bias in science. As Antony (1991) argues, “an empirical study of biases” is needed to “tell us something about the reliability and corrigibility of biases of various sorts. It may turn out that we can on this basis get something like a principled sorting of biases into good ones and bad ones, although it will be more likely that we’ll learn that even a ‘good’ bias can lead us astray in certain circumstances” (p. 216). Similarly, an empirical study of the social factors at work in philosophy may help to reveal which kinds of social influence resulted in interesting and original philosophical ideas, and which have led us down more frustrating or less fruitful paths. These observations about the beneficial or detrimental effects of various kinds of social influence would help to inform practical decisions about admission, employment, tenure and promotion, even publication and citation in the field. The claim here is not that naturalized metaphilosophy will reveal determinate principles that ought to govern our policies. Instead, the hope is that by understanding more fully the ways in which the field operates, we might in turn revise our practices to enhance argumentation in the field.

Even if social influence is present in philosophy, it need not be deleterious. Philosophers can rely on reason to justify (or reject) beliefs that they have adopted primarily for social reasons, even if confirmation bias makes this difficult. By awarding

opportunities for social influence to individuals at least partly in accordance with their good use of rational argument, social forces can enhance the role of reason in philosophy. Achieving a level of understanding sufficient to shape philosophical practice wisely will require careful empirical study, to which we turn in Sect. 5.

4 The possibility of progress

Another fear raised by naturalized metaphilosophy is that the influence of social factors undermines the possibility of progress in philosophy. If philosophical theories rise and fall because of purely social factors, their dominance and decline may have more to do with fashion than with truth. The ruffled collars of Elizabethan gentry may be hard to find today, but this is for social reasons, not because fashion has “progressed” beyond “false” beliefs about the attractiveness of ruffled collars. If the explanations for the general rejection of the teleological argument for God’s existence are sociological or psychological, it might be similarly difficult to explain its rejection as “progress.”

If we understand progress in philosophy to consist in coming to have more true beliefs and fewer false beliefs, there are three arguments for the claim that the influence of social factors will undermine progress. Two of these arguments come from the literature on scientific realism: Putnam’s (1978) metainductive argument (MI) and Sklar’s (1981) unborn hypothesis argument (UH). The third is more specific to the field of philosophy because it depends on the fact that philosophers return constantly to ancient, medieval, and early modern sources.

MI takes as its starting point the observation that past scientific theories have been wrong in their fundamental posits about the world. It amasses cases of luminiferous ether, phlogiston, the four humours, and even Newtonian mechanics—all of which were leading candidates in their day yet subsequently rejected for newer theories—and then concludes that it is highly likely that our present theories will suffer the same defeat. According to MI, we have no reason, at any time, to believe that current science is correct, given that it has been incorrect in many past cases. As Putnam puts the objection, “just as no terms in the science of more than fifty (or whatever) years ago referred, so it will turn out that no term used now...refers” (p. 25). The appearance of progress throughout the history of science is merely illusion foisted upon a series of errors. There is, of course, a good Kuhnian explanation for why these views survived for as long as they did: each was regarded as authoritative during its day and passed down to successive generations of scientists in the course of their training in the discipline.

In a parallel way, one might run MI on the naturalized view of the discipline we have proposed: Many past philosophical views are regarded today as crude, incorrect, oversimplified, and so on. They were believed, in large part, because of social factors, rather than the light of reason. Though some take this judgment to reflect philosophical progress—the approximation of truth in the limit of inquiry—the failure of these past cases should invite skepticism of, rather than inspire confidence in, our current philosophical views. We have every reason to believe that our current views are as erroneous as our past ones, and naturalized metaphilosophy explains our shifting allegiance to different ideas based on the pull of social factors. The difficulty posed by MI is even

greater in the case of philosophy than it is in the case of science. At least scientists can point to greater technical success than their predecessors and can more easily explain their predecessors' failures in terms of limited experimental abilities (Devitt 1991). Aside from marked improvements in formal logic, what technical results can modern philosophy boast over the eighteenth century? What tools of inquiry do we have, which Hume or Kant lacked, that better equip us to find the truth?

Another line of objection takes a less devastating, but equally pessimistic view of social influence in philosophy, drawn from Sklar's discussion of "unborn hypotheses." Sklar worries that if unborn hypotheses (theories that have not seen the light of day) were to be born, it is likely that some would replace our current hypotheses, and by their lights, our current hypotheses would be wrong. "[W]hat credibility," he asks, "can accrue to the victor in a battle for survival which, by historical accident and paucity of imagination, simply keeps nearly all of the competitors out of the arena?" (p. 20). Adapting UH for the present context, one might worry that certain philosophical ideas have gained (or lost) prominence only because their adherents (or opponents) have been well-connected, influential figures in the historical network of philosophers. Were *others* to occupy those nodes in the network, their ideas would have entered the arena, displacing the views we hold in the present to be correct, interesting, fruitful, etc.

The common worry posed by MI and UH in this context is that naturalized metaphilosophy will reveal our current views to be massively erroneous or unimportant—or, to put the point more mildly, so highly contingent that our credence in them would be severely threatened. Those who think that philosophy is a purely rational enterprise carried out according to the normative requirements of reason will deny that our philosophical judgments are contingent on social factors. Rather than consigning philosophy to the flames, they will argue, we should sooner reject the naturalized metaphilosophy that called it into question. This line of argument, however, is highly irrational. It is an empirical question whether philosophy is a *purely* rational enterprise devoid of social influence, a question that can only be settled by empirical investigation. To avoid empirical evidence on the matter is to succumb to confirmation bias, or perhaps just to irrational wishful thinking, thereby undermining the claim that philosophy is a purely rational enterprise. If this claim is correct, the rationalist should welcome empirical evidence, for it will reveal a lack of social influence.

If, however, there is evidence of social influence, we would do well to follow Bloor's (1976/1971) advice on the sociology of science:

It is quite possible to sweep this empirical observation [that scientific disputes are often social "priority disputes"] aside and declare it to be irrelevant to the true nature of science. Science as such, it may be said, develops according to the inner logic of scientific enquiry and these disputes are mere lapses, mere psychological intrusions into rational procedures. However a more naturalistic approach would simply take the facts as they are and invent a theory to explain them. (p. 22)

Nothing inherent to the methodology of naturalized philosophy lends support to the skeptical threats of MI and UH. The rational path of inquiry is to gather whatever evidence is available and develop a theory to explain that evidence.

As part of this task, philosophers must explain a curious phenomenon in the discipline concerning progress: Unlike the sciences, which focus on current research, philosophers routinely study, cite, and engage with theorists who were writing several centuries, if not millennia, earlier. These sources remain en vogue, presumably, because teachers continue to assign them to students as canonically important texts in the field. But the mere fact that a text is well regarded does not entail that its claims are true, so we are again presented with a case in which social influence possibly leads the field astray, thereby hampering philosophical progress. To put the point more finely, either philosophers turn to old sources because they think those sources contain enough true claims to be worth reading, even at the expense of reading contemporary philosophers; or they devote attention to old sources for reasons other than to gain true beliefs. Either interpretation threatens the possibility of progress. On the first interpretation, this phenomenon suggests that we have not come far enough since Plato and Aristotle to set them aside. How often do contemporary scientists invoke or grapple with Newton, much less Ptolemy or Archimedes, in their work? On the second interpretation, it suggests that philosophers are so consumed with a search for something besides truth that any progress we make, as we understand ‘progress’ here, would be almost coincidental.

Rather than fueling this skeptical fire, naturalized metaphilosophy suggests that we need not worry about it. On the first horn of the dilemma—the one on which we return to Aristotle, for example, because we think Aristotle was right with respect to some claim—we may indeed believe that Aristotle’s claims were correct because of social factors involving in our training in the field. But it would be a genetic fallacy to suppose that, because social influence directs us to read him, we err in thinking that he was right about some claim. The process by which we acquire our beliefs about which texts contain important truths does not determine whether those texts do contain such truths. On the other horn of the dilemma—the one on which we return to Aristotle, say, for reasons other than truth—tracing the lines of social influence may help us discover *which* aspects of a text we find valuable over and above its possible truth. These criteria for evaluating theories may include parallel characteristics from the philosophy of science: consistency, whether a theory contradicts itself or other currently accepted theories; scope, the degree to which a theory extends beyond the particular domain it was initially designed to explain; simplicity; and fruitfulness of further research programs (Kuhn 1977). This metaknowledge about the discipline may be useful, upon reflection, for revising our present views of the field and making evaluations that better fit that reality.

Naturalized metaphilosophy, as we suggested in the last section, may be able to provide us with tools to better decide what counts in making progress in the field and adjust decisions of awards, promotions, and honorable posts in light of that understanding. Far from calling philosophical progress into question, naturalized metaphilosophy may be able to explain and enhance that progress through its empirical investigations.

5 A research proposal in naturalized metaphilosophy

In Sect. 2, we hypothesized that social factors influence the development of philosophy, and in Sects. 3 and 4, we hypothesized that positions of status and reputation—and

Table 1 Some formal relationships and sources of documentation for them

Type of relationship	Sources of documentation
Student/teacher relationships	
Advisor/advisee	Dissertation front matter
Classroom student/teacher	Various sources ^a
Peer/peer relationships	
Fellow students	Degree dates in dissertations, CVs, etc. ^b
Fellow faculty members	University catalogs ^c
In-print relationships	
Citation (i.e., who cites whom)	Publications
Acknowledgements	Publication front matter, footnotes, autobiographies
Editor/contributor	Anthologies/journals

^a These can be very roughly inferred by comparing students' dates of attendance in a program with the lists of faculty teaching in the program at the time, narrowed according to the students' and faculty members' areas of interest

^b At some institutions, doctoral candidates traditionally include a CV at the end of their dissertation, which gives dates and institutions for undergraduate degrees. Some institutions' alumni records also contain information about doctoral students' undergraduate degree

^c University archives maintain copies of all annual or biannual university catalogs

thus social influence—will tend to be awarded to philosophers who offer rationally compelling arguments for their views. These hypotheses about the role of social influence in philosophy are speculative; their confirmation will require comprehensive data about the discipline that is not available in any existing representation. Such a test will require substantial data about relationships between philosophers, theoretical methods for extracting useful information from that data, and techniques for rendering those inclusions in cognitively salient ways. All three of these elements are currently available, but they have not yet been brought together in a form sufficient to evaluate our hypotheses.¹

Investigating social influence in philosophy requires compiling data about the social relationships, both formal and informal, between philosophers. The availability of this data varies by era and geographical location. Since around 1880 in North America and longer in Europe, institutions have maintained useful records about personal connections between philosophers. The written record of the field stems back even further. Relationships of particular interest are explained in Table 1.

The relationships people establish through correspondence, conference attendance, scholarship on the same topics, membership in the same professional associations, and mutual friends are less formal, though at least as important as these

¹ At present, there are several projects that address portions of this task: Josh Dever's Philosopher's Family Tree (Dever 2008), David Chalmers' Australasian Family Tree (Chalmers 2008), the Indiana Philosophy Ontology Project (Allen et al. 2008), various departmental placement pages (vid. Blatti 2008), and Phylo (Morrow and Sula 2008), which explores the historical network of individuals, institutions, and ideas in the history of philosophy.

formal relationships. Informal relationships can sometimes be documented, though with greater difficulty and less regularity.

Data on philosophers living and working before the late nineteenth century may prove more challenging. In many cases, institutional records do not exist, and written works are lost or inaccessible. However, Collins' achievement in tracing relationships among Western and Eastern philosophers over two millennia does give some hope that relationships parallel to those listed in the table above can be established, even for very ancient figures. In part, this process relies on the role of experts in the field who have catalogued and examined the documents pertaining to individual philosophers, including where they lived, with whom they communicated, and what they wrote about. Moreover, several digitization projects, such as the Universal Digital Library and Google Books have made explicit commitments to preserving older texts and making them publicly accessible. The digital form of these materials can allow for automated extraction of keywords and references, which greatly simplifies the task of analyzing millennia of data from the field. In some cases, these methods, too, must be enhanced with the input of experts. An eighteenth century moral philosopher, for instance, may make reference to "the author of a treatise on moral sentiments," without naming Francis Hutcheson or David Hume explicitly. This reference might escape a computer algorithm, but it would be transparent to any historian of ethics. Identifying which (pieces of which) materials are purely machine-readable from which require expert input could be made more efficient with filtering algorithms trained to recognize possible references and queue them for further analysis. Requests for input could also be distributed across a group of experts, with each contributing small but important pieces of data, thereby minimizing the amount of work required by any single expert.²

If technology-enhanced methods of data collection are successful, the vast amounts of information gathered will require analysis in ways that make them cognitively salient to interpreters who wish to evaluate our hypotheses. One particularly promising method of analysis lies in the "science of networks" (Barabási 2002; Buchanan 2002; Watts 2003), a burgeoning field of research that has its roots in the field of discrete mathematics known as graph theory. Each network is made up of a set of discrete elements (i.e., vertices, nodes, actors) and a set of connections (i.e., edges, links, relationships) between them, both of which vary by discipline. In recent decades, studies have included populations and disease, chemicals and reactions, people and information, and publications and citations. Of particular interest for our purposes are social networks, which can be used to study decision making, belief systems, diffusion and adoption of innovations, coalition formation, and other phenomena (Wasserman and Faust 1994). Network analysis can yield information about *individual* actors, including their prominence and the roles they play as isolates, liaisons, bridges, etc.; *pairs* of actors, including distance and reachability; and *group-level* properties, including centralization, density, prestige, and recurring structural patterns (equivalence classes and blockmodels). Current work by network analysts includes the study of multiple relations, dynamic networks, and longitudinal network data.

² Our own project, Phylo, makes wide use of this distributed input from experts, combined with automatic methods of data collection and analysis.

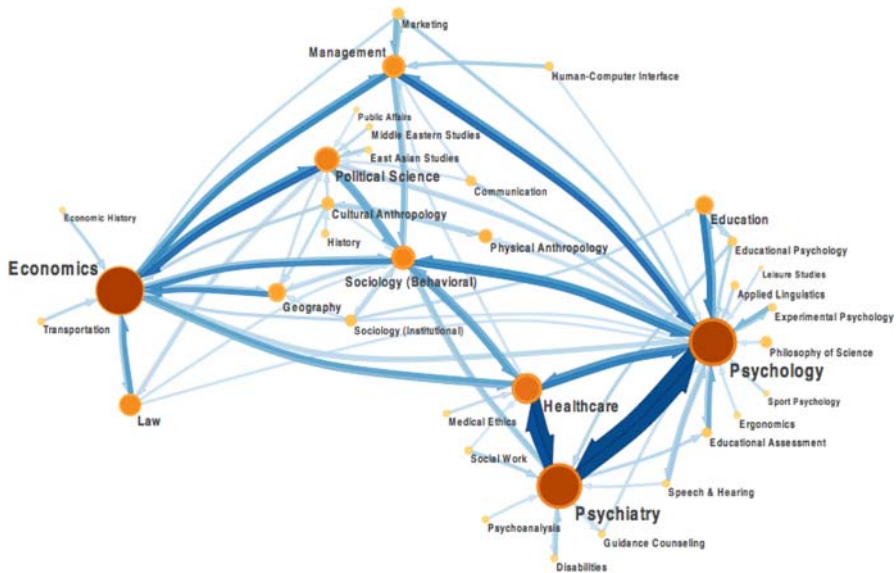


Fig. 1 Rosvall and Bergstrom's map of the social sciences based on 1,431 journals and 217,287 citations. The nodes correspond to different subdisciplines. The direction and thickness of the edges correspond to the flow and number of citations between those subdisciplines. (Rosvall and Bergstrom 2007)

By now, the prospects for network analysis in testing our hypothesis should be clear: the science of networks can serve as a tool for studying empirical observations of the field in a rigorous fashion. As an example of this application, consider Rosvall and Bergstrom's study (2007) of citation patterns in more than 1,400 social science journals during a single year. Using tools of network analysis, they generated a map of the social sciences showing the relative density of citations in disciplinary clusters of journals, as well as the weight and direction of citations between them (see Fig. 1).

The resulting map of the social sciences reflects some of our intuitive understanding of these disciplines and their connections. But its empirical foundation provides a measure of accuracy and reliability that is lacking in the nonempirical work. Similar applications of network analysis to citations in philosophy, as well as other forms of data discussed above, could yield representations of the field that go beyond traditional models in the kinds of information they reveal and the empirical support they provide. Such representations give us a rigorous way of identifying prominent individuals, institutions, and ideas; discovering patterns and trends in the way that the field develops over time; locating important divisions within the discipline and points of contact between subdisciplines; and tracking demographic changes in the discipline.

Good representations in naturalized metaphilosophy will strike a balance between uncovering important structural features of the discipline and excessive amounts of detail. Emerging technologies of information visualization have an important role to play in creating intuitive displays that convey the results of network analysis in accessible ways to human users.

6 Conclusion

At the heart of naturalized metaphilosophy lies a recognition that philosophy is an activity carried out by human beings—social creatures who, despite being able to carry out complicated chains of reasoning, are deeply influenced by their social context. Traditional representations cling to a myth of pure rationalism—the myth that arguments, and nothing else, drive the course of philosophy, according to their objective logical merit. If this were true, it would be possible to understand the history of philosophy strictly in terms of the arguments on which these representations focus, and it would be possible to identify the important arguments (and the individuals through whom they were articulated) from a neutral, encyclopedic standpoint. But casual observations like Cohen’s remark about the analytic/synthetic distinction and more rigorous empirical studies like Collins’ suggest that there is more going on in the discipline. Naturalized metaphilosophy extends this investigation to discover what other forces shape the development of philosophy. In defending this approach, we do not wish to advance social factors at the exclusion of the rational ones on which the field is consciously focused. The latter deserve a central place in representations of the discipline, for they play important causal roles. We are simply pointing out that they cannot tell the whole story; if we want to understand the practice of philosophy in a broader sense, we need to look beyond ideas and arguments alone.

Naturalized metaphilosophy is more than a study of error. It is not an attempt to debunk philosophy or scandalously reveal socially induced failures of reason. Instead, we prefer to think of it as a study of success. If we can identify circumstances in which social forces cause philosophers to stumble in their pursuit of truth, we can also identify circumstances in which social forces nurture good philosophy—circumstances in which philosophical inquiry is most successful or most productive or most transformative. And then, equipped with that knowledge, we can improve the social structure of the discipline to ensure that the circumstances in which we actually do philosophy are as close to those ideal circumstances as possible.

References

- Allen, C., Buckner, C., & Niepert, M. (2008). InPhO on the web. Retrieved May 31, 2008, from the Indiana Philosophy Ontology project’s Web site: <http://inpho.cogs.indiana.edu>.
- Antony, L. M. (1991). Quine as feminist: The radical import of naturalized epistemology. In L. M. Antony & C. Witt (Eds.), *A mind of one’s own* (pp. 185–225). Boulder, CO: Westview Press.
- Barabási, A.-L. (2002). *Linked: The new science of networks*. Cambridge, MA: Perseus.
- Blatti, S. (2008). Placement data. Retrieved May 31, 2008, from De Dicto: <http://dedicto.blogspot.com/2008/05/recent-placement-data-of-top-30.html>.
- Bloor, D. (1976/1991). *Knowledge and social imagery*. Chicago, IL: University of Chicago Press.
- Buchanan, M. (2002). *Nexus: Small worlds and the groundbreaking science of networks*. New York, NY: Norton.
- Chalmers, D. (2008). Australasian philosophy family tree. Retrieved May 31, 2008, from David Chalmers’ Web site: <http://consc.net/tree.html>.
- Cohen, G. A. (2000). *If you’re an egalitarian, how come you’re so rich*. Cambridge, MA: Harvard University Press.
- Collins, R. (1998). *The sociology of philosophies: A global theory of intellectual change*. Cambridge, MA: Harvard University Press.

- Dever, J. (2008). The philosophy family tree. Retrieved May 31, 2008, from The University of Texas at Austin Web site: <https://webspaces.utexas.edu/deverj/personal/philtree/philtree.html>.
- Devitt, M. (1991). *Realism and truth*. Princeton, NJ: Princeton University Press.
- Ditto, P. H., & Lopez, D. F. (1992). Motivated skepticism: Use of differential decision criteria for preferred and nonpreferred conclusions. *Journal of Personality and Social Psychology*, 63(4), 568–584.
- Ditto, P. H., Scepansky, J. A., Munro, G. D., Apanovitch, A. M., & Lockhart, L. K. (1998). Motivated sensitivity to preference-inconsistent information. *Journal of Personality and Social Psychology*, 75(1), 53–69.
- Festinger, L. (1950). Informal social communication. *Psychology Review*, 57(5), 271–282.
- Fischer, P., Jonas, E., Frey, D., & Schulz-Hardt, S. (2005). Selective exposure to information: The impact of information limits. *European Journal of Social Psychology*, 35, 469–492.
- Fung, Y. (1948). *A short history of Chinese philosophy: A systematic account of Chinese thought from its origins to the present day*. New York, NY: The Free Press.
- Grafman, J., & Goel, V. (2002). *Neural basis of reasoning and thinking*. *Encyclopedia of cognitive science*. London: Macmillan.
- Hays-Gilpin K., & Whitley D. (Eds.). (1998). *Reader in gender archaeology*. New York, NY: Routledge.
- Hrdy, S. (1981). *The woman that never evolved*. Cambridge, MA: Harvard University Press.
- Hull, D. L. (1988). *Science as a process: An evolutionary account of the social and conceptual development of science*. Chicago, IL: University of Chicago Press.
- Keller, E. F. (1983). *A felling for the organism*. San Francisco, CA: W. H. Freeman.
- Keller, E. F. (1985). The force of the pacemaker concept in theories of aggregation in cellular slime mold. In E. F. Keller (Ed.), *Reflections on gender and science* (pp. 150–157). New Haven, CT: Yale University Press.
- Kitcher, P. (1990). The division of cognitive labor. *Journal of Philosophy*, 87, 5–22.
- Kitcher, P. (1993). *The advancement of science*. Oxford: Oxford University Press.
- Kruglanski, A. W., Pierro, A., Mannetti, L., & De Grada, E. (2006). Groups as epistemic providers: Need for closure and the unfolding of group-centrism. *Psychological Review*, 113(1), 84–100.
- Kuhn T. S. (1977). Objectivity, value judgment, and theory choice. In T. S. Kuhn (Ed.), *The essential tension: Selected studies in scientific tradition and change* (pp. 320–339). Chicago, IL: University of Chicago Press.
- Latour, B., & Woolgar, S. (1979/1986). *Laboratory life: The construction of scientific facts*. Beverly Hills, CA: Sage.
- Lloyd, E. (1993). Pre-theoretical assumptions in evolutionary explanations of female sexuality. *Philosophical Studies*, 69, 139–153.
- Longino, H., & Doell, R. (1983). Body, bias, and behavior: A comparative analysis of reasoning in two areas of biological science. *Signs*, 9, 206–227.
- Morrow, D., & Sula, C. A. (2008). The phylosophy project (phylo). Retrieved May 31, 2008, from The Phylosophy Project's Web site: <http://www.phylosophy.net>.
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2(2), 175–220.
- Putnam, H. (1978). *Meaning and the moral sciences*. London: Routledge & Kegan Paul.
- Rosvall, M., & Bergstrom, C. T. (2007). Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences of the United States of America*, 105(4), 1118–1123. doi:10.1073/pnas.0706851105.
- Russell, B. (1912/1997). *Problems of philosophy*. New York, NY: Oxford University Press.
- Sklar, L. (1981). Do unborn hypotheses have rights. *Pacific Philosophical Quarterly*, 62, 17–29.
- Solomon, M. (2001). *Social empiricism*. Cambridge, MA: MIT Press.
- Spanier, B. (1995). *Im/partial science: Gender ideology in molecular biology*. Bloomington, IN: Indiana University Press.
- Strevens, M. (2003). The role of the priority rule in science. *Journal of Philosophy*, 100, 55–79.
- The Biology and Gender Study Group (1989). Importance of feminist critiques for contemporary cell biology. In N. Tusana (Ed.), *Feminism and science* (pp. 172–187). Bloomington, IN: Indiana University Press.
- Waring, M. (1990). *If women counted*. San Francisco, CA: Harper Collins.
- Wasserman, S., & Faust, K. (1994). *Social network analysis: Methods and applications*. New York, NY: Cambridge University Press.
- Watts, D. J. (2003). *Six degrees: The science of a connected age*. New York, NY: Norton.