Conventional Norms of Reasoning

ADAM MORTON *University of Alberta*

ABSTRACT: There are conventional norms of reasoning. That is, we have conventions about which patterns of reasoning we encourage *or* disapprove of.

RÉSUMÉ: Il existe des normes conventionnelles du raisonnement. C'est-à-dire que nous avons des conventions quant aux modèles de raisonnement que nous encourageons et désapprouvons.

I am going to argue that there are conventional norms of reasoning. That is, that there are conventions, with the element of arbitrariness that is essential to a convention, that we use to regulate our thinking. We could have used other conventions. I expect that, stated this way, the idea seems implausible. It sounds like the suggestion that given that one believes that it is snowing and also believes that when it is snowing gloves are a good idea, one then arbitrarily goes along with the convention that gloves are a good idea. Other conventions also are possible, so that one's conclusion might have been to go bare-handed, or to put both gloves on one hand. This seems crazy, especially in winter. Well, my suggestion is not this crazy thing. It is rather that we have conventions about which patterns of reasoning we *encourage* or *disapprove of*, about when and how we push one another into reasoning various ways, and to some extent how we react to these pushes. Some thinking works better than others for many reasons, some of which will not be changed by any amount of convention. But with which aspects of thinking we nudge one another, and how we direct our nudges, is something for which there is a large element of convention.

***Dialogue* 50 (2011), 1–14.**

doi:10.1017/S0012217311000357

Norms and Theories

We criticize one another for foolishness and commend one another for sharp thinking. We recommend ways of solving problems, designs of experiments, strategies for investigating phenomena. There are principles of logic, statistics, and the assessment of risk that we hold one another to. It is not obvious how we do this: what we comment on, and what effect it has. Some of our comments are part of the everyday lore and practice of educated people, and some are based on very sophisticated theories. These theories have developed in a way that is very significant for our purposes. The original normative theory of reasoning is Aristotelian logic, which while certainly not trivial, and extending to inferences that stretch our capacities, is only scratching the surface of the complexity of deductive logic. We also have formal theories of evidence, expressed probabilistically, and theories of expected utility. I shall call all of these N-theories, and I'll call information contained in them N-facts. They all describe facts that are often relevant to successful thinking, in a form that makes them relatively subject-neutral. They typically are theories hard to master fully, and anyone will find aspects of them that they typically misunderstand or balk at. Suppose an unending series of human people of increasing reasoning capacity ("intelligence" we might say, if we trusted the word). Then for each person there would be a question for which he or she could not get the answer that a more capable person would get from logic or statistics. Everyone hits his or her limit.

The complexity of these theories is one of several reasons not to take them as prescribing patterns of thought which one *must* follow if one is to succeed, because we never fully follow them and we do sometimes succeed. And it is just not true that the nearer we come to following the thinking of some marginally conceivable agent who fitted the prescriptions perfectly, the better off we would be. So our attitude should be to take them for what they are, rich theories giving relevant information. We do not need to water them down to fit our capacities, as in very different ways Cherniak, Pollock, and Weirich have suggested. What *does* have to fit our capacities is the pressure we put on one another to pay attention to an N-theory, and the parts of the theory that we draw attention to. That is the topic of the next section.

Conventional Norms of Advice

We often comment on one another's thinking. We are aware of some of what we think; we monologue to some extent, and we answer questions about our reasons for what we say and do. And we know the situations of others enough to be able to think of lines of reasoning that might be relevant. As a result, we intervene critically in our own and other people's thinking in a number of ways. We suggest conclusions and routes to them ("Here's something you should consider."). We object to the way conclusions are arrived at ("That doesn't follow – what if . . . ? "). We commend bits of good reasoning ("That's

clever, of course given that . . . "). Sometimes we bluntly reject thinking ("Something's wrong there."). And when we are being sophisticated, we focus on the reasoning in contrast to the conclusion ("I agree, actually, but I don't think that's a good reason.") The effect is that we try to inhibit some trains of thought, encourage others, and nudge ourselves and other people into profitable patterns.

An interesting relation exists between comment-giving and N-theories. There is usually some connection between what someone is thinking and other relevant possibilities, which consists in a deductive or probabilistic relation between the thought and the possibility. So drawing attention to the connection is sometimes useful. But because such connections often are irrelevant, drawing attention to them also can be anti-helpful. When someone's belief is unfounded, drawing attention to a correct inference to an opposed conclusion, or redoing their reasoning in ways that depart from what they were thinking, will sometimes undermine it. Sometimes this can take an informal model-theoretic form, where one shows how the conclusion could turn out to be false, or very unlikely. I do not want to speak as if thinking consists of chains of deduction, or anything similar, which are then compared to standard paradigms (I take it that this is a subject of serious psychological controversy, as shown by the essays in Elio). All I need is that drawing attention to a fact drawn from an N-theory often changes people's thinking.

Though people everywhere comment on one another's thinking, doing so by reference to an N-theory is restricted to a fairly small circle. It is not as exotic as one might think, though, since closely connected activities, such as providing counter-examples and alternative possibilities, are much more common. And there is a stock of shared habits and opinions about reasons for belief and action. We have doctrines of how to make convincing arguments and to identify and resist other people's misleading ones. We have moral doctrines and theories of how to be socially effective. Particularly important from an epistemological point of view are doctrines about the design of experiments, an amalgam of statistical theory and reflections on scientific practice that has evolved into a normative lore that is at the heart of science, though philosophers do not discuss it enough. I discuss experimental design in a separate section, below. There is a creative aspect to this: we can find novel ways to direct recommendations at people, and we can find recommendations that will impact particular people's activities in ways that fit their thinking. We can mine the richness of statistics for ideas that we can apply imaginatively to our thinking as it is. Real normative activity is in general much less an application of fixed rules and much more a creative a process than philosophical descriptions sometimes suggest.

Many ways of commenting on our reasoning share a feature that makes them particularly effective. Let me call it the persuasion phenomenon. It involves nudging people into changing their beliefs or intentions, or revising their reasoning, just by presenting an alternative. One doesn't say "You're

wrong" but rather, "Consider this." For example, a confident denial that switching is a good strategy in a Monty Hall case can be countered by showing how over a long series of trials, switchers will come out ahead. Or a belief that most water is liquid can be countered by the argument that "Most of the water in the universe is not on Earth; the temperature of almost none of the universe is between 0 and 100 degrees; therefore, most water may well be either steam or ice." What happens is that evidence is brought to a person's attention, then he or she weighs it and miraculously changes his or her mind. Very often what I have called evidence consists in an item from an N-theory. So the persuasion phenomenon is in part about the way that the patterns of reasoning that are normatively endorsed are capable of changing our opinions, so that mere presentation of them, rather than the force of social approval and disapproval, does a lot of the work.

Not that any presentation of any theoretically relevant N-fact will do the trick. My take on what is going on is that N-theories give facts, often basically mathematical facts, that are relevant to questions of truth, probability, or effectiveness. When a fact is easily enough understood and when it combines with what a person already thinks, it points to a conclusion. That conclusion is often then adopted. We can say this without assuming that the thinking that lies behind this persuasion resembles a mathematical proof, a probabilistic argument, or a logical derivation. (And without assuming lack of resemblance, either. We can be neutral on the psychology of thinking.)

Very often there is a tiny gap between the absorption of the N-fact and the change of opinion. One has learned a higher-order general fact: that in many cases p will be true, or that if some assumptions are true then p is true, or that if a means is achieved an end will follow. Then one has to make the leap to believing p or intending the means. One can hold back, seeing that one "should" believe or intend, but not being able to. In paying attention to the N-fact one has done something to change one's opinion, something which as a piece of human psychology usually works, but which in this case has failed. This manoeuvring of oneself or others into a position where one's psychology combines with a general high-order fact to bring about a valuable outcome is at the heart of many traditions of normative comment. (Compare the moral case. One is brought to believe that one should do or want something. Then one has to get oneself to want or intend it. Usually the connection is pretty immediate, though it is mysterious quite how it works: see the last chapter of Michael Smith. It seems plausible that considerations about what is generally best and considerations that get one to shift one's perspective to that of another serve a similar role to the higher-order N-facts.)

It is here that I think an element of convention enters. Comment on our thinking can be helpful or not. Good advice is best taken, and it is usually best for the advice-giver, too, that it be taken. This is partly because we usually care for the people we interact with, so that what is best for them is in that way good for us, and partly because that interaction is usually co-operative, so that it is in

our individual interest that they do their part in achieving common goals. (And there is the matter of forming an identity together with others, which one could extract from David Velleman's work; see ch.7.) Suppose for example that when people utter explicit outright contradictions, others point this out to them in a tone of disapprobation. If these comments are ignored there is no benefit to those they are addressed to, nor either to the comment-givers, who may be trying to share activities with them and so need definite consistent reports. There is no benefit also if the comments are not ignored but do not get the right reaction: if the recipients do not pause and reconsider. But a pause and reconsideration when a contradiction is pointed out, typically is often beneficial both for the giver and for the receiver of the comment.

Conventions of normative discourse tend to make people coordinate their activities, in ways that help them succeed. As Allan Gibbard puts it:

> Normative judgements tend towards consensus—shakily, but not by accident. With some other judgements consensus is automatic: we easily agree on the layout of surrounding rocks and trees. . . . We nudge each other to agree on norms too. We do this in a cluster of ways and, agreement achieved, we treat norms like rocks and trees, more or less. . . . In normative discussion we influence each other through conversational invitations and pressures. We demand consistency one of another, and try to meet these demands. These socio-psychic mechanisms combine, at times, to make norms as interpersonal as trees. (249)

So the comment-giving is something like a Lewisian convention, in the sense that it is a solution to a coordination problem. It is in the interests of both the comment-maker (advice-giver) and of the recipient of the comment, that the advice is taken. (Gibbard refers in this context not to Lewis but to Lewis's source, Thomas Schelling.) This is so, given that a list of conditions is met. There has to be a general practice of heeding such advice, involving both suitable reactions (pausing and rethinking, but perhaps not too obstructively or too long) and acceptance of these reactions by others. There have to be suitable longer-term follow-ups, and these must take place in a suitable intellectual-cultural context. And there may be other conditions. It would be very hard to specify in full detail what is needed for a practice of normative comment to be a solution to a coordination problem among a number of interacting thinkers.

(The issue of whether other features of Gibbard's mild anti-realism in meta-ethics can transfer to norms of reasoning in general is interesting and important, but I shall not discuss it. See chapter 5 of Cuneo.)

Four linked points are worth making. The first is that it is the giving and taking of advice or comment, at suitable moments, that is the mutually beneficial equilibrium, rather than the behaving in accordance with the normative principles that underlie the advice. For example, it may be a generally good idea to point out to someone that since she agrees that she is pregnant, taking a certain medicine may be a mistake, given the printed warning that "if you are

pregnant or have a weakened immune system, then use of this drug is contra-indicated." But anyone who went around generally adding *p or q* to her beliefs whenever she realized she believed *p*, would clutter her mind so as to obstruct her own projects and interfere with her co-operation with others. (This is a Gilbert Harman theme; see part 1 of Harman.) Giving good advice is less of a menace than actually taking it, but the real benefit comes from selective giving and taking at suitable moments.

Second, the convention is beneficial because the N-facts that are drawn attention to are true. The convention is not a source of those truths: they're just facts. But it is a source of their having normative status, that is of the attention directed at them having a certain authority.

Third, the pressure to conform is generally gentle. There is no clear line between holding one another to norms of good thinking and co-operative problem solving. Many norms, including moral norms, are in part enforced this way, by the pressure, encouragement, and correction we exert when we operate together. The picture of coordination as enforced by explicit rewards and sanctions ignores essential factors. Besides the intrinsic benefits of coordination, there are demonstrations of how to coordinate, practice in the moves required, and indications of which coordinations others are aware of and likely to fall into line with. (When driving over a section of road under construction, it is not always obvious what now are the right and left sides of the road, so one picks up cues from the other cars. Only as a last resort do they honk at you.)

Fourth, the conventional advice-giving and advice-considering acts are hedged with many conditions. For there are many situations in which pointing out that an alternative hypothesis is supported by available evidence, or even that a set of beliefs is inconsistent, is time-wasting or obstructive. But all conventions are like this. There are times when referring to a cat by "cat" (rather than, say, by "dog," if that is what is called for in a game) is unco-operative, and times when it is dangerous to drive on the conventional side of the road.

To sum this up, I shall say that a pattern of thinking conforms to a rational norm among a group of people when:

- It is the object of a convention of giving and responding to commentary about thinking when circumstances are appropriate
- It is a solution to a coordination problem: everyone is generally better off if everyone generally heeds the commentary
- The commentary centres on drawing attention to relevant content of an N-theory
- Giving the commentary typically evokes the persuasion phenomenon: it gives the recipients the resources to change their own attitudes

The sense in which the norm is a solution to a coordination problem is not that it requires everyone to do the same thing in similar circumstances. (Do not be misled by the diagonal of high values in a textbook normal form coordination

problem.) A coordination problem can require that people do what are intuitively very different, but complementary, acts. I might add that although it does not matter for this discussion, I would prefer not to characterize the situation in terms of the agents' beliefs and preferences. I would rather say that there are vital interests of the agent that are in fact furthered if everyone conforms given that everyone else does. My intention is closer to that of Ruth Millikan in her discussion of the origins of language, or of evolutionary biology: there are vital interests of the agent that are in fact furthered if everyone conforms given that everyone else does (see Millikan; Demeter).

Variety

Standard examples of conventions involve a choice of equilibria. We can drive on the right or drive on the left; call dogs "dogs," "gǒu," or "hunden"; it does not matter as long as we all do the same. Or, more precisely, the advantages of all doing the same outweigh the differences between the alternatives. Perhaps there are no choices in some areas of normative theory. Perhaps there is only one logic, or only one statistics. Perhaps there is a single best way to design experiments, though we do not have it yet. Perhaps even there is only one ethics. So it may seem problematic to speak of convention here.

The conventions, though, are in the giving and not in the content of the advice. So the idea is that the advice-giver has been socialized into a certain style of normative comment, the hearer has been socialized into a certain style of response, and that under suitable circumstances the combination is good for both. Can there be alternatives here?

It certainly seems that there can be alternative ways of presenting the content of such a theory. Perhaps deductive logic is the least plausible case, so I shall concentrate on it. Contrast relatively minor differences such as those between systems in which a principle of deduction can have one, two, or more premises, and those in which each principle works on only one premise. The latter are found in many systems of semantic tableaux, and require one to represent the two-premise *modus ponens* (given p, and *if p then q* derive *q*) in two stages: first one derives the alternation *not p / q* from *if p then q*, and then one notes that its first alternative contradicts the remaining premise, leaving only the second. This has its advantages, but when teaching logic you find that students can grasp one style or another, but not both at once. We can imagine two logic-wielding cultures, in one of which arguments are standardly produced in two-premise form and in the other, in one-premise form. One might be better than the other for general use, but the advantages would be slight in comparison with those of uniformity. Given that everyone else is presenting arguments as connections between single premises, tied together at a higher level with principles such as *reductio* and conditional proof, it would be in an individual's interest to do so too. A similar example would be the use of a logic in which only principles true in absolutely all domains, including empty domains, are valid, so that most derivations require an additional premise of non-emptiness.

More cumbersome, clearer in the abstract, but considerations either way are outweighed by those of uniformity.

A more extreme example would be Christopher Cherniak's imaginary species who find our rules of logic too cumbersome to use, but have their own, which are manageable for them but too cumbersome for us (see Cherniak.) In terms of either one can reconstruct the other, but only by long arguments that strain or exceed the bounds of intelligibility. Logical advice in the style of Cherniak's beings would be a definite alternative to ours, albeit an alternative that we would be ill-advised to switch to. What is not clear is that if everyone—every real human—but one used this system, it would be in the interest of that one holdout to switch to it. We can avoid this problem by considering instead an alternative system of rules that is complicated enough, from a human perspective, to be a real burden, but just manageable if one is forced to use it. There are many such systems; they would have just the right degree of difficulty in terms of half-tutored human capacities. (Of course, it is hard to assess this without a lot of data; thus hard to give an example in detail.) A plausible example would be a system like that of *Principia Mathematica*, in which, because there are no schematic letters, one must use complicated substitution rules to get from axioms involving atomic sentences to derivations with complex instances.

The choice of an N-theory could be yet more different. The standard response to an invalid inference might be a geometrical counter-example, and the N-theory might be a kind of geometry especially focused on this purpose (Euler or Venn diagrams generalized; roughly in the spirit of cylindrical algebra). Or discussion of deductive matters might be made part of probability theory. Persuasion that a conclusion follows might be a special case of the evidence being overwhelming, and inconsistency might be assimilated to probability zero. There would be a different mix of advantages and disadvantages, but no one would be advised to switch to it unilaterally.

A culture could operate with no practices of logical or statistical comment at all. Most human cultures have operated without any. (It is a contested question in psychology whether people who get to conclusions without any help from a normative theory are using their own innate versions of such a theory. See Rips.) So one alternative to any system of logical commentary is no commentary at all. And just as it would be foolish to insist on staying on the right side of the road if others were choosing lanes at random (like pedestrians in England), it would probably be foolish to stick to any logic-presenting convention if others used none.

A wild question is whether a non-standard set of normative principles—logic, statistics, utility—might function well with a suitable complementary set of intellectual virtues. Perhaps a convention of correcting one another's reasoning in terms that would strike us as utterly bizarre might support a science and an economy if generally adopted, given suitable virtues of when and how to give advice in these terms and when and how to heed it. I am going to ignore this question. There is a tamer version: Could an unorthodox variant on standard

normative theory work as well, given general adoption and suitable virtues of compliance? Intuitionistic or relevance logic, for example; Fisherian rather than Bayesian statistics. Perhaps a more manageable issue, but I have no intention of getting involved. Enough for our purposes is the extra credence that the bare conceivability of such a story gives to the idea of conventions of normative advice.

Design of Experiments

People reason well enough without having learned logic, and there is little evidence that teaching them logic improves their reasoning. (Philosophy departments are in the business of teaching "critical thinking." There is an important empirical debate about how much skill is imparted.) But one place where a lore of rational practice deeply affects our practice is in the design of experiments. This lore has grown up slowly over the past two centuries or so, and learning it is an essential part of the education of every scientist. Different sciences teach different versions of the lore, so that for example astronomers and botanists learn different, overlapping, parcels of wisdom. The traditionally core problems of experimental statistics in physics concerned compensations for the results given by an instrument for different observers on different occasions. It was these that led to the doctrine of the "personal equation." The traditional core problems in botany concerned the effects of unknown factors on the development of plants placed in known conditions. It was these that led to Fisher's classic development of the idea of randomized trials. If philosophers are to discuss and help develop doctrines of how we can most profitably think, then one major focus of their interest ought to be the topic of how to set up an experiment and how to interpret its results, given that one is choosing between a set of alternative hypotheses and has given practical objectives, and given constraints on what one can do. I take it to be *the* most important normative lore of thinking in our culture. But few philosophers know enough about it to write helpfully about it as a body of recommendations that shape much of what we do and know. I am no exception: what I say about experimentation is not based on any particular expertise on the topic—as an epistemologist I take this as a rebuke—so the remainder of this section has a definitely amateur quality. (One philosopher who is an exception to this is Kent Staley. See, for example, Staley. And for what I wish I knew more about, see Bailey, Radder.)

The design of an experiment is affected by several competing desiderata. (They are also interlocking and overlapping.) First, the experiment has to be something that can be done in the time allowed with the resources allowed. Second, it has to yield data that can be analyzed by statistical methods that are understood by—or, at any rate, available to—the experimental team. Third, the results of the analysis have to be interpretable as evidence that can adjudicate between hypotheses of interest. It has never been obvious how best to do this. When should we randomize treatments and the subjects for them; how large should samples be; when should we use blind or double-blind methods;

how many hypotheses should we test simultaneously? Attractive answers to some of these conflict with attractive answers to others. Some of these questions obviously are not relevant to some kinds of experiments: you wouldn't double-blind an experiment in astronomy. And, most importantly, they are a very varied and open-ended lot. Some of them are relatively new (e.g., double-blind techniques) and no doubt there are many techniques yet to be introduced. Some are essentially mathematical problems, some turn on human psychology, and some are sensitive to the sources of variation in the natural world.

The lore of experimental design is, I argue, a perfect example of a normative theory of thinking. It is routinely used, unlike deductive logic, in praise and condemnation. We commend people for correct experimental designs, and when we understand them we commend even more innovative and interesting designs. We criticize faulty designs and we try to point out their faults. And this commentary comes with sanctions. We reward with academic and scientific distinctions: promotion and funding. And we punish by non-citation, non-promotion, ridicule, and even lawsuits. But look how selective we are. Different parts of the lore get applied to different experimenters, depending on the discipline they are part of. Parts of the lore that are norms of good experimentation in a given (sub)discipline are applied authoritatively, with sanctions and rewards. And other parts, plus various doses of statistical wisdom, can be presented to experimenters as useful information—relevant truths—to be used or ignored as the experimenter judges. (There is also a norm of taking to heart relevant information given to you by someone who should know. See the next section.)

Within a discipline, the selection of experimental lore to apply as a norm, rather than as useful relevant information, is to a large extent a matter of coordination. If a standard is not generally applied within the discipline then it is not treated as a norm, but can count as a useful truth. So which experimental standards get applied in a discipline? It depends on the nature of the subject matter, the general level of statistical sophistication, the associated practical purposes, and other local factors. Probably the application of some standards is rather accidental, but this results in the coalescence of the coordination on those standards rather than others.

Virtues

If a standard hasn't the force of a norm it can still be mentioned as useful information. Should it be heeded? Often not. Most truths are irrelevant to most projects, and most relevant truths will disrupt most projects if they are just shoehorned in. So a good thinker pays attention to relevant truths at moments when doing so will help. There is a virtue of paying attention to helpful relevant information, and a virtue of ignoring probable truths when it does not seem that the information outweighs the distraction.

There also are virtues of collecting relevant evidence and getting good advice about the design of experiments, and about the amount of evidence one should collect and the attention one should pay to experimental design. Indeed

there are many virtues, more than we can name and certainly more than we ought to remind one another of at most times. *Some* of those figure in norms of rational advice, in that we have conventions of bringing attention to them and of criticizing their neglect. And as with all conventions, there is an arbitrary element. It consists in part of the choice of virtues to name and press on others. In treating some virtues as normative, a culture makes a choice of intellectual strategy, largely invisible but with enormous consequences.

Issues about intellectual virtues are rarely far away in this paper. There is obviously a lot more to say about defining intellectual virtues and about their use in our encouragement of one another's thinking, as well as our selection of which virtues to encourage in a norm-like way. But not here.

Conclusion: The Most Likely Alternative

I have been arguing that we have normative conventions of giving advice about thinking. Normative in that they are in various ways sanctioned, and conventions in that they are in various ways arbitrary. The most plausible alternative to this conclusion is not the claim that there are a priori principles of thought which all thinkers, or at any rate all human thinkers, should adhere to. The problem of finding a defensible meaning for the "should" here pales before that of finding plausible candidates for such principles. Tradition favours the principles of deductive logic. But as I not very originally argued above, "believe conclusions that follow from premises you believe" is just terrible advice. It would have you believing irrelevant and obstructive consequences, relevant consequences whose complexity would obstruct your thinking, and consequences whose probability is much lower than that of any of the premises.

The negative injunction "do not believe contradictions" generates fewer of these problems. So does the negative injunction "do not believe that the Earth is flat." But it is not clear what these ask of us. Searching through one's beliefs on finding that one believes that arithmetic is consistent, to see if "arithmetic is not consistent and penguins cannot fly" would be as silly as inferring "either arithmetic is consistent or penguins can fly." The injunction may require just that when we notice that we believe something which we have reason to believe is a contradiction—and therefore false—we should take very seriously the option of changing the situation. Yes, that sounds roughly right: but it is not exactly instructions on how to think. Still, it does describe the content of advice that it is sometimes worth giving, and which there might even be a proper conventional norm of giving at suitable times.

The most plausible alternative to the norm-convention position is the simple view that there are no intellectual norms. (We might associate this view with David Papineau. See Papineau for a clear early statement.) Some beliefs are true, and some desires are satisfied, effective thinking leads to true beliefs and satisfied desires, and sometimes we can point out to people general facts that will help them get true beliefs and satisfied desires. But none of this adds up to a system of norms. There's no "ought" involved, except case-by-case when

one says to a truth-seeker, "Here's something you ought to take into account." Compare that utterance to the words of the safecracker's assistant, who says, "You ought to make the fuse a little longer, so we can stand around the corner." She need not be committed to a norm of blowing bank safes efficiently (she may be a police agent). She'd just rather be away from the blast.

I do not think it is a trivial matter to refute this alternative. In fact, I think that it is possible that some groups of people operate with no norms of thinking at all. There may be deep general psychological processes operating in them that tend to truth and success. But there may be nothing more normative about these processes than there is about the efficient operations of our circulatory or language processing systems. There may be many occasions on which one such person advices, commends, or criticizes another's thinking. But this may not add up to anything like a system of norms, with general patterns of what is encouraged and criticized, shaping the way we think by its presence.

In contrast, I have been arguing not that we have conventional norms of how to think—to repeat, heading off the easiest misunderstanding—but that there can be conventional norms of how to *evaluate* thinking. And it is consistent with what I am arguing that some people at some times operate outside the influence of any such norms. In fact, I think it is likely that there have been many such people. (What is harder to settle, and more interesting, is whether most people do most of their thinking beyond the influence of any norms.) But "normlessness" is an option on any conventional view: the norms of economic life do not apply to nomads operating in a barter economy, and the norms of the highway do not apply on the speedway or in the parking lot. Norms can be more global or local in their extent, and the norms I am concerned with seem to be strikingly more local than one might gather from what philosophers often say. Consider the very best example, norms of experimental design. It seems evident that we do encourage and discourage our thinking in systematic ways here, and that the ways we do vary from one scientific community to another and from one time to another. They improve, I trust, giving us results that are more reliable and better tuned to the needs and situations of particular disciplines. Our financial conventions and our traffic rules also evolve, generally for the better. As a result, there is no set of conventions of experimental design that is applied to more than a small proportion of experiments carried out in science. (And if we include experiments carried out in everyday life, of course the proportional constituency for any set of conventions becomes even smaller.)

Two main factors drive the force of any convention of intellectual behaviour. First there is the truth. A convention will have little force if it often results in false beliefs. But we can choose which truths we encourage one another to focus on. And second there is the need for coordination. If there is little to be gained from applying pressures on our thinking in a coordinated way, then a convention will not develop. (The coordination can be quite subtle, though. As remarked above, it requires complementary rather than identical actions.) So

conventions of thought to do not allow us to think any way we please if we can persuade others to share our folly. Most such expeditions would lead out into the desert. Nor, I think, do they leave us coordinating on a case-by-case basis, as the anarchic alternative would have it. Rather they allow us to assist one another, keep one another out of trouble, and share methods we have learned. When this is available, is it not something we should go for?

References

Bailey, R. A.
2008 *Design of Comparative Experiments*. Cambridge: Cambridge University Press.

Cherniak, Christopher
1986 *Minimal Rationality*. Cambridge: M.I.T. Press.

Cuneo, Terence
2007 *The Normative Web*. Oxford: Oxford University Press.

Demeter, Tamas
2010 "Mental Fictionalism." Ph.D. diss., Cambridge University.

Elio, R., ed.
2002 *Common Sense, Reasoning, and Rationality*. Oxford: Oxford University Press.

Gibbard, Allan
1990 *Wise Choices, Apt Feelings*. Oxford: Oxford University Press.

Harman, Gilbert
1999 *Reasoning, Meaning, and Mind*. Oxford: Oxford University Press.

Lewis, David
1969 *Convention: A Philosophical Study*. Cambridge: Harvard University Press.

Millikan, Ruth
1984 *Language, Thought, and Other Biological Categories*. Cambridge: MIT Press.

Papineau, David
1999 "Normativity and Judgement." *Proceedings of the Aristotelian Society, Supplementary Volume* 73: 17–43.

Pollock, John
2006 *Thinking about Acting: Logical Foundations for Rational Decision*. Oxford: Oxford University Press.

Radder, Hans
2003 *The Philosophy of Scientific Experimentation*. Pittsburgh: University of Pittsburgh Press.

Rips, Lance
2002 "Reasoning Imperialism." In *Common Sense, Reasoning, and Rationality*, ed. Renée Elio. Oxford: Oxford University Press, 215–35.

Smith, Michael
1994 *The Moral Problem*. Oxford: Blackwell.

Staley, K. W.
2004 "Robust Evidence and Secure Evidence Claims." *Philosophy of Science* 71 (4): 467–88.
Velleman, David
2000 *The Possibility of Practical Reason*. Oxford: Oxford University Press.
Weirich, Paul
2004 *Realistic Decision Theory: Rules for Nonideal Agents in Nonideal Circumstances*. Oxford: Oxford University Press.

AUTHOR QUERIES

1	Please provide the running head for this article.