

# On rationality

*David Hitchcock  
Department of Philosophy  
McMaster University  
Hamilton, Ontario  
Canada L8S 4K1*

**ABSTRACT:** In this address marking my retirement after 46 years as a faculty member at McMaster University, I share ideas about rationality that I acquired during my career—some from others, some from my own investigations. The thoughts concern the relation between myth and rational argument in Plato’s dialogues, moral issues, critical thinking, informal logic, morality, argumentation and computation, logic and critical thinking in medicine, and the existence of God.

**KEYWORDS:** argumentation and computation, artificial intelligence, critical thinking, God, informal logic, medicine, missing premises, moral issues, morality, myth, Plato, rationality

In this address marking the conclusion of my 46 years as a McMaster faculty member, I would like to share some thoughts on a central theme of my research and teaching throughout those years: rationality. I shall frame these thoughts within a rough chronology of my time at McMaster. In this written version of the address, I will **highlight** at first mention the names of the six philosophers and six scientists whose busts form the corbels of the vaults of Convocation Hall, where I am to present excerpts from this address at my retirement reception on June 13, 2014.

## 1. MYTH AND RATIONAL ARGUMENT IN PLATO’S DIALOGUES

I joined the McMaster faculty in July 1968 without having defended my doctoral dissertation—indeed, without even having started working on it. Nowadays such an appointment to a university like McMaster would be impossible, but in those days of rapid expansion to accommodate the baby boomers, universities took what they could get. And I was a known quantity, from my previous work as a McMaster undergraduate.

I spent the first six summers working on the dissertation, defending it in December 1973, at about the same time as I was given tenure and promoted to assistant professor. I chose to write about **Plato**, partly under the influence of my undergraduate teacher George Grant, who once remarked in class that, when he was quoting Plato, he was quoting absolute truth. He was mistaken about that, but I am sure that Plato’s dialogues will continue to be read for what John Randall called their “dramatization of the life of reason”. It is the practice of dialectical discussion that they exhibit that Alfred North Whitehead was referring to when he wrote that “the safest general characterization of the European philosophical tradition is that it consists in a series of footnotes to Plato”.

What puzzled me about Plato’s dialogues was **Socrates’** story-telling. Here was Plato’s main character, his admired friend, the “gadfly of Athens”, who challenged his fellow Athenians, one on one, to state the principles by which they lived their lives and to defend them against Socrates’ attempts to refute them. Such examination, he says, is necessary if one is to live well: “the unexamined life is not worth living for a human being” (*Apology* 38a). Yet Plato has Socrates and

his other spokespeople follow such examinations with fanciful tales: of the fate of the soul after death, of the creation of the world, of cosmic reversal, of a lost city of Atlantis—tales that are explicitly acknowledged to be fictions or at best likely stories. Why do Socrates and the other spokespersons for Plato switch from question-and-answer rational discussion to this mythologizing? What is the relation in Plato's dialogues between myth and rational argument?

At the time the dominant view in the scholarship was that dialectical discussion could only take one so far, and that myth took over where dialectic ran out—that it was a vehicle for communicating and vindicating higher truths that were beyond the reach of reason. In particular, the only book-length study in English of Plato's myths, J. A. Stewart's *The Myths of Plato*, interpreted Plato as a kind of Kantian. **Immanuel Kant** had argued in his *Critique of Pure Reason* (1781) that reason ran into contradictions when it tried to think beyond the limits of any possible experience—for example, about whether the world had a beginning in time, whether the soul is immortal, whether human beings have free will. General principles like the principle that every event has a cause or that the fundamental realities are individual self-subsistent entities that persist through time have legitimacy only as organizing principles for the phenomena that we observe, i.e. as conditions of any possible experience. These principles are impositions of the mind on phenomena, not objective principles of reality. Kant called his reversal of the status of such principles his “Copernican revolution”, an allusion to the substitution by **Nicholas Copernicus** in his *On the Revolutions of the Heavenly Spheres* (1543) of the idea that the Earth revolves around the sun for the received idea that the sun revolves around the Earth.

The neo-Kantian interpretation of Plato's myths got some plausibility from the fact that the myths pronounce on issues that Kant had declared to be beyond the reach of reason, especially the immortality of the soul and the origin of the world in time. But it flatly contradicted the theory of knowledge and the theory of reality that we find in the dialectical parts of Plato's dialogues. According to Plato's theory, the phenomena that we observe are so evanescent and constantly changing that we only have approximative beliefs about them. What we can know are pure forms, stable and unchanging, that are grasped by pure reason. In contrast to Kant, Plato had great faith in the power of pure reason, a faith vindicated by such discoveries of pure mathematics as the discovery in his lifetime that there are only five regular solids—the so-called Platonic solids. Plato does not hesitate to have his characters argue for principles that Kant declared to be beyond the reach of rational inquiry, such as the immortality of the soul and the beginning of the world in time.

What then is the point of the myths, if it is not to communicate higher truths that reason cannot grasp? To answer this question, I developed a set of criteria for extracting from a Platonic myth the core message that it was designed to communicate. Such demythologizing seemed justifiable in view of the fact that Plato's myths are the self-conscious creation of a systematic philosopher, not the unreflective product of a story-telling culture. The results of applying these criteria bore out my initial scepticism about the received interpretation. Sometimes, as in his account of the creation of the world, Plato uses myth as a vehicle for proposing proto-scientific explanations of perceived phenomena, explanations whose mythical and thus tentative status reflects Plato's position that we can have only approximative beliefs about what we observe. More often, as in his myths of the fate of the soul after death, Plato uses myth to communicate principles about the soul, God and the forms that have been supported already by argument, usually in the same dialogue. The myths show how these principles might be exemplified in areas beyond our immediate experience.

The vivid presentation of such a possible exemplification acts as a charm or incantation for the non-rational desires and emotions of the soul, a supplement to the charm of the reason exercised by rational argument. In this respect Platonic myth is not an independent access to a higher reality but a means of reinforcing claims for which Plato's readers have already encountered supporting rational arguments.

## 2. MORAL ISSUES

While I was working in the summers on this dissertation, I introduced in collaboration with my colleague and former teacher John E. Thomas a course that we called "Moral Issues", in which we read with our students philosophical articles on such controversial topics as abortion, euthanasia, war, and capital punishment. We taught this course together for several years in a row, each of us attending the other's lectures—an experience that I recall with great pleasure. John, who is alas no longer with us, became a good friend.

The introduction of this course reflected a growing trend in Anglo-American ethics to break loose from the strictures imposed by the analytic philosophy of the 1950s and 1960s. Under the influence particularly of the Vienna Circle of the 1920s, which held that the only meaningful statements that were not logical truisms were those that could be verified by observation, analytic philosophers supposed that the only role in ethics for philosophy was to discuss the meaning and justification of moral statements, whose status the Vienna Circle had made problematic. Were they expressions of personal emotion, or assertions of personal attitude, or commands? It was no longer the business of the philosopher to propound general ethical principles like Kant's categorical imperative or John Stuart Mill's utilitarianism. Still less was it the business of the philosopher to engage in casuistry, exploring what was the right thing to do in particular cases or how to resolve some controversial moral issue.

The upsurge of philosophical work in applied ethics in the early 1970s was not the consequence of a theoretical reinstatement of the status of moral statements as making meaningful substantive claims. The meta-ethical problems remained. Despite their intractability, philosophers simply moved on. They started thinking and writing about things whose legitimacy was still under a philosophical cloud. John and I rode this wave.

After a few years, the course became so popular that John and I divided it into two, John keeping the topics that concerned personal decision-making, like abortion and euthanasia, and I taking the topics that had a more social and political dimension, like capital punishment and war. John kept the title of the original course, and I called mine "social and political issues". Both courses remain on the books, with substantial enrolments by students from across the university.

John's course evolved into a very popular course in bioethics, in which John was a pioneer in Canada. He and I attended together in the mid-1970s the same summer workshop on bioethics, put on by and at the Hastings Center, a bioethics research institute. For many years, we were both members of the clinical ethics committee of the McMaster University Medical Centre, where we contributed to the discussion of such issues as the hospital's development of a do not resuscitate policy, the ethics of initiating a program of *in vitro* fertilization, the ethics of randomized clinical trials with just one subject, and the criteria to be used in selecting who was to be admitted to the intensive care unit if there were not enough beds available for those who needed them.

It might be asked what contribution philosophy can make to such discussions. Kant supposed that pure reason in its practical function could establish the ultimate principle of morality, from which in combination with information about the facts of a particular case one might hope to deduce what is the morally correct thing to do. Kant noticed that the dictates of morality differ from the dictates of prudence and of the crafts in being unconditional. Prudential norms are conditional: if you want to be healthy, eat moderately. So are the norms of crafts: to get a square corner, use the 3-4-5 rule. Moral norms have no such ifs: do no harm. In Kant's terminology, the commands of morality were not hypothetical but "categorical". Kant supposed that the reason for this unconditional or categorical character was that morality was something that reason dictated to itself, independently of any particular desires that a person happens to have. What could such a self-legislation be like? Since reason is deprived in this self-legislation of a goal that the person wants to achieve, the only basis for giving a rule to itself is the form of a rule: its universality. Hence the first formulation of Kant's so-called "categorical imperative": Act only on that maxim by which you can at the same time will that it be a universal law. By a maxim, Kant meant a subjective principle of action with three components: the circumstances in which it applied, how one would behave in such circumstances, and the goal to be achieved by behaving this way. An example would be the principle: when it is necessary, make a promise that one does not intend to keep, in order to get someone else to do what one wants. Could a person with this approach to life at the same time will that everyone in such circumstances make a promise that they did not intend to keep? No, because the foreseeable effect of such a universal law would be undermining the institution of promise-keeping. If it were a universal law that people who found it necessary made promises that they did not intend to keep, people would not believe promises and promise-making would become ineffective for the intended goal of getting the person to whom one makes the promise to do what one wants.

A second formulation of the categorical imperative, alleged by Kant to be equivalent to the first, is the principle: Always act so as to treat humanity, whether in your own person or in that of another, as an end and never merely as a means. To treat the humanity in another person as an end is to avoid short-circuiting the ability of that person to decide for themselves what to do. Kant's second formulation rules out most obviously deceit and coercion of other people—indeed, any behaviour where one rides roughshod over the other person's wishes. It does not rule out treating other people as a means, provided that at the same time one treats them as ends.

These two formulations connect with recognizable features of our common morality. The first formulation, sometimes called the principle of universalizability, links up with a common way of teaching morality to children: What if everybody did that? It calls to mind the "golden rule", which one can formulate positively as the rule to do unto others as you would have them do unto you, or negatively in Confucian fashion as the rule not to do unto others what you would not have them do to you. Unfortunately, all these appeals to universality have obvious exceptions. The maxim of never being the last to leave a party, in order not to wear out one's welcome, is not universalizable, since it is impossible for nobody to be the last person to leave a party. But there is nothing immoral about this maxim. And there are occasions when it is morally wrong to do unto others what you would have them do unto you; think of sexual desires for example. Likewise, there are occasions when it is morally permissible, or even obligatory, to do unto others what you would not have them do unto you; think for example of a dentist filling a cavity. It is quite difficult to

articulate without generating absurd consequences the core of truth in such appeals to universality.

The second formulation of Kant's categorical imperative, which enjoins treating humanity in one's own person and in that of others as an end, is commonly accepted nowadays in the form of a principle of respect for persons or of respect for personal autonomy. In Kant's formulation, which allowed no exceptions, it has troubling consequences. If a friend takes refuge in your house from a man who is out to kill him, and this man comes to your door asking if the friend is within, is it morally permissible for you to lie to him? Kant says no, and defends this answer in a short article entitled "On a supposed right to lie from altruistic motives". Few would follow him that far. A more controversial exception to the principle of respect for personal autonomy concerns people with impaired cognitive or emotional functioning that we label as "mentally ill". Currently, such people can be forced into treatment only if they are at risk of harming themselves or others. But there is a case for forcibly treating them if their life is going very badly for them as a result of their disorder and would go significantly better after treatment.

The more basic difficulty with Kantian ethics is the way that Kant derives his categorical imperative. He leaves out the social dimension of morality, treating morality as the product of a solitary pure reason giving a law to itself. I am convinced that this is a mistake, although I cannot prove it. Morality is a social institution, an informal universally recognized constraint on the way humans deal with each other and with non-human nature. Although morality needs to be internalized, so that one acts morally out of inner conviction rather than merely out of social pressure, it is fundamentally an institution for social control, an informal analogue of the law.

John Stuart Mill in his *Utilitarianism* (1861) propounded as the foundation of morality the principle of utility, or usefulness, "that actions are right in proportion as they tend to promote happiness; wrong as they tend to produce the reverse of happiness. By happiness is intended pleasure and the absence of pain; by unhappiness, pain and the privation of pleasure." Utilitarianism has attracted many objections, to some of which Mill replies in his essay. The most telling objection, in my view, is that taking promotion of the general happiness as the sole principle of morality justifies the subordination of individual rights to the general welfare. Another telling objection is that our common morality does not impose on us a general duty to promote the positive happiness of other people. If I pass up an opportunity to bring a little more joy into the life of a stranger that I pass on the street, for example by smiling and saying hello, have I done something morally wrong? Most people would think not. Negative utilitarianism, which takes as its touchstone the reduction of harm suffered by others, is much more plausible.

What is the rational foundation of the principle of utility, however formulated? Mill argues that, as a first principle of conduct, it is no more susceptible of proof than a first principle of knowledge, like the principle that equals added to equals are equal. But he offers as "the sole evidence it is possible to produce" in support of taking happiness to be the only thing desirable as an end that people do in fact desire their own happiness. To the objection that some people desire other things than happiness as ends, and not merely as means to achieving happiness, he replies that they desire them as a part of their happiness—thus emptying the concept of happiness of any determinate content. Further, he seems completely unaware of the need to establish that morality requires each of us to do whatever we can to contribute to the ultimate good of all other people—a rather strenuous principle that does not fit the rules of right and wrong conduct that we learned at our mother's knee.

Variations of Kantianism and utilitarianism are the main contenders in contemporary philosophical ethics for ultimate principles of morality. The doubtfulness of any such theory means that philosophers working as ethics consultants play a different role than most expert consultants. Whereas a cardiologist who is consulted about a patient's heart condition can give a well-informed diagnosis and prognosis and recommend the appropriate course of treatment, a philosopher consulted about a moral quandary on a hospital ward cannot pronounce authoritatively on the morally correct course of action. What the philosopher can bring is a certain conceptual clarification and a sensitivity to the nuances of argument. Consider a nurse on a neo-natal intensive care unit who is uncomfortable about the decision to deprive an anencephalic newborn of artificial feeding. Is her discomfort reasonable? Depriving a terminally ill patient of artificial feeding is in a grey area. On the one side, it can be analogized to depriving a terminally ill patient of aggressive forms of treatment that serve only to prolong the dying process, like intentionally not resuscitating a dying patient whose heart has given out. On the other side, it can be analogized to failing to give a patient the basic necessities of life. Beyond highlighting the problematic moral status of the treatment decision in this case, the philosopher *qua* philosopher cannot conscientiously go.

Likewise, in consulting on professional codes of ethics, the philosopher *qua* philosopher cannot go beyond conceptual clarification, elucidation of functional requirements of the profession, articulation of commonly recognized professional obligations, and elaboration of alternative positions on controversial issues along with their accompanying rationales. There is no philosophical right answer, for example, to the question whether the Hippocratic oath taken by graduating medical students, at McMaster and elsewhere, should include an undertaking not to help a woman have an abortion or give anyone a lethal drug, as did the original oath of **Hippocrates**, the father of medicine.

Much later in my career, I came across a very wise book about the nature and justification of morality, whose account I take to be basically correct, and which connects morality to what it means to act rationally, or more specifically to act irrationally. I will discuss this book later.

### 3. CRITICAL THINKING

During my teaching of the moral issues course, one of my teaching assistants, Dale Roy, brought to my attention that the students were having a great deal of difficulty extracting the main point from the articles that they were being asked to read. He thought that it would be beneficial for students to take as a pre-requisite a course that developed their skill at analyzing such extended reasoning. I introduced such a course in the mid-1970s, with the title "Reasoning", and like the moral issues course it attracted a substantial enrolment of students from across campus. It was however not made a pre-requisite for the moral issues course. Subsequently the course was offered as well under a humanities label, at first as half of a first-year course in writing and informal logic, then as a stand-alone second-year course called "Critical Thinking", which for a couple of decades met a requirement of honours programs in the sciences and social sciences, and thus attracted almost 2,000 students a year. A similar course was required in the first year of McMaster's flagship Arts & Science program, from its very beginning in 1981, first in combination with the teaching of writing, and more recently as a separate course. I have had the pleasure of teaching critical thinking in all these incarnations, and in particular the pleasure of team-teaching it with teachers of writing from the English department, Sylvia Bowerbank and Tom Cain, who like John Thomas are no longer with

us. I learned a lot about writing and the teaching of writing from Sylvia and Tom.

An introductory course in critical thinking or reasoning or informal logic—call it what you will—is a valuable service to undergraduate students. Philosophy departments typically assign this course to the most junior member of the department, thinking that anybody with philosophical training can teach it. This is a mistake. Such courses work best, other things being equal, when the instructor has a good background in the specialized scholarship that underlies such an introductory survey. Although students taking such courses are unlikely to teach their instructors anything of theoretical interest, and although such courses provide little opportunity for sharing the fruits of the latest research in a highly specialized field, there is a reward to the instructor in being an instrument for the improvement of thinking skills that will help the students academically, in their everyday life, in their role as citizens of a democracy, and in their chosen career. As well, and importantly, the instructor can nurture, through explicit mention and modelling, the attitude and behavioural dispositions of a “critical thinker”, a person who is willing and able to reflect in a fair-minded and open-minded way on an issue that needs such attention. I have been privileged recently to supervise a dissertation by my doctoral student Ben Hamby on this attitudinal side of the critical thinker. Ben defends his thesis on “the virtues of a critical thinker” on July 9.

The reasoning course that I introduced was different from a standard introduction to logic of the sort that I had taken as a McMaster undergraduate. In response to demands from students of the late 1960s for something that was more relevant to their social activism than the categorical syllogisms that have come down to us from Aristotle or the truth tables of propositional logic, a new type of textbook appeared, one that used examples for analysis and evaluation that were taken from current debates on social and political issues. Some of these textbooks, such as Howard Kahane’s *Logic and Contemporary Rhetoric: The Use of Reason in Everyday Life* (1971), took a fallacies approach, sensitizing students to find in discourse on controversial issues such mistakes of reasoning as distorting the views of one’s opponent (*straw man*), substituting personal attack for argument (*ad hominem*), assuming the point at issue (*begging the question*), and trading illegitimately on a shift of meaning (*equivocation*). My experience with this approach was that students were delighted by it, but that it was highly problematic whether a given piece of rhetoric committed a fallacy and if so which one. Indeed, students became able to find a fallacy in anything. And, having found what they thought the fallacy was, they stopped thinking about the issue under discussion. So I have generally steered away from a fallacies approach in my teaching of critical thinking, letting the fallacies emerge in the context of more positive lessons about how to reason well—the fallacy of *biased sample*, for example, involving a failure to choose a method of sampling that was either genuinely random or likely to produce a representative sample.

The textbook I chose was Michael Scriven’s *Reasoning*, initially self-published but then taken on by McGraw-Hill (1976). To supplement the book, I developed a set of exercises that I sold through the bookstore, exercises that eventually turned into my own textbook, *Critical Thinking: A Guide to Evaluating Information* (Methuen, 1983). Like Scriven, I framed this book as an aid to analyzing and evaluating arguments and presentations of the sort that students will encounter in both academic and non-academic contexts. I developed the acronym OMSITOG for a checklist of steps to take as part of such an analysis and evaluation:

Get an OVERVIEW of the passage.

Clarify MEANING.

Portray STRUCTURE of argumentation, if any.

Check whether INFERENCES are sound.

Evaluate the TRUTH of claims not support by argument.

Consider OTHER relevant evidence and arguments.

GRADE the passage.

These steps formed the chapters of the book. I used it as a textbook for many years, both in the Arts & Science course and in the reasoning course offered to other students. Occasionally I got encouraging feedback indicating that the book and the course had proved useful in later life. For example, one graduate of the Arts & Science program wrote to say that, although he thought that he was done with OMSITOG after his freshman year, it turned out to be very useful in law school.

In the textbook I argued that the most appropriate standards for deciding what to do and what to think are the standards of reason—ultimately, the avoidance of inconsistency and improbability. Relying on the standards of reason, however, does not mean making all our choices and adopting all our beliefs on the basis of explicit reasoning. Such a policy would be paralyzing; we don't have enough time for it. Most of the time we make decisions and adopt beliefs on the basis of intuitive hunches, feelings, tradition, authority, direct observation, habit, and so forth. That is perfectly fine, since those sources are often reliable. It is perfectly fine, too, that we are often unaware of how we came to make a certain decision or acquire a certain belief. If such short-cuts, popularized by Daniel Kahneman as fast thinking in his recent best-seller *Thinking: Fast and Slow*, were generally unreliable, we human beings would be in more of a mess than we are as we make our way through the world. But the short-cuts sometimes let us down. So we need to be on the alert for such failings, and to appeal to reason to correct for them. If we discover that our intuitive hunches are often false, we had better stop relying on them. Emotions are necessary for actually making a decision and carrying it out, as the neuroscientist António Damásio points out through a vivid case study in his *Descartes' Error: Emotion, Reason and the Human Brain* (1994). But emotions can distort or be inappropriate, and reason should step in on such occasions to correct them—something that is more easily accomplished indirectly than directly. The advantage of reason is that, unlike these other sources, it is self-correcting. Reason in this sense is not opposed to experience. Indeed, an important requirement of being rational is to get good data by observation where these are relevant to an issue into which one is inquiring.

It is a legitimate question whether stand-alone critical thinking courses using textbooks like mine do much good. Students get practice and feedback on their arguing and reasoning in the context of their academic work generally. On average, they improve their critical thinking skills as they take undergraduate courses of any sort, typically ending up with scores on standardized tests of critical thinking skills that are half a standard deviation higher at graduation than when they began their undergraduate education. It is quite possible that stand-alone critical thinking courses add no extra improvement. In fact, most university teachers think that they are teaching critical thinking in their courses. Maybe they are, and maybe that is enough.

I had an opportunity to put this hypothesis to the test when I began to teach the large-enrolment service course in critical thinking that met a requirement for science and social science majors. In response to budgetary pressures, the tutorials for the course had been replaced by self-



testing on-line exercises that provided feedback on one's answers—exercises developed by my colleague Jill LeBlanc. The course ran with no tutorials, with classes of about 450 students per section, and with one or two teaching assistants who had office hours and conducted review sessions before tests and examinations. All testing was in multiple-choice format. I was sceptical that teaching critical thinking in this format was effective. So I administered as a pre-test and a post-test a standardized test of critical thinking skills, and discovered to my pleasant surprise that there was an improvement of half a standard deviation in the score, which as a rule of thumb is a moderate gain, large enough to be noticeable. In fact, the gain was greater than that found in classes of 30 students or so in which the instructor graded written assignments and tests.

The reader might be wondering about the absence from my study of a control group. Such wonder reflects possession of an important critical thinking skill, assessing the design of studies of causal effectiveness. My students' improvement by itself shows nothing about the effectiveness of the critical thinking course in improving their critical thinking skills. We need to rule out other possible causes of the improvement, such as intellectual maturation or the other courses that they were taking. Ideally, as in clinical tests of new drugs, one would randomly allocate students to an intervention group receiving the critical thinking instruction and to a control group receiving an alternative treatment not specifically designed to improve critical thinking skills. Such ideals are unattainable in most educational research. As a substitute, one can use as a control group similar undergraduate students taking a full load of courses without taking a critical thinking course. As a stand-in for such a control group, I used the improvement measured in students at the level of my students (2 and above) over one semester—on average, 0.5 standard deviations, or about one-tenth of my students' improvement. So my study showed that the computer-assisted instruction in critical thinking was indeed effective in improving students' critical thinking skills.

Critical thinking is an important educational goal, as I argued in a paper published recently in China (in Chinese) in a journal of higher education. It is worth clarifying the goal and investigating what methods of instruction are most effective in helping students to achieve it. A recent preliminary meta-analysis of 117 studies of the effectiveness of various kinds of instructional intervention on critical thinking skills and dispositions found that one-third of the variation in improvement could be attributed to two factors: the training of the instructor and the basic design of the intervention. Other things being equal, students improved more if the instructor had some training in critical thinking. And, somewhat surprisingly, of four possible basic designs, the most effective, other things being equal, was a unit on critical thinking embedded in a subject-matter course—more effective than stand-alone instruction in critical thinking of the sort I pioneered at McMaster. I only hope that the results of such studies are incorporated in the way critical thinking is taught in the future.

#### 4. INFORMAL LOGIC

Soon after I began to teach the reasoning course, I received an announcement of the "First International Symposium on Informal Logic", to be held in June 1978 at the University of Windsor. This symposium was a career-changing experience for me. I listened with great interest to the opening address by the co-organizers, Ralph Johnson and Tony Blair, on "The recent development of informal logic". They drew to my attention a body of work subsequent to the Second World War

in which philosophers tried to work out an approach to argument and reasoning that differed from the mathematizing symbolic approach that dominated logic in the wake of the work of Gottlob Frege and Bertrand Russell. As an appendix to their paper, they listed some problems and issues in informal logic, thus setting a research agenda for this newly self-conscious sub-field of philosophy, a sub-field that promised to provide the theoretical underpinnings for the new type of introductory course on reasoning. It was exciting to be in on the ground floor, and I began to contribute to the journal *Informal Logic* that emerged from this conference, initially as a newsletter, then as a peer-reviewed scholarly journal. At the Second International Symposium on Informal Logic, held five years later, I was honoured to be elected as the founding president of the Association for Informal Logic and Critical Thinking, which since then has sponsored scholarly meetings three times a year in conjunction with the divisional meetings of the American Philosophical Association.

In my own research, I took as my main problem for investigation number 7 on Johnson and Blair's list: "the problem of assumptions and missing premises", which they formulated as follows:

What exactly is a missing premise? What different kinds of assumptions can be distinguished in argumentation? Which are significant for argument evaluation? How are missing premises to be identified and formulated? Are these just practical and pedagogical questions, or are they theoretical as well?

The kind of missing premise that interested me came to be called a "gap-filler". The idea, going back to Aristotle, was that people often leave unexpressed a premise of their argument. Aristotle explained this curious tendency by the fact that the unexpressed premise is known to the audience, which can supply it for themselves. Leaving it out makes one's argument more concise, and thus easier for the audience to understand. This explanation is still repeated in contemporary work on logic and argumentation.

The problem for theorists, textbook writers and instructors is how to supply the missing premise. It was assumed that its addition would make the argument formally valid, in the sense that the conclusion would then follow logically from the supplemented set of premises in accordance with a purely formal rule of inference. And it should not make any stated premise redundant, just fill the supposed "gap" left by the fact that the argument as stated was not formally valid. But there was more than one way to meet these two requirements. At a minimum, one would add the conditional statement that, if the stated premises are true, so is the conclusion. But this "logical minimum", it was claimed, merely reiterated the claim made in inferring the conclusion from the stated premises. Something more informative, and logically stronger, should be attributed to the argument. But by what criteria should one select from an infinity of possible gap-fillers of this sort?

To take an example, suppose someone defends the inclusion of Lord Kelvin among the six scientists whose busts adorn the west wall of Convocation Hall, by advancing the following argument:

Kelvin discovered the value of absolute zero, so Kelvin was a great scientist.

This argument is not formally valid, so by Aristotle's criteria it has an unexpressed premise. The logical minimum whose addition would make the argument formally valid is the argument's

“associated conditional”: If Kelvin discovered the value of absolute zero, then he was a great scientist. But this logical minimum just reiterates the claim made by the word ‘so’ in the original argument: that the conclusion follows from the premise. What is needed is some informative claim about *how* it follows. One possibility is the so-called “covering generalization” that whoever discovered the value of absolute zero was a great scientist. A more ambitious covering generalization is that anyone who discovers the value of a fundamental physical constant is a great scientist, a generalization that needs to be supplemented by adding the truism that absolute zero is a fundamental physical constant. How to choose between these and other alternatives, especially if one does not have the arguer around to explain how they thought that the conclusion followed?

As I thought about the problem of missing premises, I developed the idea that the gap-filler of an argument that was not formally valid was some covering generalization of its associated conditional, that more than one such generalization was always available, and that there were plausible criteria for selecting which of the unlimited number of covering generalizations to attribute to the argument. At the same time I became aware of a minority tradition in logical theory according to which these arguments had no gaps at all. Like the emperor’s new clothes in the fairy tale of Hans Christian Andersen, the missing premises were simply not there. The argument was complete as it stood.

This idea solved the problem of missing premises much more neatly than the traditional view coming to us from Aristotle. The only reason in this tradition for thinking that an argument had an unexpressed premise was the demand that a complete argument be formally valid, with the conclusion following in accordance with a purely formal rule of inference. But this demand was simply a prejudice. In fact people typically argue and reason in accordance with material rules of inference, rules that are not purely formal but have some content. In the Kelvin argument, for example, the conclusion follows in accordance with the rule: from the premise that *x* discovered absolute zero, you may infer that *x* is a great scientist. The covering generalization of the argument is not an unexpressed premise, but the articulation in statement form of this rule of inference. To support the claim that no premise is omitted in such arguments, I invited the reader of my article on this issue to try a simple phenomenological test: The next time you find yourself drawing a conclusion for yourself from information at your disposal, reflect on your reasoning, see if it was formally valid, and if not consider whether you were aware at the time of leaving something out. Of course, it makes no sense to suppose that one is leaving out a premise when one is reasoning for oneself, since there is no gap between the step from reasons to conclusion and awareness of the reasoning.

I have continued to hold that reasoning and argument can be “materially valid” as well as formally valid, and that material validity is just as good a form of inference as formal validity. In both cases, if the argument has a sound rule of inference, the conclusion follows necessarily from the premises. A rule of inference is sound, I came to think, if and only if the covering generalization that expresses it is true or otherwise acceptable and supports counter-factual instances. For example, the rule entitling one to infer that someone is a great scientist on the basis that they discovered the value of a fundamental physical constant is sound if and only if anyone who discovers the value of a fundamental physical constant is a great scientist and in particular that Avicenna, for example, would have been a great scientist if he had discovered the value of a fundamental physical constant. Incidentally, I am not sure that this generalization is true, and I was reassured on investigation to

learn that Lord Kelvin made a number of important contributions to science besides discovering the value of absolute zero. (I had no similar doubts about the propriety of including **Isaac Newton** among the six scientists chosen for portrayal in Convocation Hall. Newton's *Philosophiae Naturalis Principia Mathematica* (*Mathematical Principles of Natural Philosophy*), from which we get Newton's law of universal gravitation and indirectly the three laws of "Newtonian mechanics", is a signal achievement in the history of science. I was not so certain about the inclusion of **Michael Faraday**, whom I personally would rank below Galileo among the greats in the history of science, but was reassured to discover that among his many experimental discoveries was important pioneering work on electromagnetism.)

I believe that the approach of allowing material validity as well as formal validity can be extended to reasoning and argument whose premises do not necessitate the conclusion drawn but merely make it probable, or establish a presumption that it is correct, or provide some reason to accept it, or establish it as a real possibility worth further investigation. Such weaker inferences are good if the reasoning or argument has a qualified covering generalization that is true or otherwise acceptable and that supports counter-factuals. I have applied this approach to a number of such types of reasoning, including weighing up the pros and cons, reasoning by analogy, extrapolating and generalizing from observed instances, and practical reasoning from a goal to a means for achieving it. There is much more to be done along these lines, and I expect to make such contributions as I continue my research in my retirement.

## 5. MORALITY

In the year 2000, while browsing in the publishers' display at the annual meeting of Canada's learned societies, I came across a book entitled *Morality: Its Nature and Justification* (1998), written by a philosopher I had never heard of, Bernard Gert. Leafing through the book, I found that Gert had spent 30 years developing a description of what he took to be the common morality of humanity: the moral code to which all adult human beings subscribe whose views about right and wrong are not influenced by specialized religious or scientific doctrines or distorted by serious cognitive or emotional dysfunction. The morality he described sounded like the morality that guided my behaviour. I bought the book, read it, and followed through by arranging for Gert to come to McMaster for a week as a distinguished Hooker visiting professor, and by using a condensation of Gert's views, a paperback entitled *Common Morality: Deciding What to Do* (2004), in teaching our introductory ethics course. Gert's account stood up well to critical examination, and I recommend the book wholeheartedly to anyone who wants to deepen their understanding of morality.

Gert claims to be simply describing the morality that humans generally use in deciding what to do and in judging their own and others' behaviour. His description has nothing new to offer, nor does he attempt to give our common morality a philosophical foundation in some highly abstract general principle. In fact, he thinks that all attempts to reduce morality to principles are a mistake. Our morality, Gert thinks, has two components, a set of moral ideals and a set of moral rules. We are rightly praised for living up to a moral ideal, and rightly blamed for violating a moral rule without adequate justification. The moral ideals are vague and general, but involve attempts to lessen harm suffered by others; they are the ideals that inspire people who enter such helping professions as medicine, nursing and social work. The moral rules are more precise. Gert formulates them in

easily memorable fashion:

1. Do not kill.
2. Do not cause pain.
3. Do not disable.
4. Do not deprive of freedom.
5. Do not deprive of pleasure.
6. Do not deceive.
7. Keep your promises.
8. Do not cheat.
9. Obey the law.
10. Do your duty.

Each of these rules receives clarification of its central concept. None of them is absolute; each has exceptions. So each should be read as preceded by the qualification, “unless you have adequate justification to the contrary”. To determine whether there is adequate justification to the contrary, we implicitly follow a two-step procedure. First, we classify a particular action that we have in mind by its morally relevant features, for which Gert provides a comprehensive list. Second, we ask whether the type of action so defined can be publicly allowed. If everybody could publicly allow it, the action is not immoral. If nobody could publicly allow it, the action is immoral. If some but not all could publicly allow it, the action is in a grey area. Whether a person can publicly allow a certain type of behaviour depends on whether in their estimation the consequences of everyone knowing that it is allowed are no worse than the consequences of everyone knowing that it is forbidden. For example, we could not publicly allow that a physician can deceive a patient into giving consent for surgery, because we would estimate that the consequences for the breakdown of trust from everyone knowing that such deception is allowed are worse than the consequences of everyone knowing that such deception is forbidden. This somewhat intricate test fits well with the truism that a type of behaviour that you tell yourself is morally OK is in fact morally questionable if you feel a need to conceal from others the fact that you think it is morally OK.

We can group the 10 rules of this “common morality” into a smaller set. The first five amount to the rule ‘Do not harm’, the types of conduct that they prohibit being the different ways that a person can cause harm. Thus Gert takes our common morality to include the commonly recognized injunction, ‘above all, do no harm’ (*primum non nocere*), which people falsely think is part of the Hippocratic oath taken by physicians. Rules 6 through 8, enjoining promise-keeping and prohibiting deception and cheating, can be grouped under the more general rule, ‘Be trustworthy’. Rules 9 and 10 amount to the rule ‘Do your duty’, if we take our duty to include the duty to obey the law as well as the responsibilities that go with our occupation and our social roles. Gert includes under the rule ‘Do your duty’ the duty to help those in need when one is in a unique position to do so and can do it at little cost. It might be clearer to make this a separate rule, in which case we could boil Gert’s 10 rules down to four:

1. Do no harm.
2. Be trustworthy.

3. Do your duty.
4. Help those in need.

What is the point of this moral system, with its precise set of moral rules and its vague set of moral ideals? Gert explains it as an informal device for humans to reduce the harm that they suffer. He thus works in a tradition going back at least to the time of Plato according to which morality is an agreement neither to harm nor to be harmed. Given this understanding of the point of morality, it would be irrational outside of certain special circumstances for a person to refuse to endorse morality publicly. By an irrational action Gert means an action that causes the agent or someone the agent cares about to suffer harm or be at increased risk of suffering harm, without an objectively adequate reason for doing so. Refusal to endorse morality publicly is in general irrational in this sense, because it puts one at risk of suffering harm at the hands of others who regard one as a dangerous person because of this refusal.

For myself, I have seen no reason to doubt the accuracy of Gert's description of our common morality. On the basis of it, Gert has co-authored *Bioethics: A Systematic Approach* (2nd edition, 2006), which I would be tempted to use as a textbook if I ever get the opportunity again to teach a course in bioethics. Gert's account of the purpose of morality and of the extent to which it coincides with rationality, however, strikes me as only part of the truth, not the whole of it. Trustworthiness, conscientiousness in performing the duties of one's stations in life, and helpfulness to those in need serve more than the reduction of harm. They bring the benefits that are the opposites of the harms that Gert identifies, especially the benefit of freedom. Further, our motivation to obey the moral rules and to live up to the moral ideals is not just the desire to avoid being regarded by others as a danger if they find out that we are ready to act immorally or have no inclination to lessen the suffering of others. Any person who is brought up in a reasonably loving environment comes to feel empathy with other human beings, and even with non-human animals that are capable of suffering and enjoyment. This empathy is a powerful source of morally decent behaviour. Without it, a human being is in danger of becoming a calculating sociopath.

Nevertheless, Gert has clarified an important concept of rational action. He sets out to define rationality in such a way that it makes no sense to question whether one should refrain from acting irrationally. In other words, he is clarifying that concept of irrational behaviour according to which irrational behaviour is behaviour that it makes no sense to perform. Harming oneself or someone one cares about, or risking such harm, makes no sense if there is no adequate reason for doing so. So Gert has put his finger on a sufficient condition for objectively irrational action.

Is it also a necessary condition? It is hard to think of any other condition than causing or risking harm to oneself or a loved one, without good reason, that would make it nonsensical to perform an action. For example, on some accounts of rational decision-making, the rational thing to do is whatever in the situation satisfies one's preferences to the greatest extent. In decision-making under risk, where the outcomes of the various mutually exclusive options one is considering are affected by states of the world to which we can only assign probabilities, the rational action is said to be the one that has the greatest expected utility, or usefulness, where the expected utility of an option is the sum of the products, for each of the mutually exclusive and jointly exhaustive causally relevant states of the world, of the probability of that state and the utility of the option's outcome given that state. For example, suppose that I am deciding whether to take an umbrella when I go for

a walk, knowing that the forecast is for a 30% probability of rain. If I take the umbrella, I have a 70% chance of having the inconvenience of carrying it without any compensating protection from rain but a 30% chance of getting the protection along with the inconvenience. If I leave the umbrella at home, I have a 70% chance of keeping dry and avoiding the inconvenience of carrying the umbrella but a 30% chance of getting wet. Let's suppose that I can assign a value to each of the four possible outcomes on a scale of 0 to 100. (There are techniques for discovering what people's so-called "utility functions" are and assigning such numbers to their valuations of outcomes.) I rate getting wet at 10, inconvenience with no benefit at 30, inconvenience with the benefit of being kept dry at 50, and keeping dry without inconvenience at 80. Then the expected utility of taking the umbrella is the sum of (.7 times 30) and (.3 times 50), or the sum of 21 and 15, which is 36. The expected utility of leaving the umbrella at home is the sum of (.7 times 80) and (.3 times 10), or the sum of 56 and 3, which is 59. So the rational thing to do is to leave the umbrella at home. The strategy of maximizing one's expected utility can be shown to maximize the satisfaction of one's preferences in the long run. In the example, it fits our intuitive judgment that whether to take the umbrella depends on how likely it is to rain, how inconvenient I find it to take the umbrella, and how much I dislike getting wet in the rain.

Is it irrational to choose an option in decision-making under risk that does not maximize one's expected utility? Hardly. Speaking for myself, I cannot recall a situation where I made a decision by calculating the expected utilities of the options I was considering. It is too much bother, and the numbers assigned to the possible outcomes give an artificial precision to the exercise. The judgment in such cases has to be qualitative, especially if the decision is an important one. There is a story that a world-renowned expert in decision theory was struggling with the question whether to accept an offer of an attractive position at another university. His colleague is supposed to have said to him, "You're a world expert on decision-making. Why don't you just calculate the probabilities and utilities, and do whatever maximizes your expected utility." "You don't understand," came the reply. "This is serious."

Anyway, there is nothing irrational in choosing less when you could choose more. We human beings have limited time and limited resources, and it is reasonable in most decision-making situations to do no more investigation than is needed to get a satisfactory outcome—a strategy that Herbert Simon christened as "satisficing" as opposed to optimizing. There is a lot to be said for satisficing. Aside from the computational effort that might be required to maximize or optimize, there is the effect of such maximizing on one's future desires. What is important in life is to like what one gets (as well as what sort of person one is and what one does), not to get what one wants. Stimulation of one's desires increases the probability of future dissatisfaction, and makes one more a hostage to fortune. Better from this point of view to keep one's desires moderate, which may mean choosing less rather than more.

Thus Gert's conception of irrational action seems to capture what we mean when we think of an irrational action as one that it would make no sense to perform. It provides a basis for psychiatric diagnosis of someone as suffering from a mental disorder because they are behaving irrationally, if they habitually do things that they believe are harmful to themselves or to someone they care for, without adequate subjective reason for doing so. The theory of rational decision-making under risk provides no such basis—it would be ridiculous to recommend psychiatric treatment for someone who was habitually not maximizing their expected utilities.

It is important to recognize that actions that are objectively rational in Gert's sense are not necessarily rationally required. An action is objectively rational if it is not objectively irrational to perform it, i.e. it does not cause or risk harm to oneself or a loved one without adequate reason. Most such actions are rationally permitted but not rationally required. An action is rationally required if and only if it is objectively irrational not to do it—that is, failing to do it causes or risks harm to oneself or a loved one, without adequate reason. The vast majority of courses of action open to us are neither irrational (rationally forbidden) nor rationally required. The realm of rational behaviour contains many types of mansions.

## 6. ARGUMENTATION AND COMPUTATION

In the same year that I picked up Gert's book on morality at the conference of the learned societies, I participated in a week-long invitation-only workshop in Scotland on argumentation and computation. The task of the 22 invitees to this workshop was to write a book during the week. We were divided into five groups, each responsible for a chapter. Each group included both philosophers and specialists in artificial intelligence. I was assigned to a group working on a chapter on practical reasoning, for which I wrote in advance a lengthy review of the current state of philosophical thinking on this topic. This review formed the starting-point of the chapter that the four of us eventually produced, on decision support systems: computational systems that help people make decisions in their professional capacity, such as physicians, lawyers and airplane pilots.

In our chapter of the book, which came out under the title *Argumentation Machines: New Frontiers in Argumentation and Computation* (2003), we recommended that decision support systems be used for the guidance of agents, singly or in groups, in deciding in a wide range of domains what is to be done. Whatever the system, it should be based on argumentation, and transparent in that respect to any user. The system should use a variety of argumentation schemes and techniques. It should reflect the richness of quality argumentation, and should use the techniques appropriate to the domain in which it gives advice. There should be an open-ended approach to advising, and users should be able to deliberate jointly with the system about advice and how it is generated. The interactive interface between agents and the machine should facilitate the giving of advice and the joint activities of system and agents.

The invitation for me to attend this workshop resulted from work I had done on the computational architecture for a rational agent developed by the American philosopher John Pollock. Pollock's model of practical rationality assumes a much richer psychology than the traditional analysis, going back to Aristotle, of practical reasoning as combining a belief with a desire. On Pollock's account, practical reasoning, understood as having the function of making the world more to its possessor's liking, requires seven distinct types of states: beliefs, situation-likings, feature-likings, intentions (construed as adoptions of plans), and three kinds of desires (primitive, instrumental, present-tense action). This more complex psychology is the basis of an extremely complex and subtle computational architecture for a rational agent, for whose construction much sophisticated thinking and refinement of originally attractive but over-simple ideas was required. Complex as it is, Pollock's architecture is incomplete in at least three important respects. First, it is solipsistic, in the sense that there is no provision for verbal input from, or verbal output to, other autonomous rational agents, still less for back-and-forth discussion, whether argumentative or non-



argumentative. Second, it is egoistic, in that the function of the entire system is to make the world more to the liking of that system itself, without regard (except instrumentally) to whether its actions make the world more or less to the liking of other systems which have situation-likings and situation-dislikings; morally speaking, Pollock's rational agent is a monster. Third, it is unsocial, in that it does not (and cannot) belong to any groups of autonomous rational agents with governance structures for making decisions about the actions of the group; it is a citizen of no country, belongs to no professional associations, owns no shares in any joint-stock company, has no immediate family, does not belong to a recreational bridge-playing group, etc. A really comprehensive system of rational guidance for human action would have to remedy all three of these lacks.

The workshop in Scotland led to fruitful collaborations with the two members of our group from the field of artificial intelligence. I collaborated with Peter McBurney in developing a framework for a formal system for deliberation dialogue, construed as a discussion among two or more participants about what to do, where doing something is understood in a broad sense as including the adoption of policies and plans as well as performing particular actions. The framework allows discussants to raise many types of considerations, including pros and cons, general norms, means-end relationships, and moral judgments. It has been implemented computationally in a decision support system for physicians in Spain deciding on which patients will get organ transplants that become available. The empirical research on how they made their decisions turned up 50 different considerations that they took into account, each of which was incorporated in the formal system for deliberation dialogue that underlay the software.

The other member of our group from the field of artificial intelligence, Bart Verheij, proposed that we put together a collection of new articles on a model of argument proposed by the English philosopher of science Stephen Toulmin, in his influential book *The Uses of Argument* (1958). This model was part of the inspiration for my own work challenging the myth of missing premises. Toulmin called the generalization that licenses an inference from one's data to one's conclusion a "warrant", and he allowed that warrants, and thus the inference to one's conclusion, could be qualified by such terms as 'probably', 'presumably', or 'possibly'—in which case the inference or the conclusion was subject to rebuttal even if the data went unchallenged. His model has turned out to fit the way people actually reason and argue better than deductive formal logic, which however has its own legitimate place in understanding reasoning and argument. Bart and I issued a call for papers in the various fora open to us, received a large number of submissions, had them peer-reviewed, and eventually accepted 24 articles which became chapters of our co-edited book, *Arguing on the Toulmin Model: New Essays in Argument Analysis and Evaluation*. These articles came from 27 scholars residing in 10 countries on three continents, working in disciplines ranging from artificial intelligence to philosophy to speech communication. They addressed Toulmin's model from the perspective of their contemporary work, not from a purely historical perspective. They thus testified to the abiding influence of Toulmin's elegantly simple model.

Another result of the contacts made during the Scottish workshop was an invitation to give a keynote address to a workshop on argumentation in multi-agent systems. The paper, "Instrumental rationality", subsequently appeared in Springer's Lecture Notes in Artificial Intelligence series, as part of a collection of revised, selected and invited papers from the workshop. Philosophers tend to take the concept of instrumental rationality as quite straightforward, consisting simply in finding an effective means for achieving a posited goal. In contemporary analytic epistemology, for example,

there is a vigorous debate about whether to construe the rationality of a belief as its instrumental rationality in contributing to the achievement of some epistemic goal, such as grasping the truth. Debate centres on whether epistemic rationality can be reduced to practical rationality in this way and on how to characterize the epistemic goal, not on the intricacies of instrumental rationality. Instrumental rationality is however more complicated than most philosophers acknowledge. The scheme for solo reasoning from end to means that I developed in this paper includes an initiating intention to bring about some goal, as well as five types of premises: an immediate means premise, an achievability premise, a permissibility premise, an alternative means premise, and a side effects premise. All six components are subject to challenge and may be in need of justification. Further, the concluding intention to bring about a means to the intended goal may itself trigger a new application of the scheme, in which the intended means becomes an intermediate goal.

Many initiatives in the artificial intelligence community aim to facilitate interactive argumentation, on the Web and elsewhere. Anyone who reads the comments posted at the end of news articles and opinion pieces on the Web will soon recognize the need for improvement in the quality of online discussion and debate on contemporary controversial issues. An ideal motivating many of the researchers in this field is the development of a thriving online deliberative democracy. To what extent this ideal will be realized should become clear in the next years and decades, as well-designed frameworks for argumentation go live on the Web.

## 7. LOGIC AND CRITICAL THINKING IN MEDICINE

In the mid-2000s, I was approached by a clinical epidemiologist, Milos Jenicek, for a review of some chapters on logic that he had written for a textbook on the foundations of evidence-based medicine. As a sequel, Milos proposed to write a book on logic for physicians, and invited me to collaborate. The book that resulted, which the American Medical Association Press published under the title *Evidence-Based Practice: Logic and Critical Thinking in Medicine* (2005), was primarily the work of Milos, who wrote the entire first draft before I looked at the manuscript. I persuaded him to use the Toulmin model as the logical framework for the book, and it worked well with his medical examples. I also reworked extensively a chapter on critical thinking, explaining the development of this concept, identifying its component skills and attitudes, and articulating a seven-component checklist somewhat like the OMSITOG approach of my earlier textbook:

1. Identify and analyse the problem.
2. Clarify meaning.
3. Gather the evidence.
4. Assess the evidence.
5. Infer conclusions.
6. Consider other relevant information.
7. Reach an overall judgment on the problem.

This checklist is more constructive and less reactive than the OMSITOG procedure, but otherwise similar. We illustrated it with reference to the problem of how modern medicine should respond to the challenge raised by so-called “complementary” and “alternative” medicine. Our judgment at the

end of a process of reflection on this problem was that claims by proponents of complementary and alternative medicine should either be appraised by the standards of evidence-based medicine or relegated to the realm of faith and belief.

Our textbook was part of an effort to incorporate training in critical thinking in the education of medical students and nurse practitioners. Since the textbook appeared, we have collaborated with an academic physician who specializes in emergency medicine, Pat Croskerry, in advocating that critical thinking be made an explicit part of medical education. Croskerry, with a background in cognitive psychology, is particularly concerned about the way in which cognitive biases can cause physicians to make errors of diagnosis, especially in contexts like the emergency room and the office of the general practitioner where they see patients with a wide variety of complaints. A particular danger is so-called “confirmation bias”, which leads people who have in mind a certain hypothesis as the explanation for some phenomenon to ignore evidence that counts against it. This tendency is very strong, even in people who are not strongly committed to the explanation that they have thought of. Both empirical psychological studies of reasoning and historical-textual studies have found that the most common source of reasoning errors is premature closure on a favoured conclusion that leads one to ignore or underplay evidence pointing to an alternative conclusion. There is useful ongoing research on how to identify and overcome this bias, which is a source of mistakes not only in medical diagnosis but also in detective work and in intelligence work, and indeed anywhere one is trying to find the correct explanation of some puzzling phenomenon.

Contemporary research on such biases has a precursor in the four “idols” identified by **Francis Bacon**, who was both Lord Chancellor of England and a visionary defender of empirical natural science and its potential for improving the lot of humanity. Bacon argued that the human mind was not a blank slate on which experience could write truths, but came with false beliefs that it had to get rid of in the process of learning. They were of four types: idols of the tribe, idols of the cave, idols of the marketplace, and idols of the theatre. Idols of the tribe are false beliefs rooted in human nature, such as the tendency to pay more attention to positive instances of a general principle than to the exceptions that falsify it, as when people notice that something they dreamed about comes true but ignore the cases where what they dreamed about did not come true. Idols of the cave, named after Plato’s allegory in his *Republic* of human education as like the experience of prisoners chained in an underground cave, are false beliefs acquired by custom and the peculiarities of our upbringing, such as the superstitions that Friday the 13th and days of a full moon are unlucky—a superstition not noticed in setting the date of the reception marking my retirement (on Friday the 13th, on the day of a full moon). Idols of the marketplace are false beliefs imposed on us by words, as when people assume without thinking that the ministry of defence, formerly called the ministry of war, has something to do with defending our country; it would more appropriately be called the ministry of attack. Idols of the theatre are false beliefs coming from dogmatic philosophical systems and received but tenuous scientific theories, such as the view still current in Bacon’s time and inherited from Aristotle that heavier bodies fall faster than lighter bodies, since weight is a tendency to fall downwards. We can still learn from Bacon’s taxonomy of idols.

## 8. GOD

In 2004 I was asked to give the so-called “proemial lecture” at the beginning of the academic year

to the faculty and graduate students in the joint doctoral program in philosophy that my department offered in conjunction with those at Guelph and Wilfrid Laurier. As my topic, I chose a perennial philosophical question: Does God exist? I tried to approach this question in an open-minded way, without any pre-conceived assumption about what God would be like if such a being existed. So my first task was to clarify the meaning of this word ‘God’? What characteristics would one need to have in order justifiably to be called ‘God’? What does it take to be a god?

The answer, I thought, had to come from religious observances. A god would have to be an appropriate object of religious worship, which I took to be a kind of devotion that held up its object as an ideal to be emulated. Absolute perfection was not necessary. But a god should be active in some way, with a cognitive and emotional life that would make a human who lived that sort of life without flaws in those respects.

Having clarified what we were looking for, I then considered whether we could find some actual entity that fit the bill. I looked at the physical universe as a whole, at what if anything is beyond the physical universe, and at parts of the physical universe.

Some thinkers—such as Plato in the *Timaeus*, the ancient Stoics and Spinoza—have held that the physical universe as a whole is divine. Spinoza, for example, refers to the single substance that comprises all of reality as “God or nature” (*deus sive natura*, *Ethics* I) and refers to the eternal aspect of this reality as “nature naturing” (*natura naturans*). But advances in astronomy have disabused us of the idea that the physical universe as a whole is a single living entity with a governing soul that is divine. Spinozistic pantheism is not a defensible position.

As to what transcends the physical universe, a transcendent god would have to be a mind without a body. The idea of such a being gets its purchase from a dualistic conception of a human being as composed of a material body and an immaterial mind or soul, a mind or soul that can continue to exist after the person dies and is perhaps immortal. The workings of the human mind have however turned out to be inextricably tied to brain processes and through them to the human body as a whole. A human mind without a body makes no sense. So the postulation of a divine being as an active being with a cognitive and emotional life but without a body makes no sense. To act, to think and to feel require organs.

It remains to consider parts of the physical universe. Various religions have worshipped parts of the physical universe, such as an ear of corn or the sun. These things certainly exist, but do not meet the requirements for being an appropriate object of religious worship. The sun, for example, on which we depend for our very existence, is just a huge mass of hot gases undergoing constant thermonuclear reactions, with no cognitive or emotional life.

What about the possibility that some particular human being is divine? All but one of the world’s major religions make it clear that they do not regard their founder—Abraham, Lao-Tse, Siddhartha Gautama, Vardamana Mahavira, Mohammed, Guru Nanak Dev Ji, Bahá’u’lláh—as divine. The exception is Christianity, which claims that its founder, Jesus of Nazareth, is divine. There is good reason to think that Jesus existed. Was he without flaws? The Gospels represent him as demanding all-consuming devotion, thus promoting an exclusivism that has been responsible for enormous suffering by millions of people over the subsequent centuries in such contexts as the Crusades and the Inquisition. I judged this demand for all-consuming devotion to reflect a cognitive flaw. When I gave a version of this talk some years later to a Christian apologetics group at Redeemer University College, one student reasonably objected that, if Jesus really was the Son of

God as he claimed, then his demand for all-consuming devotion was not a cognitive flaw. I would respond by arguing that the transcendent Father to whom Jesus refers does not exist and that Jesus could not have had enough evidence to make him as confident as he is reported to have been of the existence of that Father. He would thus have violated a fundamental principle of the ethics of belief, that one should proportion one's confidence in a belief to the strength of the evidence that one has for it.

The consideration of the Christian claim that Jesus is divine opens up the question of whether some other human being is divine. Could a human being be a god? Such a person would have to be without cognitive, emotional or (I would add) moral flaws. That is a high demand, even if the absence of flaws is not a counsel of perfection. A flawless knower, for example, need not be omniscient or even, I would suggest, infallible. What is required are unshakeable mental habits of proportioning one's confidence in one's beliefs to the accumulated evidence at one's disposal and of readiness to seek out and consider dispassionately new evidence that might lead one to change one's mind. That is, one must have the dispositions of an ideal critical thinker; no more is needed. Similarly, a flawless emotional life does not require that one's feelings are in accordance with the actual situation that stimulates them, only that they accord with what one reasonably believes the facts to be. A flawless moral life does not require supererogatory or saintly virtue; it is enough if one does not violate the prohibitions and prescriptions of the moral rules that constrain every human being, even if one fails sometimes to take advantage of opportunities to perform acts of kindness that are morally praiseworthy. A human being might well be without flaws in these senses.

However, to worship such a person, we would have to be reasonably certain that this person had no flaws. About those who are now dead, we do not have enough evidence to reach such a conclusion. We are certainly well aware of our own cognitive, emotional and moral flaws—at least I am of mine. So our predecessors probably had foibles too. As to people now alive, there is a danger in worshipping someone who is discovered to be so far without flaws. Such a person might develop flaws after being made an object of worship. In fact, being treated as a god in one's own lifetime might well induce a certain *hubris* and arrogance—which is an emotional and moral flaw.

So the upshot of my investigation was that the physical universe as a whole is not divine and there is no god transcending the physical universe. There might be human beings who were sufficiently flawless to deserve our worship, but we cannot be confident enough of the flawlessness of any particular human being to set them up as a god. For practical purposes, there is no god.

Two years later, having heard of this talk, the local Campus for Christ club asked me to debate the question “Does God exist?” with the Rev. Joe Boot, a Christian apologist who was making a tour of Ontario universities. The organizers massively promoted the event, for which they charged \$2 admission. They filled the biggest lecture theatre they could find, with 400 students, and turned another 200 away. For a time a video of the debate was available on the Web, on a site with ‘doesgodexist’ in the URL, but the site seems no longer up. I was pleased to see such an interest in a fundamental question, and also positively impressed by the care taken by the organizers to make the debate fair. Anyone arguing for the negative in a debate on this question sponsored by a religious group has a fear that they are being set up. I was reassured when the first question asked of Rev. Boot came from a student who identified himself as a Muslim and asked why Rev. Boot, who had argued that God was trinitarian, did not accept the revelation of Mohammed. Not a setup, obviously.

My reflections on whether there is a god did not take the standard form of considering

standard arguments for and against the existence of a god. Joe Boot, for example, ran the changes on a form of argument that postulated the existence of God as a necessary condition for the objective status of morality, for the truth of logical principles, and for the correctness of the scientific method. This argument strikes me as postulating something that is less sure as a reason for believing something that is more sure. Morality, as I have already remarked, can stand on its own as a human social institution. As to logic, nowadays there are many rival systems of logic, and one's choice of a system may depend on one's purposes; the days of the one true logic are over. Other standard arguments for the existence of a god, such as the argument that the universe must have a first cause or that the obvious purposiveness in the design of living organisms points to an intelligent designer, have been undermined by the progress of scientific investigation. The question of what led up to the big bang from which our universe originated is a scientific question, not a theological one. The purposiveness in the design of living organisms can now be explained by the theory of evolution developed from the ground-breaking contribution of Darwin in his *On the Origin of Species* as the product of natural selection operating in a physiological framework of what Darwin called "descent with modifications", i.e. inheritance of the genetic basis of each organism from its parent or parents, with occasional modifications. In fact, the theory of evolution has put the attribution of purposiveness on a scientific basis. We know scientifically that a certain organ or behavioural trait has a certain purpose if we know that it was selected historically for that purpose.

## 9. SUMMARY

Rationality has been a main theme of my research and teaching during my 46 years as a faculty member at McMaster. I joined the faculty while still working on my doctoral dissertation, which concerned the relation between myth and rational argument in Plato's dialogues. In it I argued that, contrary to most of the then current scholarship on Plato's myths, Plato did not suppose that myth was a vehicle for the communication of higher truths that were beyond the reach of rational argument to grasp. Rather, it was generally a way of reinforcing, in the captivating framework of a story about matters beyond human experience, truths for which arguments had already been given. In the *Timaeus* Plato used myth as a vehicle for conveying a proto-scientific astronomy, physics, chemistry, biology and physiology—the mythical framework reflecting the tentativeness and approximation that Plato attributed to all beliefs about phenomena that we observe.

During those early years, I joined with my colleague John Thomas in introducing to the undergraduate curriculum the philosophical treatment of issues in applied ethics—a change that reflected a general movement in Anglo-American ethics away from exclusive attention to meta-ethical issues of the meaning and justification of moral claims. Moral, social and political issues were not just matters of opinion, but matters to which philosophical tools of conceptual analysis and argument could make a real contribution. That contribution is not like that of a cardiologist brought in as a consultant to diagnose a patient's heart condition and recommend a treatment. The philosopher cannot pronounce definitively *qua* philosopher on what is morally right or wrong in some controversial particular situation, but can only clarify concepts and marshal relevant arguments on various sides of the issue.

Prompted by problems my students in the moral issues course were having in getting the main point in the articles they were assigned to read, I introduced a course in reasoning, whose goal

was to improve the ability of the students to analyze and evaluate the arguments and presentations they were going to encounter, in their academic work, in their future careers, in managing their life, and in their engagement as citizens of our democracy. Teaching this course prompted me to write a textbook entitled *Critical Thinking*, framed by a seven-step procedure for thinking critically about some issue or claim. In my textbook, I defended reliance on the standards of reason—namely, the avoidance of inconsistency and improbability—as the only self-correcting basis for making decisions and adopting beliefs.

Attendance at the First International Symposium on Informal Logic in 1978 prompted me to investigate the so-called “problem of missing premises”. I eventually worked out that the premises were not in fact missing, and had in fact never been present. The assumption that arguers omit premises of their arguments was due to a mere prejudice, that the only way a conclusion could follow from premises was formally or logically. In fact, it can also follow materially. Over a number of decades I have worked out a theory of this sort of material consequence. I have also applied it, with appropriate qualifications of the material inference-licenses involved, to the evaluation of a number of forms of non-conclusive reasoning, such as weighing up the pros and cons, reasoning by analogy, extrapolating and generalizing from observed instances, and reasoning from end to means.

At a publishers’ display one year, I came across work by the philosopher Bernard Gert describing what he took to be the common morality of humanity. I believe that his description is accurate, and am pleased that he was able to come to McMaster to deliver a series of lectures in ethics and bioethics, and also that I was able to use a paperback version of his account of the nature and justification of morality in our introductory ethics course.

An invitation to collaborate with specialists in artificial intelligence on a book chapter on practical reasoning led, besides the chapter, to joint authorship of a framework for deliberation dialogue, co-editing a collection of original papers on a model for argument analysis developed by Stephen Toulmin, and sole authorship of a paper proposing a scheme for instrumental reasoning from end to means. It has been a rewarding experience, unusual for a philosopher, to contribute to work that promises to have useful applications.

I have been privileged as well to collaborate as second author with Milos Jenicek on writing a textbook on logic and critical thinking in medicine, as well as on some subsequent articles.

In a proemial lecture in our doctoral program, I explored the perennial philosophical question: Does God exist? I construed this as the question whether there exists in reality a being who deserves religious worship, which I argued was appropriate only for an agent that was without cognitive or emotional flaws. I argued that the physical universe as a whole was not such an agent, that it made no sense to suppose that there was such an agent transcending the physical universe, and that such parts of the physical universe as the sun were not such agents. Although theoretically a human being could deserve to be religiously worshipped, in practice we cannot be sure enough that any particular human being is without cognitive or emotional or moral flaws to be justified in setting them up as a person to be worshipped. For practical purposes, then, there is no god.