

Realism in Ethics

An interview with Peter Railton

by Darlei Dall’Agnol and Hans-Hürgen Link

Ethic@: Would you please introduce yourself to the readers of *Ethic@* and tell us why you chose to study philosophy and what your main interests are?

Railton: I have taught philosophy for a good many years at the University of Michigan, Ann Arbor, having come here not as an ethicist, but as a philosopher of science with some interest in moral and political philosophy. Given the strong tradition in ethics at Michigan, and the interest and importance of questions about morality, I soon found myself devoting more and more time and research to ethics. As a philosopher of science, I was particularly drawn to issues about the foundations of ethics, the place of ethics in a natural world, the nature of objectivity, and the justification of moral theories.

My first experience of philosophy came as a young teenager, when I began to worry about free will. I read the existentialists, and, enlightened but not satisfied, felt driven to formulate my own views. Fortunately, entering Harvard College in 1968 enabled me to explore other philosophical traditions, including analytic philosophy. The compatibilist approach to freedom of the will struck me as the first real promise of a solution to what had previously seemed to me an inescapable paradox. I eventually concentrated in philosophy—though, in truth, most of my time was devoted to political causes. After graduation, I felt I should explore

the world of work rather than go on in academics. A year spent in a factory, another in an office, and another as a commercial fisherman convinced me that philosophical questions are entirely real—they are raised by every life and every form of life. My desire to do philosophy was stronger than ever, and so I returned to graduate school at Princeton, and then had the good fortune to be able to go on to a career in philosophy.

Ethic@: Let us then start with a methodological question: What kind of analysis have you practiced in your ethical investigations? What are its advantages over the other approaches?

Railton: I am somewhat suspicious of the notion that philosophy has a distinctive, analytic method. It seems to me—as it did to Hume and Aristotle, as far as I can tell—that philosophy is inquiry in the most general sense, not simply the analysis of concepts. Properly philosophical questions range from empirical matters (say, the nature of human psychology or its place in the natural world) to issues in metaphysics (say, the nature of the self or of fundamental physical reality) to issues about the concepts we use (say, the concept of freedom or objectivity). So in my research I often combine issues from these different domains. The advantage of such an approach is that it does not confuse understanding

a concept with arriving at substantive truths about how things are or what we ought to do. These latter are, as Kant insisted, synthetic, and cannot be the product of analysis alone. The disadvantage of such an approach is that it is very hard to practice well—one cannot simply substitute empirical results for philosophical reflection, for example.

Ethic@: What role does the “idealized self” exactly play in your naturalistic approach to ethics?

Railton: Suppose you want to understand the nature of justification. One way to begin is to ask, “How is it that we go about challenging or defending our beliefs and values?” We then can study which processes are involved, and how these processes gain authority, step by step. For example, when very young, I strongly desired to drive racing cars—I spent inordinate amounts of time reading about automobile racing, thinking about it, etc. But I had never driven anything other than a go-kart. Once I started driving, I began to realize what driving is really like. It was not nearly as thrilling as I had imagined, nor did I find I made great progress in becoming an excellent driver. Other things proved to interest me more, and engage me more deeply. So my desire to race cars diminished markedly. I think of this as a process of *learning what I really want*. My later self was far from ideal, but it was closer than my earlier self, and I think it had much more authority regarding what to do with my life. So the “idealized self” is simply the extrapolation from one’s actual, more informed and experienced actual self—something perfectly at home in the natural world. Since we understand

how each step of idealization can contribute greater authority, we understand why appealing to an “idealized self” might have normative authority, even if we never ourselves reach this limit.

Ethic@: It is interesting to note how much convergence there is nowadays among metaethical views. For instance, your kind of realism, or even Boyd’s and other Cornell Realists, has much in common with expressivists such as Gibbard or quasi-realists (e.g., Blackburn). What exactly are the differences? Is it possible to come even closer?

Railton: This is a very large question! As you know, by the time philosophers have been working away at an issue for decades, positions that began in stark form have become very elaborate, and have had to evolve in the face of similar critical pressures. So characterizing the differences becomes much, much harder. When I wrote “Moral Realism” in the early 1980’s, moral realism was widely considered to be an *obviously* mistaken position. Now even expressivists like my colleague Allan Gibbard want to be claim that their view is a form of realism. Yet this is not because we have agreed on the fundamental issues. Rather, the *fundamental* differences remain but as the views have been articulated, each has attempted to accommodate the same data. It is rather like evolution. Very different species present in similar environments may evolve under the pressure of natural selection to have quite similar features appearances—such as the bat wing and the bird wing, or the tail of a tuna and the tail of a porpoise — even though the underlying

physiology is profoundly different.

The key issue dividing cognitivists like myself and the Cornell School, on the one hand, from expressivists, on the other, is the way each sees the phenomenon of moral deliberation and judgment. For expressivists, the primal form of coming to a moral opinion is akin to taking sides—one isn't attempting to discover an independent *fact*, but rather is resolving one's feelings, attitudes, loyalties, or commitments. It is hard to deny that this captures an aspect of moral deliberation and judgment. And from this starting point, as my colleague Allan Gibbard has shown, an amazing amount of moral thought and discourse can be accommodated. But for cognitivists, the primary character of coming to a moral opinion *is* like attempting to discover a fact—one that is independent of how one feels, what stance one takes, and so on. Historically, this has been the dominant view. The Greeks tended to think in terms of the discovery of a *telos* for action, Kant thought in terms of determining whether one's will could be universalized, the Utilitarians thought in terms of determining whether an act most contributes to the good. Although feelings and attitudes obviously enter into everyday moral judgments, still, the cognitivists think that such judgments are subject to a distinctive kind of reflective scrutiny. Any sort of positioning or taking sides is, so to speak, logically "downstream" from the business of moral judgment. For cognitivists of a naturalistic bent, like myself and the Cornell School, moral learning is of a piece with other forms of empirical knowledge, which explains why moral discourse fits seamlessly with other areas of factual language.

I differ from the Cornell School because we see the particular way moral terms function somewhat differently. For them, moral terms function like natural kind terms—they serve up a referent via causal links. I don't think this view can do justice to moral thought and deliberation, since moral terms seem to have more meaning of their own to contribute. Moral terms for me (as for philosophers in the Canberra School) have considerable descriptive content, presenting the world under a distinctively normative mode of presentation.

Will all these views come even closer? All of these philosophers are committed to a kind of *methodological naturalism*, trying to place morality within the natural world without introducing any *sui generis* metaphysics or epistemology. All are also committed to preserving as much as possible of the "propositional surface" of cognitive discourse. So there may be further evolutionary convergence yet to come.

***Ethic@*: Do you think it is possible to reach some convergence in normative ethics as well? For instance, is it possible to formulate a principle that Kantian constructivists and Consequentialists could agree upon?**

Railton: There is already a great deal of convergence in normative ethics. To take one example, Rawls' theory of justice, put forward as Kantian, is in fact very similar to what a rule utilitarian might come up with under the sorts of empirical assumptions Rawls makes. Even Rawls' famous strategy of argument, the hypothetical contract that constructivists have

taken as their touchstone, was first introduced by John Harsanyi, an economist, to give a new way of demonstrating that utilitarianism is consistent with demands of justice. For my money, Harsanyi's argument is actually stronger than Rawls', since it avoids a subtle fallacy that creeps into Rawls' derivation. But that is a long story, and our focus here is simply on convergence.

Few Kantians these days are absolute rigorists, and most consequentialists are "indirect", that is, they apply the principle of maximization to systems of principles, not to individual acts directly. The result is very broad agreement. Other moral camps, however, hold out for quite different basic starting points: libertarians and others start off with natural law; virtue theorists and particularists shun the principle-based approach of Kantian and Consequentialist alike; and, of course, there are the religiously-inspired moral theories, which often require fidelity to a sacred text.

Ethic@: What happens if we have a conflict, let us say, between a deontic notion such as respect and a distinct notion such as the maximization of welfare? What kind of procedure could you use to avoid dilemmatic situations which lead ultimately to inaction?

Railton: *Respect*, unlike *well-being*, is a "moralized" notion. That is, what it is to show proper respect for someone will depend upon the moral point of view taken. In an honor culture, for example, addressing a young woman from a different clan politely will be seen as a failure of respect. In libertarian ethics, a hungry person

picking up a fallen apple on another's property—an apple that otherwise would rot unnoticed—is a failure of respect. For Kant, lying to a tyrant to cover up for a dissident is a failure of respect.

A utilitarian, for example, might say that lying to the tyrant is the more respectful act, since it shows recognition of the interests of those who suffer at the hands of the tyrant. Or that taking the apple is more respectful, since it places human well-being ahead of unquestioning obedience to a property convention.

Well-being, on the other hand, is generally speaking not a moralized notion. Sometimes very pious people say things like, "An unjust person suffers more from his injustice than those he cheats". But if this were so, why would people feel outrage when someone who has behaved unjustly escapes legal punishment and lives a long, luxurious, and healthy life, while those he has cheated suffer in poverty? Why would do people demand that the unjust individual be denied the spoils of his cheating and placed in prison? Clearly, people recognize the difference between living well or living badly, whether this is done by just or unjust people. Being healthy and wealthy contribute to well-being, even for the unjust; being in prison is living badly, so it should it is *appropriate* for the unjust.

So I'm not happy with the usual way that respect is contrasted with concern for well-being—for consequentialists, concern for well-being *is* a kind of respect, the most important kind, while following certain rules (don't lie, don't take another's property) may be a sign of disrespect.

Dilemmas are, to my mind, inevitable, because values can come into irreconcilable

conflict. Do I spend extra time with my very talented daughter, since she might blossom into a great talent? Or do I spend extra time with my less talented son, since he needs more help getting through the basics? I can't do both. Doing neither might be worse than doing either one, though. Consequentialism seeks to focus us clearly on the values at stake—including the values at risk if we do nothing.

Ethic@: Going back now to metaethical questions, what would you respond to a critic who says that you don't explain in a satisfactory way normativity if you do it in causal terms? In other words, if you proceed in a reductionist way, are you missing something in the process? By the way, how do you define 'normativity'?

Railton: With only a few exceptions, metaethicists agree on the thesis of *supervenience*: there can be no normative difference between two acts or states unless there is a non-normative difference. Ordinarily, this is interpreted as meaning that the moral supervenes on the natural. If that is granted, then everyone must agree that something we do *as natural beings* constitutes being subject to normative guidance. By analogy, if the mental supervenes on the physical (as most believe), then everyone agrees that something we do as physical beings constitutes being in a mental state. Now I can reverse your question: Since everyone in metaethics needs to explain how causal beings, living in a causal universe, manage to engage in normative thought and practice, then there won't be an "satisfactory way" of understanding

normativity *without* a causal account.

So I'm happy to proceed in the following way. I will develop what I think is the proper causal account, and then others in metaethics, who need such an account as much as I do, are free to use it to solve *their* problem with supervenience. Once this is done, we then can intelligently ask what more might be needed for a complete account of normativity. Until then, I'd better keep at work on the causal account!

One can't give a single definition of 'normativity', since the term applies to two very different phenomena (at least!). First, the "normativity" of values—the sense in which they are appropriately end-setting. Second, the "normativity" of rules—the sense in which they can appropriately be binding upon us. I have tried to give a good causal explanation of the normativity of value in a number of writings on ethics and aesthetics, some of which are collected in *Facts, Values, and Norms*.

I have only just begun to make headway on the normativity of rules—I think the essential idea is that, by *disciplining* ourselves to a norm or standard, we are able to give our acts a meaning or significance they otherwise would lack. Stay tuned.

Ethic@: Do you still hold that instrumental rationality is enough in ethical theorizing? Why?

Railton: I've never really insisted that instrumental rationality must be enough. Rather, I've tried to do as much as I can with instrumental rationality, since it is a tolerably clear notion. I now am inclined to think that what we call

‘instrumental rationality’ might better be viewed as a kind of “agent competence”, and that we must add a more substantive notion of reasons for action. I do not, however, subscribe to the current “buck-passing” view about reasons—I do not think they are ultimate. Rather, they are always *relational*: x is a reason for y relative to F (this fingerprint is a reason for us to believe Jack was at the scene; the aesthetic value of this painting is a reason for the museum to protect it from vandals, etc.).

Ethic@: What kind of empirical evidence have you added to your view in the last 20 years, since the publication of your famous paper, “Moral Realism” (recently reprinted in *Facts, Values, and Norms*)? After all, you do take ethics to be an empirical enterprise.

Railton: In that paper, I argued that we could see certain patterns in moral theory and practice that cohered with my picture of the nature of non-moral value and moral rightness. For example, moral codes had become more universal and inclusive over time, and there has been convergence over elements of well-being. I now am able to give much more careful accounts of these phenomena. In the original paper, I presented a fairly simple picture of the progression of moral thought and practice. I now would qualify this, but also point to several phenomena, each in its way historically unprecedented, that indicate how institutional forms bear this increased inclusiveness out. For example, there now are recognized (though hardly always honored!) “universal human rights”. The list is very close to what a rule

utilitarian might devise, since it includes positive “welfare rights” as well as classical rights against interference. Moreover, there now are international courts and non-governmental organizations that monitor violations of these rights on a worldwide level. Equally strikingly, there now are international environmental accords—people are actually including the scope of morality to future generations and to our relationship with nature.

More surprisingly, perhaps, the idea I advanced of developing an empirical science of well-being has (no thanks to me, I’m sure) now become a reality. A very large amount of psychological research has emerged concerning the nature and experience of well-being, as well as the possibility of learning through experience about one’s well-being. I draw very heavily upon this body of research in my recent work. Even in aesthetics, we now see a burgeoning empirical literature on the nature and appreciation of aesthetic value.

Finally, the sort of offhand evolutionary speculation I engaged in back in “Moral Realism” now can be replaced with a vast amount of empirical research on questions like the evolution of cooperation and moral sentiments. Barely a week passes without someone in evolutionary theory publishing a book on “evolution and morality”, often drawing upon game theoretic simulations and hard evidence from the animal world. I am delighted to see this, and work closely with several active researchers in this field.

Ethic@: Thank you for this interesting interview, Professor Railton. Anything else you would like to add?

Railton: Thank *you* for your probing questions, to which I have given only the barest sketches of answers. I continue to be fascinated by the problem understanding the problems that arise at the intersection of facts and values, and am immensely gratified that so many others—in philosophy, psychology, evolutionary theory, anthropology, etc.—are now bringing their insights to bear on these questions.