
DISCLOSIVE ETHICS

INTRODUCTION

Philip Brey (2000, 2010) believes we underestimate the effects of technology in general and information technology in particular. Computer systems and software have consequences which, by and large, we have failed to notice; consequences which are embedded in the technology and which carry important ethical implications. Along with other commentators, Brey suggests we need to adopt a combination of Value Sensitive Design and Disclosive Ethics. This combined approach would, or so it is hoped, reveal commitments we are unaware of and hence avoid some of the more deleterious effects of information technology. Lucas Introna has used Brey's work as well as that of Bruno Latour to develop an account of the deeply political character of digital technologies (Introna 2007)

We have reservations about the arguments which Brey and Introna offer. We also believe there are plenty of practical reasons why the approach Brey suggests wouldn't work. However, these and like views seem to be gaining currency within the research communities concerned with the human and social aspects of computing. Efforts are underway to promote them more broadly (Cummings 2006, Manders-Huit 2011). Although in some ways we find this popularity unsurprising — many colleagues engaged in research within the HCI and CSCW communities seem particularly drawn to what appear to be new and controversial ideas coming from the social sciences — it *is* a little puzzling. As far as we can tell, the analysis on which Disclosive Ethics rests is a species of very conventional (not to say traditional) Sociology. Moreover, the arguments for an ethics based on that analysis appear to be grounded in a common fallacy, one often associated with the adoption of social science findings by other disciplines. The possibility that a particular sociological description of some phenomenon can be provided does not mean that this is the sociological description which must be used.⁵¹ Both Brey and Introna base the need for Disclosive Ethics on this fallacy. The result is that we are left in the uncomfortable position of being asked to adopt a particular approach to ethical judgments simply because a sociological analysis intimates we should.

In this essay we will set out our reservations. We will focus first on Brey's (2010) extended account of his position and the general analytic scheme he suggests for identifying what he calls "embedded consequences". Having set out our views on Brey's position, we will turn to Introna's use of Brey's scheme. In both cases, we will mostly be concerned with their general character and the ways the scheme is used to demonstrate the ethical import of specific information technologies. Having set out what might be thought of as "disciplinary

⁵¹ If it were not so cruel, you might label this a form of the naturalistic fallacy; one peculiar to cross border trading in the social sciences.

issues", we will move on to a summary of more practical matters and propose there are good reasons to suspect the proposals Brey and Introna make will not work, or not in the way which they suggest they should.

Before we start, though, we want to make two important points. We are not here setting out to challenge or otherwise dispute the style of sociological analysis on which Brey relies. Posing questions of it should not be read as a subterfuge for arguing for some other form of Sociology. We simply ask: Will this Sociology do what Brey wants? Second, we are not denying that some systems and software do indeed pose ethical challenges which we should all reflect on. As will become clear, we are less than convinced that such challenges should be said to be embedded in the technologies themselves.

EMBEDDED CONSEQUENCES

Brey wants to distinguish those cases where the use of a particular system or piece of technology may turn out to have ethical implications from those where those implications are embedded in the technology itself. Both are said to be consequences of the technology. For Brey, the use of geo-locational applications to stalk people for example, whilst definitely an ethical matter, is not an embedded consequence. The ethical challenge comes, rather, from the use of the technology by particular individuals. In his view, almost all discussion of the ethical challenges of information technology is of this kind. Debate on such issues is valuable and important, but not what he has his eye on.

Brey wants to focus on consequences which derive from the way the technology has been designed or from the way it plays in the network of infrastructure and correlated systems within which it is deployed. In this he follows Langdon Winner (1985) and proposes that such consequences may be intentionally or unintentionally built-in. He goes on to assert they may also be obvious and generally recognised or unrecognised and hence unacknowledged. These distinctions which form the core of his analysis, are set out in the table below. Three of the four examples in the table are ours.

	Intended	Unintended
Recognised	Traffic Light Management	SMS & Texting
Unrecognised	Quality of Service	ATMs

Taking each cell in turn. The software that manages traffic light systems uses allocation rules to produce orderly traffic flow. These might simply be constant or variable time allocation or, as is often the case with lights used when road works are in place, by relative volume of traffic as sensed by motion sensing cameras. An orderly flow is intended, and it is obvious when it is produced. The case of SMS and texting is a little different. The Short Messaging System (SMS) was provided by mobile telephone companies to enable engineers to communicate with each other and for the companies, themselves, to be able to send messages to their customers. However, once the facility began to be used outside the engineering environment, it rapidly took off and became the mainstay of growth in mobile phone use. Whilst this was not what the engineers intended, the creation of a product and its market, "texting", was obvious.⁵²

Quality of Service management is necessary in any environment where demand can outstrip supply on a moment by moment basis. In managing access to broadband, for example, suppliers have introduced functionality which enables them to tune the access individual users receive. This tuning can be on the value of the account, the type of application being used, or a number of other parameters. Only very occasionally will users understand the effects of such management, usually experienced though the slowdown of delivery of video, for instance, or the inability to access certain other bandwidth hungry applications. Quality of Service

⁵² See Taylor & Vincent (2005) for the history of SMS

management allows broadband companies to conserve their resources and limit their investment in bandwidth. Thus, they manage their costs and increase their profits.

With ATMs (this example is Introna's & Whittaker's (2005)), the story is slightly different. ATMs are designed to be High Street located dispensers of cash. They are, therefore, designed for the mass of the population. However, because they are designed for the majority, they assume certain physical and psychological capacities: the ability to see and to read; the ability to manipulate a keyboard; the ability to understand pin numbers and remember them; and so on. When the deployment of ATMs coincides with the closure of High Street banks themselves, those who cannot operate ATMs lose ready access to cash. This combination of circumstances leads to discrimination against them by the retail banking sector. Such discrimination, of course, violates the social value of equal treatment for all. For Brey, and here again he follows Winner, of all of the types of consequences, it is unintended and unrecognised consequences and their correlated values which are the most important.

First, let us look at the distinction between recognised and unrecognised unintended social consequences. For sociologists, this is a very familiar distinction. It is the fulcrum around which Robert Merton's (1949) classic paradigm for functional analysis turned. Merton termed the two kinds of consequences "manifest" and "latent" functions. The paradigm set out how, for any designated system of action (social institution, society, sub-culture, social practice), the consequences of courses of action could be shown to be either functional or dysfunctional; that is, they could either contribute to the adaptation of the encompassing system or to its disruption. Tracing through the latent functional and dysfunctional consequences of social phenomena became the central motif of much sociological thinking and analysis.⁵³ Certainly, in one way or another, it underlies the sociology of science and technology⁵⁴ and as our sketches indicate could easily be applied to software and computer systems. As can be seen, the snippets we have given are functional in form.

Anyone unfamiliar with the history of sociological disputation (which runs wide and deep) might wonder why this matters. We are not, here, engaging in arguments for and against functionalist arguments, but attempting to draw the attention of those not versed in the debates over functionalism to the unresolved and highly contested form of that analysis. When argument forms openly designate themselves as functionalist, they are apt to be discounted by many sociologists simply on the grounds that functionalist arguments are held to be inherently conservative in character. In many cases, however, functionalist forms are not openly identified as such, often because, as with Brey, those who put them do not recognise their functionalist character. Both those who advocate and those who dismiss functionalist analysis recognise that there are real difficulties in setting functionalist theories out in a valid form and in securing those functionalist schemes as empirically sound descriptions of their designated phenomena. Brey offers no solutions to these problems and simply begs the question with respect to the very serious difficulties inherent in the type of approach he has taken.

Because his analytic structure is so clearly functionalist in form, it is open to a number of the considerations which we outlined in Essay 7. In fact, Brey's framework displays most, if not all, of them. As a consequence, the grounding for the requirement to introduce Disclosive Ethics turns out to be extremely weak. In the rest of this section, we will outline why.

On Brey's account, the embeddedness of consequences in an artefact depends on the extent to which such an artefact can have autonomous effects and the relative specificity of its context of use. A bridge, for example has more autonomous consequences than a power drill because the drill depends upon an operator for its

⁵³ This is not the place to elaborate on the advantages and disadvantages of using terms like "perspective" to characterise different sociologies. Nonetheless, as we explain in the previous essay, we would want to say that those perspectives which give primacy to the explanation of social structures and processes at the global level are all, inevitably, functional in form.

⁵⁴ Given Merton's role on the development of the history and sociology of technology, where he has become the leading example of what more recent sociologies of science and technology are *not* doing. Brey's failure to acknowledge the style of sociology his own analysis deploys is more than a little ironic.

use.⁵⁵ Furthermore, because the bridge is fixed in its location and has a single purpose, very particular consequences can be said to be built into it. The traffic flows it allows determine certain outcomes.

The question to be asked here is just what we mean by saying something is built into the bridge. Obviously there are the physical characteristics. These constitute the bridge; they are the bridge. Because it has a certain set of physical characteristics, it enables some kinds of traffic (private cars, say, but not buses) to pass underneath. Let us allow, for the moment, the suggestion that because only some people own cars, the bridge can be said to discriminate against non-car owners who use buses. But do we want to say that this discrimination is built-into the bridge in precisely the same way that the physical features are? If so, isn't it part of the constitution of the bridge as well? If we want to say that this is so, that discrimination is part of the constitution of the bridge, then there is no logical distinction between the bridge and its consequences. Brey has taken over this weak argument and further weakened it.

Of course, saying that the discrimination is built in could really be just a way of reinforcing the dramatic tie between the bridge's physical characteristics and its social consequences. In this way it would be, so to speak, a rhetorical flourish which helps to add force to the story being told. Such a device might be necessary because the lack of close coupling between the elements in the story. With Winner, we go from physical characteristics to traffic flows to social demographics to discrimination without securing any of the steps on the way. Certainly no evidence is offered by Brey or Winner to show that the suggested consequences did actually come about.⁵⁶ Without the assertion that the social consequences are built-in, the functional account loses its plausibility. As a result, it can only be secured only by connections which have an "It stands to reason...." ring to them.

So, what are we left with? The argument that consequences are built into artefacts falls on two grounds. First, it collapses the logical distinction functional analysis requires between a social phenomenon and its consequence. This distinction is needed for a consequence to be a function of the social phenomenon. Second, the collapse of the analytic space compounds the failure to offer evidence that the consequence did indeed come about. Here, in adopting his example, Brey simply incorporates Winner's lacuna. Finally, given the weakness of the functional case, all that is left to support the contention is an appeal to intuition.

FROM EMBEDDED CONSEQUENCES TO EMBEDDED VALUES

We turn now to the way Brey moves from consequences to values. Brey begins by suggesting that values are often realised only in part. For example, the value "freedom" is only realised in full if everyone in the world is "completely free". Given the "constraints and limitations", as he calls them, which keep people from being "completely free", freedom is only realised "to a degree" (all quotations from p 46.). What are these constraints and limitations that limit us simply because we live in a social world? Limitations to do what, and where, and under what other circumstances? What would we be free to do but for the fact we live in a social world? What rights others have over us and what obligations we owe to others has been the central motif of ethical theory. Trying to define and understand them has been a key debate in philosophical ethics (see for example Nagel 1991). Brey simply glides over this debate confident in the assumption that we can say *tout court* that values such as freedom are always realised only in part. Another way of looking at this is to say that Brey simply mischaracterises the actual value of 'freedom' as it is held at least in western society, where 'freedom' is actually treated as a qualified value, one to be redefined relative to and not pre-emptive over other values. As we will see, Brey himself doesn't want a company such as Google to be free to exploit its advantage as maximally and persistently as it might like.

⁵⁵ This is not a chance example. Bridges, it will be remembered, were at the core of Winner's argument.

⁵⁶ To be fair, Winner does provide some sort of case. But, ironically, it is (a) not the case he actually wants to make and, as we have said, is (b) very weak.

Having made this claim, Brey goes on to suggest that the partial realisation of a value and, in his terms, its unjust limitation, can be the consequence of a piece of technology. Where this occurs "systematically" (and again we get no help in understanding what that might mean), the limitation on full realisation of the value can be taken to be an embedded consequence. In other words, in such circumstances the realisation of partial value outcomes can be the manifest or latent functions of technology.

There is a small but important point to be considered here. The consequences which Brey cites are all construed negatively. Presumably, though, embedded values need not always be negative. To turn to the bridges of New York example again. Had they been built 2 metres taller presumably the fact that they allowed a wider variety of traffic and hence a greater array of social groups to use the highway, would be regarded as a positive consequence. Improving integration would, we assume, be a positive value outcome. Or, to take another example: computer controlled milking parlours. This technology has significantly reduced the financial and administrative burden on farms, especially small farms, as well as allowing improved animal welfare. What otherwise would have been unprofitable or marginally profitable farms have been able to stay in production. In turn, this has kept families in the countryside and so allowed schools and other services to be maintained. Where this has not happened, rural areas have been "cored out" and so lost their sense of community spirit. Milking parlour software, or so it might be argued, has had the consequence of modifying and preserving a way of life.

What are we to say now? The introduction of the milking systems, like the improved power of tractors, has reduced the need for farm labour and introduced mass production techniques into dairying. For many this can only be a bad thing. But it has kept farms running and this presumably is a good thing. How are we now to decide whether the identified embedded consequences sum positively or negatively?⁵⁷ A significant part of the problem is actually the use of a kind of circular logic. Analysts such as Brey turn to Sociology in the hope it will provide strong (perhaps even scientific) evidence that will support, if not vindicate, their ideological, political and/or moral preferences. And yet, it turns out that they depend on their original preferences to decide how the sums are to be totalled.

Our purpose in labelling the lack of clarity is to emphasise just how loose, and hence slippery, the steps in the analysis are and to emphasise that an argument which is essentially a preaching to the ethically converted, shows no awareness of the fact that valuations in a society are not necessarily uniform and consequently offers no means of resolving these differences into unified judgements. Neither firm conceptual connections nor strong empirical evidence are provided. But the analytic security of each step in the description is what gives functional analysis its plausibility. Brey simply waves his hands at how these connections are to be made. To try see just what could be meant by the suggestion that technology can have value outcomes, we will look at an example which Brey cites, where a clear set of ethical outcomes generated by a specific technology is said to have been demonstrated. The case is that of web search engines as described by Introna and Nissenbaum (2000). This example has another useful feature. Like Winner's bridges, it has become totemic in the literature.

The web, as we all know, is gargantuan. The task for search engines such as Google is to reduce the search space of web pages to manageable proportions. This is done by a combination of page indexing and page ranking. The ranking of indexed pages enables them to be presented to the searcher according to some order of relevance. The search engine, then, reduces the number of pages to be sorted and then sorts them by relevance criteria. At its core, the algorithm Google uses defines relevance in terms of importance with the number of pages linked to a particular page being taken as a proxy measure of importance. Once a page has

⁵⁷ There is yet another irony here. If, as many want to do, you turn to Heidegger for guidance on how to make these determinations, you are likely to end in an impasse. Heidegger's (1977) rejection of the "enframing" character of modern technology reaches its apotheosis in computer-controlled mass production. At the same time, that rejection is rooted in prioritising the (agricultural) community and its way of life.

been defined as important those pages linked to it both derive importance from it and serve to bolster its importance (a kind of increasing marginal return to page rank power). Thus, what emerges is a highly configured landscape with a (relatively) few high ranked pages and a (relatively) large number of low ranked ones.

The key term here is "at its core". In addition to this page link criterion, Google uses other criteria which, it says, underpin its competitive advantage. For this reason, it will not divulge them. These other criteria working together with page linking and ranking enable Google to provide the service it does and hence attract the users it does. In turn, this allows it to charge a premium for advertising and other services. In business-speak, the search algorithm is the heart of Google's business model.

How do matters of ethics enter all this? Well actually for Introna and Nissenbaum, they don't. Following Winner's lead, they want to talk about the politics of technology not its ethics. This is not a trivial point. The relationships between certain orders of valued ends (e.g. freedom, to use Brey's example once again) and certain political means (for example, representative democracy) have long been the subject of "moral" debate. That the practice of politics is entangled with the advocacy of values is undeniable. What is very deniable is that there is agreement on a clear and well understood way of demarcating them. In citing the search engine example, Brey provides no indication of how to translate Introna and Nissenbaum's "politics" into his "ethics", unless we are to assume that all politics *is* ethics.⁵⁸

Introna and Nissenbaum believe the internet and the web are likely to follow the same development path as the mass media and commercial broadcasting. Relying on McChesney's (1997) review of the media, they propose a scenario where commercial interests could become "woven in to the very fiber" (p169) of the internet. If this happens, what was a public good will have been suborned by vested, private interests. Further, since, they assert, web search engines have biases built into them, we can already see the beginnings of this happening.⁵⁹

It is Google's refusal to publish its algorithm in full which is central here. Introna and Nissenbaum accept that search engines must operate in much the way Google's does. However, how they work should be transparent to the user. At present, any user of Google cannot know how the search space has been configured nor how the relevant results have been compiled. As a consequence, users cannot tell if anything *they might have felt to be important* has been left out or lowly weighted. The withholding of information about the weighting algorithm reduces transparency and so is a deliberate, political act. Lack of transparency sets the reciprocal challenge to the writers of web pages. They cannot know how to get their pages weighted highly by Google. Although there are rules of thumb for the design of web pages, where to put important information, how to tag, and so on, everyone who is designing a web page is, in fact, second guessing the algorithm.⁶⁰

Introna and Nissenbaum recognise that there are many different search engines. 'Why should we be worried about "bias" if we have choice?', one might ask. Won't a market for information access emerge which will provide everyone with what they need? For Introna and Nissenbaum, the acceptance of the web as a market (or market of markets) is just as bad as accepting the bias of search engines. However, their position can't be that search engines aren't meeting our needs, only that we cannot tell if they are meeting them as well as they might. To operate efficiently and effectively, markets have requirements for free flow of information that are

⁵⁸ This assumption, though easy to make, would be difficult to sustain. See Nagel (1991) and Taylor (1995)

⁵⁹ We have already discussed this use of "built into". Its use has become the leading term in discussions of technology and ethics.

⁶⁰ We will simply note in passing that *if* everyone knew how to get a high ranking by satisfying the algorithm, the need to provide an ordered listing would simply mean the ranking was achieved in some other way.

not and will not be provided by the web.⁶¹ And anyway, are markets the right way to distribute access to a public good? Should not public goods (access to full information) be open and available to all?⁶²

The combination of a lack of transparency and the use of market mechanisms to provide access to a public good is what, in the end, makes web search engines political. This political character is achieved at the cost of a partial realisation of a value namely freedom (to access information). Without transparency, we cannot know if the algorithms used by Google and others are sacrificing the interests of the majority in retaining open access to information to the commercial interests of a minority. And, whilst we cannot know this, given the history of broadcasting and the mass media we ought to assume not only that it is happening but that the process will accelerate. As this happens, or so the thinking goes, what had been an innovation shaped as a public good to provide open and democratic information access to all will be controlled by sets of vested interests. It will become anti-democratic. Anti-democratic social institutions are, or so it is implied, unethical.

A number of points should be noticed here. First, the analysis slides around between the various types of consequence used by Brey. Because Google won't publicise all the criteria and how they are weighted, we cannot know if, for commercial reasons, the algorithm is deliberately directed to select certain pages and not others, although, given the commercial environment that Google operates in, Introna and Nissenbaum suggest we would be wise to *suspect* that it is. Such deliberate intervention in the operation of the algorithm would be of the same instrumental type as Langdon Winner ends up describing for the bridges of New York but is, of course, not the kind of *unrecognised* consequence that Brey is hunting down. This evidence would then secure the steps from social practice to consequence. However, Introna and Nissenbaum do not even produce the most minimal evidence to support their claim. They show neither that the "bias" (as they call it) is present and intended nor that it actually does serve dominant commercial interests. What earlier we called Winner's lacuna, is repeated here too. Without such evidence, the claim is no more than vague hypothesis (at best) or an allegation (at worst).

The second point to note is the lack of clarity in the concepts such as 'freedom' and 'democracy' that Brey deploys. We have already pointed out that 'freedom' describes a very constrained value in our society. Our freedoms are limited by the freedoms of others. In addition, the notion of 'free access' is distinctively ambiguous. Free access to information meaning access open to all does not imply that access should be free of cost. There are similar confusions in the way in which interests and democracy are connected. The difficulty of determining what *exactly* are the configurations of Google's (or any commercial organisation's) interests or those of any users means that what we are offered are stereotypes of 'commercial interests' and 'democratic interests'. The contrast between them relies on a preaching-to-the-converted sense of 'anti-democratic', whereas, for large tracts of society, the fact that, for example, corporate organisations are not run democratically does not mean that they are undemocratic, let alone anti-democratic. Such simplistic preconceptions of values run up against the undetermined character of values accepted in society.

The second point concerns materiality. Even if the technology was biased and did work in the way claimed, would that matter? Or rather, would it matter for us, *the users of Google*? After all, if Google allows us to harvest relevant and useful information from the universe of web pages, does it matter to most of us most of the time that some particular pages are not selected? If Google does what we want it to do, is it material that the results might be marginally effected by commercial considerations? As users, we, and we suspect others, would think not.

⁶¹ This is an interesting failure to distinguish the presuppositions of an (or *the*) pure economic model of free markets with the conditions within actual markets. Even economists don't think that real markets are like pure markets in all respects (Kuorikoski, Lehtinen & Machionni 2010).

⁶² Of course, in many respects the internet is *not* a market, since many of its services are provided for free.

It might matter, of course, if the world wide web and other services provided by the internet were public goods. But are they? We can see that the original ARPANet could be said to be a public good, provided as it was from Governmental funds, as presumably was the original DNS service. Equally, the hypertext protocol could be said to be a public good originally released as it was by Berners-Lee and CERN for the benefit of the (scientific) public. However, we find it hard to see the web of published documents as a public good, provided and funded as it is by a wide variety of individuals, groups and organisations. Search engines were one of the first services to be provided for that web of documents. In many ways the old metaphor of an information superhighway probably remains the best way of thinking about all this, with the internet infrastructure a public good, like the road system, around which anyone can set themselves up as a site. The parallel then becomes between search engines and wayfinding services such as maps and navigation systems. No-one thinks these are or always should be public goods. Moreover, they can be at vastly different scales and for very different purposes. No one supposes that there should be a single integrated map which showed everything (what on earth could that be?).

It seems, then, that both sides of the Introna and Nissenbaum argument fail. The arguments about the politics of a lack of transparency turn out to be assertions with no evidential base. Those about the suborning of a public good are, at best, muddled. All we are left with, then, is a commonplace but weak analysis of possible latent functions to act as a bridge from consequences to values

In the end, then, we are no further on in securing the move we need. Neither Brey nor Introna and Nissenbaum actually make the case for the embedding of values as latent functions in information technology. Yet, without accomplishing this, the need for a Disclosive Ethics is left hanging in mid air. Of course, such an ethics might be a good thing to have and highly necessary. But the arguments considered so far don't seem to give us reassurance on either count.

DISCLOSIVE ETHICS

If we were to need a Disclosive Ethics, what would that be? For Brey, the thing to underscore is that the consequences of technology are often unrecognised. He calls such consequences "morally opaque" (Brey 2010 p51). Consequently, the task of a Disclosive Ethics is as follows:

Many computer-related practices that are morally opaque are so because they depend on operations of computer systems that are value-laden without it being known. Many morally opaque practices, though not all, are the result of undisclosed embedded values and norms in computer technology. A large part of the work in disclosive computer ethics, therefore, focuses on the identification and moral evaluation of such embedded values. (Brey op cit pp 51-2)

This will be done, he suggests, by looking first at how some system jibes with a given ethical principal such as the preservation of privacy. Next, as and when the introduction of the technology coincides with apparent changes in values, tensions between ethical principles and the priority ordering among them are taken up. Finally, at what he calls the application level, the outcomes of the previous two levels of reflection are applied to technical decisions. Furthermore,

*Disclosive Ethics should not just be multi-level, ideally it should also be a **multi-disciplinary** endeavour, involving ethicists, computer scientists and social scientists. The disclosure level, particularly, is best approached in a multi-disciplinary fashion because research at this level often requires considerable knowledge of the technological aspects of the system or practice that is studied and may also require expertise in social science for*

the analysis of the way in which the functioning of systems is dependent on human actions, rules and institutions. Ideally, research at the disclosure level, and perhaps also at the application level, is best approached as a cooperative venture between computer scientists, social scientists and philosophers. If this cannot be attained, it should at least be carried out by researchers with an adequate interdisciplinary background. (Brey op cit. p.53 emphasis in original)

This raises a host of issues, some of which are disciplinary and some not. We'll start with two disciplinary ones.

Disclosive Ethics is required because of the moral opacity of systems. We ordinary citizens cannot tell if they are discriminatory, biased, anti-democratic or whatever. Sociological analysis will tell us which is what and why. But the logic of moral terms, including second order moral terms such as these, (which is what ethics is about, after all) cannot be fixed by reference to some set of descriptions of how things are. If that were possible, ethics would be a lot less challenging and intractable than it is. Calling a set of software discriminatory, biased or whatever is not a morally neutral description, but rather judgement of how it is to be viewed. Because there can be no recourse to "the facts" outside of the description given, moral judgements cannot be secured by those judgements alone. We have to consider the details of particular cases and instances. Not every selective process is discriminatory or biased. And a system which, on some occasions is, indeed, biased, may not be on others. It is all very well to say that Disclosive Ethics will provide this critical reflection. But how good and how thorough will this reflection be if the value determination is already made?

This worry relates to a second concern. We think it was best summarised by R.M Hare (1986) as follows, though the example he was considering was slavery.

Nearly everybody would agree that slavery is wrong; and I can say this perhaps with greater feeling than most, having in a manner of speaking been a slave. However, there are dangers in taking for granted that something is wrong; for we may then assume that it is obvious that it is wrong and indeed obvious why it is wrong; and this leads to a prevalence of very bad arguments with quite silly conclusions, all based on the so-called absolute value of human freedom. If we could see more clearly what is valuable about freedom, and why it is valuable, then we might be protected against the rhetoric of those who, the moment anything happens which is disadvantageous or distasteful to them, start complaining loudly about some supposed infringement of their liberty, without telling us why it is wrong that they should be prevented from doing what they should like to do. It may well be wrong in such cases; but until we have some way of judging when it is and when it is not, we shall be at the mercy of every kind of demagogy. (Hare op cit p. 165)

Mutatis mutandis the same holds for Disclosive Ethics. Until those who hold that software systems violate their rights to free access to knowledge (or are discriminatory, or invasive, or whatever) can demonstrate, first, that the systems have the effects they claim and second tell us why it is wrong that they do so, we will, as Hare says, be at the mercy of any demagogue who opposes technological innovation.

Third, there is Brey's overly sanguine view of the state of the disciplines of Ethics and Sociology. Wittgenstein once described Philosophy as a "motley". Had he been talking about Ethics, he might well have called it a "mêlée". Nothing is settled. Almost every position on any topic is under siege from some other position. There

are certainly no agreed approaches which can be used to determine the moral value of consequences and actions. Utilitarianism vies with Deontology whilst Virtue Ethics has recently gone through a renaissance. Moreover, there are even highly charged debates over what the proper basis and scope of ethics should be.⁶³ How is all of this to be reduced to clear maxims and rules of thumb for use by engineers and designers without resorting to a vacuous and principle-free pick 'n' mix approach? Appealing to ethical commitment to values which are generally accepted such as freedom or democracy is no substitute for laying out a methodology for ethical choice.

Add to this the general condition of Sociology with its unresolved debates, non-converging arguments and the fact that any functional analysis of some institution or social practice can be countered by an alternative equally functional account, often with diametrically opposed conclusions and contested links between phenomena, evidence and conclusion. Ethics and Sociology are not pacific fields of endeavour. Dispute, even on the most fundamental concepts, is endemic. Moreover, such dispute is not between positions which are secure and unambiguous and those which are not. *All* positions are loose, ambiguous and unsystematic. Erasing or short circuiting these debates to try to provide the kind of advice which designers and engineers will see as useful and would value is, as we well know, rightly only going to be dismissed by sociological (and philosophical) colleagues as trivialising, or worse.

POLITICS, ETHICS & METAPHYSICS

The idea that the means of generating economic value serves particular social interests is an old one. Whilst its most familiar formulation might be in Marx and Engle's famous apothegm about the hand mill and the steam mill, the proposal can already be seen in the political beliefs of Gerard Winstanley and the Diggers and probably has its roots in the medieval radicalism of both Britain and Europe. More recently, it has been a constant theme in almost all accounts of the development of modern technologies. It is not surprising, then, that as we have moved from breathless hyperbole about an imminent "information society" (Castells 2010) to more measured accounts of the role and consequences of information and communication technologies in contemporary society, the argument has been pressed there too. As we discussed in the previous essay, a key contribution to this discussion was Langdon Winner's (1985) discussion of the "politics of artefacts". Whilst Winner was not discussing information technologies, nor indeed technologies which could be said to be the prime means of generating economic value, nonetheless his arguments have become the bedrock on which most discussions of modern technology, especially information technology, are built. Having swallowed Winner's argument, such discussions go on in an unreflecting way under the assumption that all that is needed is yet more demonstration of its importance and further description of how political consequences manifest themselves.

We have argued that Winner's case is not very convincing. We see it more as a tract than an analysis, its aim being to raise an issue rather than demonstrate an argument. As a consequence, we find ourselves demurring from proposals that take its descriptions as their departure point, believing as we do that it stretches its claims beyond the evidence that supports them. What Winner does not do, and despite all that has been said about his paper perhaps he never intended to do, is provide *the* conclusive demonstration at all technologies are inherently political. However, because this is what he has been taken to have accomplished, unintentionally or not, Winner has acted as a Pied Piper. Researcher after researcher has set off to follow where they thought Winner was leading, only to end up in some very strange places indeed.

In this section, we will take all this up in relation to the work of Lucas Introna.⁶⁴ We will show how, because of his reliance on Winner and the weaknesses of the approach he uses, Introna has to resort to 'empirical stretch'

⁶³ See Putnam (2004), MacIntyre (1990)

⁶⁴ The papers we will concentrate on are Introna and Wood (2004) and Introna (2007). Other similar and closely related examples are Introna & Whittaker (2005) and Introna and Nissenbaum (2000)

to secure his conclusions.⁶⁵ We will then turn to the discussion of how the implication of Introna's analysis should be taken up by software engineering. Here we will be concerned mainly with the cogency of arguments made by Bruno Latour since they provide the rationale for Introna's suggestions.

THE POLITICS OF CCTV

Even without the occasional public revelations of how CCTV sometimes gets used, knowing what we do about the agencies which use face recognition technologies in public places, we are likely to be more than a little sympathetic to the allegation that they could be politically biased. Study after study has shown the widespread use of stereotypes in policing and security matters. However, this is not the argument presented by Introna and Wood (2004). Rather, following his related study (Introna and Nissenbaum 2000) of search engines, the suggestion is, first, that the operation of the relevant algorithms are not open to scrutiny, and second, the technology is a 'silent' one. We do not know when our images are being captured and processed. Because we don't know it is happening, and even if we did the process is inherently unscrutinisable, there is the possibility of bias and hence what is termed "micro politics". As we will see, the fact of this 'micro-politics' is not a discovery made by investigation and gathering evidence. It is a stipulation. Because of the very broad definition of 'politics' that is used, for any technology and its alternatives, there *must* be micro-politics in play.

Following the line of argument set out in the earlier study of search engines, the need for scrutiny is held to come from the technical requirement to reduce the information space before comparisons of face patterns can be made. Digital CCTV images are huge. Megapixel colour cameras are now widely deployed. Given the way colour is represented in image processing, this means each CCTV image could be as large as 8 million bits. Face recognition software tries to identify a face pattern within those 8 million bits in real time and then to compare it to a database of stored images or templates. Introna and Wood report that to prioritise speed of analysis and comparison, the face pattern may be encoded in as little as 84 bytes. The information reduction and determination of the array of pixels which constitutes a 'face' is done by complex statistical analysis. Based on these procedures, patterns within the image can be associated with patterns in the database. Once identified, the 'face' can be compared to the relevant database.

Information reduction combined with other features of the process lead to what Introna and Wood call "bias" in the system. Such bias leads to 'micro-political' consequences. First, the algorithmic processes *are* statistical and therefore some level of (randomised) error is to be expected. Second, the comparison to the images in the database is only as robust as the robustness of the images there. Lack of representativeness in the sample population or poor quality comparator images will all affect the robustness of the association. The possible implication of these is summarised by Introna and Wood as follows:

To conclude this discussion we can imagine a very plausible scenario where we have a large database, less than ideal image due to factors such as variable illumination, outdoor conditions, poor camera angle, etc, and the probe image is relatively old, a year or two. Under these conditions the probability to be recognized is very low, unless one sets the false accept rate to a much higher level, which means that there is a risk that a high number of individual may be subjected to scrutiny for the sake of a few potential identifications. What will be the implications of this for practice? (Introna and Wood, op cit p 189).

Such implications are defined as unanticipated and unacknowledged consequences following from possible combinations of circumstances: the suppliers may oversell the robustness of the technology; the operators

⁶⁵ We owe the identification of this particular rhetorical trope in social science to Geoffrey Marshall (1983). Others that he pointed out are: Thematic Simplicity; Model Depiction; Terminological Adaptation (something Actor Network Theory excels in); Generalisation Rush; Tangential Assertion and Dialectical Leaping.

may not understand the system's limitations; the socio-political environment may encourage a tendency to accept false positives, and so on.

One might imagine that in an environment where there is an acute sense of vulnerability it would not be unreasonable to store these false positives in a database 'just in case'. These false positive may then become targets for further scrutiny. Why? Just because they have features that make them more distinctive. We are not saying that this will happen. We are merely trying to indicate how seemingly trivial 'technical issues' can add up to political ideologies at the expense of some for the sake of others.(Introna and Wood of cit p193)

Before we move on to consider where, on the basis of this case, Introna wants to take the argument, we ought to make a couple of points. First, it is clear that no-one is saying that the bias of these systems is the result of what elsewhere has been called 'institutionalised racism' or the like. It is not the effect on the use of the technology of a cultural milieu that Introna and Wood are describing. It is, rather, the effect of the technology within a complex socio-political and cultural milieu which is micro-political. This is, of course, directly in line with Winner's argument.

Second, even though CCTV is highly visible, the "silence" of the technology is of vital import. However, for ethical evaluation, the relative harm that derives from this silence has to be calibrated against the perceived level of harm from the actions that the technology is designed to prevent. We can, indeed, imagine circumstances where such capture would certainly be felt to be unnecessarily invasive. However, we can equally well imagine cases where it would be accepted as proportionate, pre-emptive action. Yet again, it all depends on the circumstances. That the technology is silent is not, by itself, either politically or ethically significant.

Third, we can always imagine, as Introna and Wood do, scenarios where over-zealous selling, poor professional practice, and so on lead to misuse of the technology. But imagining scenarios is not describing actual instances and gathering evidence of actual cases. Once again, as with Winner, the lack of evidence from actual technologies in actual circumstances that the consequences were being realised, means all we are left is little more than scaremongering.

It is our contention that Introna and Wood end up in this position because they have taken Winner's account to be both a factual and generalisable description rather than either a political argument or the sketching of a potential research programme. For them, Winner describes not just what will happen but what must inevitably happen. Consequently, Introna and Wood see their task to be the teasing apart of the processes by which what is termed an "unauthored strategy" (or more familiarly a "hidden hand") works to achieve this general outcome. This unauthored strategy serves to ensure that technological innovation will fit and reinforce the existing dominant socio-political and moral order. Through its micro-political enmeshing with existing institutionalised norms, values and practices, technology has the latent moral and political function of reinforcing the *status quo*. This enmeshing provides information technology with its significance. Information technology is now shaping how we experience the world. It is to the analysis of significance rather than the causal description of effects to which this whole approach is given over. The finding that technologies have politics is only important because of the significance which can thereby be given to them.

THE SIGNIFICANCE OF INFORMATION TECHNOLOGY

For Introna (2007), information technology has the critical role in modern society. There is a duality to this. As a technology, it has become the pre-eminent source of images, metaphors and ways of thinking which reaffirm what Charles Taylor calls the "social imaginary" of modern life (Taylor 2004). This imaginary is

instrumental rationality or what Ellul (1964) called "technique". Indeed, by and large, rationality is defined only in means/end terms. The instrumental outlook, what Borgman (2008) terms "the device paradigm" moulds — or "En Frames" to use the Heideggerian expression — the way we think. The second aspect comes from the pervasiveness of this technology. Unlike, say, the internal combustion engine or electric power, information technology is not just to be found everywhere; it is in everything. It is all pervading. It is shaping our experience and thus our conception of who and what we are. To paraphrase Orwell: all technologies determine our social imaginary, but information technology is more determining than others.

Introna starts from Latour's (2002) assertion that instrumental rationalism is predicated on the separation of means from ends. In its view, technologies are seen only as means and often described in tool-like terms. As tools they are morally and politically neutral. It is those that design, build, deploy and use technologies who make them a means for good or harm. However, because of their significance for the shaping of our ways of thinking about and experiencing the world around us, such technologies must be drawn into moral discourse. This will only be possible if the disjuncture between neutral means and valued ends can be overcome. To do this, Introna invokes Latour's ontology of "folding" as the means by which this can be done.

As we saw in Essay 4, Latour thinks the classic binary distinctions which underpin modernism (fact v value; subject v object; representation v reality, and so on) are all suspect. They are based upon a metaphysics which, first, divides the world into human and non-human entities; and second, prioritises the former over the latter. This metaphysics permeates our imaginary and forces us to think that there can be no logical bridge between the two categories. However, from his studies of science and technology, Latour wants to argue that such binaries are unsustainable. Facts are enmeshed in values; representations are all the reality we have. Moreover, the ontology of subject and object blinds us to the moral agency that objects (especially technologies) can have.⁶⁶ Objects, tools, technologies are not just used in ways that have moral consequences, they are immersed in and constitutive of moral choices. Studies from as varied cases as nuclear power, electric cars, the development of vaccines, and many more have shown that science and technology develop within and as part of a moral order. To overcome our blindness to this, Latour proposes we should adopt a unified ontology of mutual relationships. Material culture including technology is not over against human social life but deeply entangled in and with it. Folded within these entanglements are time, space and human socio-cultural history. They are rolled up in them in a similar fashion to the higher order dimensions of String Theory. It is because of these foldings, these hidden but surveyable dimensions, that the duality of subjectivity and objectivity fails as the basis of ontology and so as the buttress for the doctrine of the moral neutrality of technology.

Given the style and level of analysis Latour provides, it is hard to know how exactly to take this set of arguments. On the one hand, he appears to want to confront the whole of mainstream Philosophy by denying its basic categories and reducing all discussion of its paired concepts to uniform dichotomies. He is doing this, not on the basis of detailed and rigorous conceptual analysis, but rather through a dazzling fire hose of empirical studies of science and technology, metaphors, and images. This is more Finnegans Wake than *Principia Mathematica*. However, since the sociologising that Latour invokes to support the dissolution of the distinctions is premised upon the methodological assumption that, for the purposes of giving a *sociological* description, we can suspend the distinction between representation and reality, fact and value (social constructivism *is a methodological assumption*), philosophically we are getting no more out of the analysis than Latour has already put in.

Second, what does this flattening of ontology actually mean? If cars, hammers, power stations, software are to be viewed as actants standing alongside humans in webs of relationships, human and non-human, what does

⁶⁶ In many ways, this can be seen as providing a moral twist to the standard actor-network theory argument that technology has material agency. see Pickering (1995)

this imply for moral (and legal) *theory*? Is he really proposing that we should scrutinise the conduct of these technologies the way we scrutinise the conduct of people and hold them accountable in the same ways? Can our concepts of responsibility, blame and approbation be extended to objects and technologies?

Third, if the bridge from studies of technology to metaphysics is secured by the former's refusal to truck with modernist dichotomies, how are we to take the results of these studies and their success in Sociology, studies of science, and now studies of ethics and technology? Are they of the same order (subject to the same strictures) as the studies they treat as their topics? Are they, too, no more than exercises in invention which are to be judged through the influence of relationships, interests, and ultimately power and (moral) domination?⁶⁷

Of course, as we saw in regard to Winner's claim about the politics of artefacts, what is under discussion might not be a proposition at all. It could be a (post-modern) joke, a rhetorical ploy pushing to the extreme a sociological account of technologies and tools. As such, it would be a trope, no more, intended to be left behind, rather as Hume advised all philosophising should, when moving from the study to the world of practical affairs.

Introna tries to resolve these difficulties by aligning Latour's claims with the view that Heidegger (1962) took on the constitution of a first philosophy. The primordial ground of such a philosophy must be our immersion in and experience of the world not our abstraction from and reflection on it. For Heidegger, the latter is not just Descartes' error, but the wrong turn taken by all western philosophy since the pre-Socratics. Following Latour, Introna suggests, the unified ontology is available through our experience of objects and technologies in the world around us. The foldings inherent within the latter are available to us as the "affordances" of this or that tool, this or that device. Such affordances are not added to the technology but "there" to be "grasped" by us in our ways of using them. Affordances constitute the possibilities and potentials of technology which are realised in use.⁶⁸ So for us, a mobile phone has a variety of uses. We can communicate with friends and relatives, store images on it, use it to manage our calendars, and so on. For a society with no concept of wireless communications and "mobile apps", the phone can be no more than a paper weight. The culture of use is folded within the phone and enables our use of it.⁶⁹

Introna takes one further step, though, and suggests that it is this constitutive character of technologies which defines their political nature. This enables him to expand the space of politicised decisions integral to technologies. As well as design and operative decisions, decisions over markets and users, decisions over implementation, roll-out and deployment, decisions about sales strategies and market entry, customer engagement, product quality, product maintenance and support, and of course product end of life, all have to be closed out and made under real conditions of time, budget and practicality. Each decision achieves closure on an issue: the inclusion and exclusion of particular outcomes and possibilities. For Introna, the cumulative effects within any socio-technical environment of such decisions and their consequences must be seen to be political in that they determine the meanings that the technology (the affordances, opportunities, possibilities) convey. In an echo of Cultural Theory's characterisation (see Essay 6), this determination, he says, is "hegemonic".

It is this ongoing, and often implicit, operation of hegemonisation - of inclusion and exclusion - inherent in all political sites which is the concern of a Disclosive Ethics. (2007 p 15)

⁶⁷ This, it will be remembered, was Woolgar's conclusion. See Essay 2.

⁶⁸ We will just note, in passing, that Gibson's (1977) notion of "affordance" is based in and motivated by a very un-Heideggerian psychology. Unless Latour and Introna are using the term metaphorically, the introduction of "affordances" can only make the theorising of this ontology less consistent not more.

⁶⁹ The sheer banality of this observation seems lost on Introna. Everyone knows this including those nations which have moved directly to using mobile phones because they lack the infrastructure of landlines.

Hegemonisation is outcome/consequence of all the design, implementation and roll-out decisions associated with ATMs, web search engines, face recognition systems, and so on. The ethical and political challenge is to reveal this hegemony and open up the possibility of its reversal. This is the role conceived for Disclosive Ethics.

There are several things to say about all this. First, despite expansion of the conceptual apparatus to include elements of post-Gramscian political theory, the structure of the analysis remains steadfastly functional in form. An unauthored strategy, a hidden hand guided by the interests of the dominant, ensures that the nexus of human/technological relationships reproduces hegemonisation. But this is not a *finding* of this way of looking at technology. It is a motivating assumption. Hegemonisation is seen to be a functionally adaptive process whereby outcomes are rationalised in terms of dominant interests.

Second, and this is critical, to make the whole approach tractable, Introna, like Latour, turns the metaphysics, i.e the dissolution of the dichotomy between human agents and material culture, into a methodological assumption. For the purposes of carrying out his (sociological) studies, he proposes to treat human agents and material culture as similar orders of (moral) being. This licenses his description of them and their consequences. But it does no more than that. This methodological move is to be justified by the evidence it makes available to us and the insightfulness, rigour, interest, fertility, or novelty of the *sociological* findings it enables, not by the popularity or radicalness of ethical (or metaphysical) stance we might choose to draw from it. To justify that ethical stance, we need to show first when, where and how hegemonisation is taking place (the close coupling of practice and function), and second why it is wrong (the ethical evaluation). For that, we would want a very different order of justification and a very different kind of argument; one which looked at detail of instances and cases. Without such an argument, the turn to Disclosive Ethics remains unjustified. There would be no more reason to hold with Introna that information technologies pose a critically important threat to the openness of our society than there would be to agree with Dr Pangloss that all is for the best in this best of all possible worlds.

This takes us to a third and very familiar point. In none of the cases, examples, specimens that Introna discusses, is there any *evidence* for his argument. As a result the steps in the analyses become very loosely connected, something, again, that Introna shares with Winner. Furthermore, there is a deep paradox to be found here. The philosophic premise which is supposed to underpin modernity, the duality of fact and value, of how things are and how we see them, is precisely the premise which Latour wants overthrown. We cannot hold the distinction between representation and reality, fact and value. Using his conceptual apparatus of foldings and hegemonisation, Introna builds a picture of how information technology systems can be seen. This picture prioritises the politico-ethical consequences they can be described as having. But, if Disclosive Ethics is to work, it has to be possible to reach conclusive judgments and finalise re-designed technologies, the picture of how things are must be *independent of how we are choosing to see them*. The task of assembling the detail of any one design process is immense and no indication is given either of how we could collect that information nor what redesigning the design process itself would entail. This leaves Disclosive Ethics as a vague portmanteau of theories and nominated but unexamined empirical domains.

GETTING DOWN TO THE PRACTICALITIES

We finish with some practical concerns. The first relates to how real life commercial software (and other) projects are lived.⁷⁰ Overriding everything else is the fact that there is never enough resource to implement the signed off specification document against the deadlines set. Some things have to go simply to stay on track. In addition, the schedules of technology development are unremittingly wicked. From the start, no-one expects to adhere to all of them. Delays are endemic and overruns normal. Yet some deadlines cannot be moved and so the project must be squeezed down to fit the available resource (time, money, manpower) to

⁷⁰ These observations draw upon our own and others' fieldwork in commercial software and development environments (e.g. Sharrock & Anderson 1996, Button & Sharrock 1994, Rooksby et al 2009)

get it 'out the door' when it was committed for. And then there is the value engineering. Project costs never undershoot estimates and so the cost component of the eventual price has to be engineered down.

Into this fraught environment, Brey is proposing to introduce sets of professionals who, as we have continually suggested, are trained not to pursue convergent thinking; whose disciplines are in an open-loop, open-argument state and the members of which suffer from what Kenneth Burke called a peculiar occupational psychosis, supposing that all life is, or should be, lived as if it were a university seminar. For project delivery, decisions need to be made and secured; designs have to be agreed and frozen. Specification and decision drift are *the* banes of the project manager. As a consequence, the average software project is the last place to conduct open-ended non-converging philosophical and sociological discussions about the significance, likely consequences, and ethics of design; and software project teams are the least likely people to stand for them.

The suggestion they should join such teams assumes sociologists, by dint of their sociologising, can foresee the likely consequences of some innovation. In fact, of course, they cannot; or at least no better than anyone else. What Sociology provides them with is a template, a structure, for relating consequences to institutional practices once those consequences are to hand. Engineers and designers might be as well turning to crystal balls or Tarot cards get the answers they need as they would to hiring sociologists.

There is more! We were part of the effort which opened up the promise of ethnographic fieldwork for software development. We still believe that it has a lot to offer. However, in the intervening years, we have watched as what can only be regarded as old fashioned customer relationship management has masqueraded as ethnography. As the consultants have moved in, the canons of fieldwork rigour have been eroded. Now, it seems, any kind of conversation with any kind of user/manager/customer can be called ethnography. The net result is a Gresham's law for research where good ethnographic work is being driven from the technology development environment. We have little doubt that if Brey were to be successful and persuade software managers to employ sociologists and philosophers to undertake the assessment of Disclosive Ethics on project teams, the consultant companies will be very fast second movers. Inevitably, as untrained and unskilled people start to ply their trade, the same dispiriting descent into banality will occur.

CONCLUSION

Philip Brey is motivated by good intentions. He wants to make sure that sufficient attention is paid during the design process to ensure that the chances of ethical breaches in the use of information technology are minimised. This is very laudable. However, the approach he takes is unconvincing in three major ways. It utilises a form of sociological analysis which was not designed to provide the sort of accounts he needs for ethical analysis. Second, he offers no clear way of reasoning from the accounts it does give to judgements concerning ethical outcomes. Third, if, as Brey intends it should, his approach were to be introduced into the development environment, it is likely to be either disregarded or damaging. Disclosive Ethics offers no panacea for the ethical challenges of information and other new technologies. Mandating it as part of the design process will only cause more problems than it can hope to solve. Because of the generalised and abstracted nature of its reasoning, it ends up advocating that general measures of (quite drastic) reconfiguration of the social order are required to counter what are, at best, only *occasional* features of the socio-technical order

That information technologies can be put to worrying uses is not in dispute. There are enough instances in the research literature and elsewhere to put that claim beyond debate. Addressing this issue is a matter of careful design and equally careful regulation of the uses and contexts of such technologies. However, to say this is not to say, thereby, that information technologies are *necessarily* ethical or *inherently* political. This claim is universal in its quantification. All technologies are political and ethical in their outcomes. The problem is that the claim can only be secured within a functional analysis, one where the purpose is to demonstrate what one means by inevitability or necessity of outcome in this regard. To put it starkly: functional analysis looks

backwards to the analytic presuppositions not forward to the facts. And yet, if one wants to intervene in the world, to impose extra strictures on systems designers and developers, to demand existing systems be re-configured and re-built, then one has to offer an account that is rooted in how things actually are. It is because they really, really are political and ethical, that these systems must be changed. Functional analysis does not tell you how things really, really are, except under the auspices of functional assumptions (to use the phrase). Faced with this impasse, Introna uses Latour's convoluted ontologising to try to unify the material and cultural worlds in a single moral frame; such a frame is how the world is. But the unified ontology is itself built on the overthrow of such distinctions. It proposes that these distinctions are nothing but ways in which we construct and enforce an account of reality. At that point, the whole project falls apart and Introna is left with nothing but exhortation and expostulation.