CHAPTER 5

Innovations

As I have already noted, there is nothing in the concept of open access that means anything must be done differently except to lower price and permission barriers to research. Indeed, Peter Suber is at pains to specify that open access is not about circumventing peer review;[1] an especially important observation given some of the misconceptions identified by the OAPEN-UK project. In this light, it might be fair to ask why it is necessary to discuss innovation, changes to editorship or modifications to peer review at all. After all, such discussions surely muddy the waters of open access. I would suggest, however, that there are three reasons why peer review and other forms of experimental innovation must not remain the elephant in the room when we talk about open access. Firstly, it is necessary to keep talking about peer review, in particular, because of the recurring misplaced belief that open access must *inherently* refer to lower standards of quality control (it means no such thing). Secondly, these shifts in publication practice allow us the space to rethink peer review and other practices and to ask whether there are analogous changes, facilitated either socially or technologically, that could be worth exploring at this time of transition. In fact, just as one of the arguments for open access is that it is culturally elitist and untenable to presume that a broader audience can neither understand nor appreciate scholarship, there are, I would argue, parallels in peer review and editorial practice that could reflect this same principle inside the academy. Thirdly, peer review is a key element for discussion because the economics of scholarly publication are intertwined with systems of value.

In this chapter, I want to flag up several ways in which the current system of academic publishing *could* be reformed in a new digital era.

Nothing written here is to be taken as necessary for open access. Instead, these musings are designed as thought experiments to determine the risks and benefits of potential modifications. It is also true that much thinking has already taken place on such issues over the last twenty years.[2] That does not mean, however, that I do not wish us to continue to think about this. As I hope has been made clear throughout this volume, I consider thinking about academic publishing as a form of reflexive critique. It is the histories of academic publishing that shape current practice and determine the possibilities for academic discourse and, therefore, communication. Thinking, in the contemporary tradition of this mode, must involve, as Michel Foucault reframed it, 'a historico-practical test of the limits we may go beyond'.[3] While this is twisting Foucault a little for current purposes, it is in a similar vein of thinking about the historically contingent practices that structure our writing that I wish to proceed.

PEER REVIEW

Peer review is often thought of as inviolable and can be held in extremely high regard. After all, what could be better than a community of self-regulating experts deciding on the merit of a piece of work, in some cases without knowing the author's identity? It sounds like the ultimate form of meritocracy. All is not, however, quite so rosy in actual implementation. Indeed, as far back as 2003, it was proposed by Carlos Alonso, Cathy N. Davidson, John Unsworth and Lynne Withey that the problems of scholarly communication cannot merely be confined to access and economics. As they note, 'the "crisis" of the scholarly monograph, then, is not merely a crisis in the economics of scholarly publishing, but also in the processes of peer review and academic self-governance, prompting reflection on practices of scholarly evaluation that we have simply taken for granted'.[4] In this first section, I want to outline provocatively the potential flaws in the extant systems of peer review before moving on to suggest new alternatives in an open-access, digital world. Some of the criticisms here will only apply in the most negative of cases; often, existing forms of review work fairly well. Nevertheless, if we value quality control mechanisms, it is important to consider the potential failings of our models, even if we eventually reject proposed

alternatives. In other words: there is always merit in playing devil's advocate, at least temporarily.

The first point to note is that the gatekeeper model – that is, the system of deciding on permissibility before publication through both publisher policies and peer-review practice – works on a series of unspoken ideological assumptions that are never wholly objective and apolitical, but are rather, at the extreme end, based on a series of exclusions and marginalisations. While much review certainly is aimed at improving work and there are often substantial efforts to bring work up to standard through iterative commentary, at high-end journals and publishers there must be a percentage of rejections based on notions of importance in order to match page budgets and preserve prestige. This is because a gatekeeper model sometimes pre-defines its audience and disregards a series of important questions. For example, how can one wholly know the value of the material that is pre-excluded given that we exist within ideologies that are not always explicitly clear from our immanent positions? How can we know what *will* be of value in the future? What do we make of the exclusions and other spaces that, under the gatekeeper model, we cannot even know at present? In other words, the model of review, as it currently stands, can only be bound by temporal norms: the conditions of the present. This makes it very good at selecting work that will be highly valued in the current instance, but contributes towards a significant weakness in our system of pre-excluding material that might be of worth at a later date.

Secondly, in one (admittedly somewhat negative) reading of peer review, it is possible to see the filter-first method as a development entwined as much with economics as with quality. Historically, one of the key functions of the gatekeeper has been to reduce the quantity of permissible material; clearly, for economic reasons, not everything that is submitted can be published. This was not only an effort to avert what is now called 'information overload' and what are perceived as low standards, but also because each issue of a print journal or each book had a specified page budget. In the world of print and physical commodities, there is a need to restrict the quantity of output because there is a material cost for each page that is printed and distributed. This is, clearly, no longer directly the case (although every output has correlative labour that must be

compensated and so has an economic scarcity) but persists through a culture that, as I have already noted, Gary Hall calls our 'paper-centrism'.[5] Under this traditional model of review and economics, the price of a subscription to a journal must cover not only the cost of the material that is printed and owned but the cost of the editorial labour that was invested in administrating the pieces that are rejected. As it is with icebergs, so it can be with editorial labour; it is easy to forget that, in a filter-first mode, the majority lies hidden below the surface.

Thirdly, single- or double-blind forms of peer review assume an honourable motivation for reviewers and provide few ways of holding readers publicly accountable for their decisions. Hypothetically and provocatively: is it right that a mere two academics, in most instances, although sometimes only one, have the private, unaccountable, final word on an article's or book's acceptability, particularly when there has been intense debate over the statistical significance of the number of reviewers, even in scientific disciplines?[6] For Early Career Researchers (ECRs) this private decision can be the difference between a lifetime of employment in academia and a lengthy period of retraining. Furthermore, reviewers and journals are often only evaluating a piece with one specific audience group in mind; if judging on the 'importance' of work, it is crucial to ask 'for whom?' Different journals and presses, of course, have different target audience constituents, but even these sub-fields may be fractured and subject to competing notions of 'importance' within a discipline. In other words, how can one accurately pre-judge, within one's own temporal, geographical and disciplinary immanence, what may be of worth to scholars free of these constraints? While it could be argued that the specificity of journal and publisher remits renders such considerations irrelevant, this lack of accountability and, as will be explored below, logic in the admissibility of papers is a problem that is exacerbated by the traditional double-blind system.

Fourthly, the term 'double-blind', as well as carrying ableist insinuations about the partially sighted in its language, can actually be a misnomer. Theoretically, the author should be unaware of the identity of his or her reviewers and vice versa, a mode used in most journal reviewing. The benefits of this are easy to articulate: it is designed to encourage an impartial assessment of the work, rather

than the author. Furthermore, reviewers are supposed to be protected from professional repercussions in cases where, for instance, the author is a prominent figure in his or her field. Often, however, this is utopian, in the sense of 'naïve'. In many small fields where work may have been presented in early versions at conferences, where authors are known for adopting a specific stance, or simply through flawed metadata erasure and/or slips of self-citation, the identity of the author can be ascertained. While it is less often that slips occur the other way around, it is sometimes possible to guess the most likely reviewer of one's work simply by dint of his or her expertise and by the idiom of his or her returned report.

Furthermore, anonymity can be problematic. The lack of accountability of readers, as above, can lead to harsh, penalising reviews, rather than to feedback that, while rigorous, intends to work in community to elevate a work to a publishable standard. Additionally, there is also something strange about the perseverance of anonymity after publication. As noted earlier with respect to the role of citation, universities and academia function on genealogies of validation; that is, on hierarchies of prestige that trace the flow of academic 'capital' and authority through publications.[7] As explored in earlier chapters, journals and presses are primarily only as valuable as the genealogies that validate their work as high quality, through submission quantity/quality and rejection rate, underpinned by the labour of peer review. However, in the current way of working, what remains is a situation where, instead of the process of review being visible in order to validate the work, the quality of the review process must be inferred from the perceived post-publication quality of the research.

To extend this argument: there are only two ways, both flawed, in which the quality of the review can be ascertained under current practice. The first of these is through trust in nominal journal or publisher brand; the problems of prestige to which much of Chapter 2 was dedicated. The second way in which journal or publisher quality is crudely measured and the one that surely most affects scholars' perceptions lies in the duplication of labour when reading an article or book; a type of second review in which academics bring their own evaluative skills to bear on already published work. During this constant re-evaluation, the blame for poor quality is often put down either to the author or to the journal/press brand. This is

interesting; what seems to have failed is actually the peer review, gatekeeping function, but this is not, in a mode of journal or publisher brand, the way in which it is perceived. While in some ways this is a fair appraisal, there could be modes through which the journal could signal the degree of delegation and trust that has been relied upon and to which I will now turn my attention.

PROPOSED EXPERIMENTS IN PEER REVIEW

The most obvious way in which these problems might begin to be addressed at the moment of transition to open access is to rethink anonymity in the review process, as has already happened in many scientific disciplines. However, it is worth saying upfront that each of the various combinations of the review anonymity matrix comes with its own problems and it may be the case that none is, in the end, as satisfactory as some forms of blind review, except, perhaps, for at least being more honest about the potential flaws. The first of these potential changes would be to remove the author's anonymity while maintaining the anonymity of the reviewers, which, although common practice in book reviewing, seems to add very little. In this scenario, reviewers could judge solely on the past reputation of the author, rather than the merit of the piece alone while remaining unaccountable for their actions. The only debatable benefit here might be that the reviewer can assess the competence of the author to write on the topic at hand, although this seems dangerous; surely the sole criterion for being able to write on a subject should be through the production of valuable work, not a track record in the area?

Conversely, one could take the opposite stance and remove reviewer anonymity (at various stages in the process, but primarily after the review and regardless of outcome) while retaining the author's veil. This mode brings absolute accountability upon reviewers while protecting the author from pre-judgements. It also gives a clear genealogy of validation and militates against corruption to some degree as any conflicts of interest would be immediately clear. The disadvantages of this approach are also obvious, though. Any system that brings unbalanced extreme accountability will result in a conservative situation of strict, normative appraisals, thereby potentially

ruling out a whole body of useful work that may be barred by fearful reviewers. While some may see such strictness as an advantage – a tightening of review standards – given the historical parallel to page budgets and evolutions in social and technological filtering processes (see below), the argument that this is solely about quality may be less solid than might be thought. Finally, although the revelation of reviewer identity in some ways helps spot corruption (as any affiliation to the author would be obvious), the extreme burden to 'make the right call' could encourage reviewers to seek the author's identity. This tactic exposes reviewers and makes a thankless task perhaps even more risky.

What, then, about completely removing all anonymity from the process? There are some advantages in this case (as outlined above) but there still remains no counter-balance to the problems that could arise as a result of exposing reviewers. Conversely, reviewers would surely also be prone to appraise the author's identity in this case.

Evidently, in each of the cases where anonymity is removed, during the review process itself, there are problems that seem, to some degree, worse than the flaws in a double-blind setup. However, this solely applies when discussing a gatekeeper model in which a paper only sees the light of day so that the journal may be associated with the most exclusive papers in order to protect its brand. Other, more radical, experiments in the sciences have worked to change this. For instance, the review criteria of *PLOS ONE* – now an enormous journal that has published over 100,000 articles since its launch in 2006[8] – reads as follows:

> Too often a journal's decision to publish a paper is dominated by what the Editor/s think is interesting and will gain greater readership – both of which are subjective judgments and lead to decisions which are frustrating and delay the publication of your work. PLOS ONE will rigorously peer-review your submissions and publish all papers that are judged to be technically sound. Judgments about the importance of any particular paper are then made after publication by the readership (who are the most qualified to determine what is of interest to them).[9]

This model implements a fix for what Clay Shirky has called 'filter failure'. In an influential 2008 keynote at the New York 'Web 2.0' expo, Shirky identified a new form of post-Gutenberg economics in which he proposed the overturn of the filter-first model

(i.e. advanced selective peer review). This is because the cost of putting something on the internet, in raw form, can be extremely low. Once the material is online, Shirky intimates, it might then be possible to create social and technological filter mechanisms that glorify the good and bury the bad.[10]

At first, this standard of publishing all papers that are 'technically sound' appears to have no analogue for many disciplines in the humanities. This may prove to be correct. As a hypothesis, though, a technically sound paper in the humanities could: evince an argument; make reference to the appropriate range of extant scholarly literature; be written in good, standard prose of an appropriate register that demonstrates a coherence of form and content; show a good awareness of the field within which it was situated; pre-empt criticisms of its own methodology or argument; and be logically consistent. These are, indeed, the exact checks that one would expect an editor to make before sending a piece out for review. While this is just a cursory stab at a definition and not meant to be finalised, implementable criteria, many of the problems of the review system as it stands could certainly be addressed through the formation of explicit consensus as to what constitutes an acceptable barrier to entry in the humanities.

Secondly, though, the inversion that *PLOS ONE* effects could leave it open, as was the very first scientific journal, the *Philosophical Transactions of the Royal Society*, to John Hill's 1751 critique: the inclusion of 'trivial and downright foolish articles'.[11] In other words, by inverting the review methodology, *PLOS ONE* could expose itself to admitting rubbish. The difference from the contemporary situation, however, lies in the economic models and technological filters at our disposal. In 2014, sophisticated full-text and social search mechanisms exist that can bury unpopular material on the furthest pages of results but without removing such items from the economy altogether. This is not, therefore, a removal of selection, but rather a different way of filtering. There is still selection of material but, at a later date, it also becomes possible to see those manuscripts that were not initially favoured. The advantage of this, as with the arguments for open access more generally predicated upon an anti-elitism, is that we do not presume to know what will be important for all time. Instead, we replace such a system twofold with the ability to ensure

that what is relevant now is found and valued while also allowing those papers in niche fields or in areas that have yet to gain any prominence to be found, if and only if the seeker desires. In this mode of post-publication review, everything would be assessed, but it would be done after the fact and the exclusion of material would not be a permanent pre-silencing, but rather a process of continuous community consensus. There are a couple of assumptions that underpin this mode. Such a system would rely upon the correct translation of the community will by software and also upon the continued participation of academics; voting systems for article/book prominence must be hardened against gaming and academics would need to be engaged enough to signal value. Of course, there is no guarantee that the peer-review criterion of 'technical soundness', however translated, would be free of abuse in itself, but this could be a step in the direction of militating against some of the perceived failings of blind review. Hurdles, nonetheless, remain, especially because there are also no current incentive systems within the academy that adequately value peer review, especially in the sense of an open, post-publication mode.

To close this section, however, it is worth once more referring to Kathleen Fitzpatrick, the most lucid thinker of these problems for the humanities to date. In her seminal book on the subject, *Planned Obsolescence*, Fitzpatrick systematically interrogates humanities' peer-review practices in the age of the digital and concludes that what is required is a mode that is less certain of the merits of 'the stability that we've long assumed in the print universe' and one that is more adaptive to generative possibilities.[12] What Fitzpatrick addresses, in essence, is the problem of the fundamentally anti-collaborative nature of humanities research in most cases. At present, review is not usually a community endeavour but rather an activity that expects to see a final artefact in which no traces of the construction remain visible, in much the same way as I have traced some of the problems of anonymity here. Experiments such as McKenzie Wark's collaboration with the Institute for the Future of the Book on his 2007 *Gamer Theory* suggest, however, that while an online collaborative model currently solicits sub-optimal levels of participation, there can be merit in the process.[13] Indeed, a draft of Fitzpatrick's own book was first published through the CommentPress system,

open for comment and revision, before being released by New York University Press. Likewise, Palgrave Macmillan ran their own form of open peer-review experiment (although, for reasons of caution, they selected titles that had already been through a traditional process first).[14] From what is visible on the archive of that project, it seems there was, likewise, a low degree of takeup as it is very difficult to incentivise something without specifically targeted readers and deadlines.

Most importantly, though, I want to use my final words here to reiterate, but modify, my opening gambit. Fitzpatrick astutely notes that, in this case (and others), 'the system that needs the most careful engineering is less technical than it is social'.[15] Bearing this in mind, it is crucial never to succumb at any point to a techno-fetishism but always to consider whether technology facilitates desirable social changes.[16] The academy has built, over many years, systems for appraising the individual rather than acknowledging the way in which knowledge is collaboratively produced and, for the first time in many years, there may be an opening through which to address this. Open access does not require any changes to peer-review practices any more than the codex meant that readers had to abandon their palaeographic antecedents. There might, however, be practical ways in which a moment of technological change could enable us to see, with apologies for inverting Winston Churchill's well-known aphorism, that perhaps our review practices are not so wholly democratic, not so entirely objective, fair or community-based; that they may only be the best that have been tried, apart from all the others.

Overlay journals: editing as social curation

A further region in which new experiments are emerging in an online, open-access ecosystem is in the editorship space. This can be seen in the rise of a specific type of formation known as an 'overlay journal'. The term 'overlay journal', originally coined by the creator of arXiv Paul Ginsparg in 1996,[17] can be specified as:

> An open-access journal that takes submissions from the preprints deposited at an archive (perhaps at the author's initiative), and subjects them to peer

review. If approved (perhaps after revision), the postprints are also deposited in an archive with some indication that they have been approved. One such indication would be a new citation that included the name of the journal. Another could be a link from the journal's online table of contents. A third could be new metadata associated with the file. An overlay journal might be associated with just one archive or with many. Because an overlay journal doesn't have its own apparatus for disseminating accepted papers, but uses the pre-existing system of interoperable archives, it is a minimalist journal that only performs peer review.[18]

To unpack this a little, the basic idea of an overlay journal is that papers are deposited on preprint servers (i.e. the work is made publicly available in a repository prior to peer review) before an editor decides to arrange for peer review of a new manuscript. The manuscript is then reviewed in the traditional way, including the upload of a new postprint version (also to the repository) incorporating revisions if required. Finally, a version is uploaded that indicates that the paper was reviewed and accepted by an overlay journal and the journal itself adds the paper to the table of contents.

Overlay journals, though, have shifted in their definition and function. With the advent of the megajournal – a single, multidisciplinary journal containing thousands of articles, sometimes operating on a post-review framework, exemplified by *PLOS ONE* – the concept of an 'overlay' has a different role. Instead of adding peer review outright, as the material is drawn from a megajournal with some form of quality control built-in, in this situation an overlay journal could bring the curational role of the editor to the fore in a fashion similar to an edited collection.

This concept of an overlay journal, which derives (as does much of the OA movement) from theoretical physics, is highly significant in specific scientific disciplines. Indeed, in physics, where the subject can change very rapidly, it is important for the field that information is available as quickly as possible and it is important for the scientists that their author precedence be recorded. Now, it could be argued that, in many humanities disciplines, these needs do not exist or are less significant. I would like to close this section with some observations to the contrary, noting that precedence, turnaround and, most importantly, the explicit positioning of the editor as the centre of prestige here, could have very real and positive effects for the

humanities and that overlay journals are one such mechanism through which this could be rethought. Once more, I will note up-front that OA does not require any of these innovations; it merely provides a historical juncture where it may be possible to implement them should we so desire.

Speed of turnaround and establishing precedence

As noted above, peer review – in whatever form – remains an important part of the scholarly publishing apparatus. However, as anybody who has attempted to publish an article or book within a deadline for any kind of assessment framework knows, it can also be a lengthy process. Delays in peer review, except in the case of ineptitude, are hardly ever the fault of the publisher. Rather, it is that overworked academics take a long time to respond to requests, then go over the deadlines and require multiple follow-up emails to ensure timely review. In extreme cases this can take years in some disciplines. However, once this process is complete, some publishers then attempt to synchronise their hard-copy journals with their online efforts, thereby re-introducing the artificial page budgets that I pointed out earlier and causing further delays. In journals this can result in a total publication delay of years from submission and often even longer in edited collections.

While some have argued that this delay is acceptable in humanities fields, there are problems with such stances. First of all, one of the most common rebuffs to a need for speed is that 'nobody dies if work in the humanities isn't published in a timely fashion', supposedly in contrast to certain medical disciplines. I would contend, despite having touched upon such a claim myself earlier, that this may not always be the case. Is it not possible to conceive of humanities works that focus on immigration and hate crime – as just one example – in historical and/or cultural terms and that could change policy, thereby saving lives? If the answer is that the humanities can never have this kind of impact, then something is going very wrong given the ethical purchase that is often claimed for humanities work. To argue, conversely, that all humanities work must have this kind of impact is also nonsensical. Secondly, in many sub-fields, such as my own area of contemporary fiction, it is deeply frustrating to have to wait

two or more years for work to emerge when there is a dearth of critical material. Certainly, much of the material surfaces at conferences, but these are frequently closed enclaves and it seems bizarre that, with the advent of the internet, our means of communicating is restricted to either flying across the globe – thereby contributing further to the crisis of global warming – or waiting two years, with a further delay of at least two years before any future work will be citing the first piece.

There is also a further reason why it could be desirable to speed up the process of scholarly communication – for we might consider it thus, rather than being termed exclusively 'publication' – and that is establishing precedence. While it may be nice to envisage a scenario where ideas were free and there was no need to attribute authorship because researchers worked in harmony rather than competition, this is the road not taken. Indeed, academia becomes ever more competitive and the imperative for novelty – no matter how contentious that word may prove – is given increasing credence. While there have been many excellent arguments for the 'slow humanities' and the value therein, even the scholar who worked slowly but published quickly would be at an advantage over those who abhor speed in both camps. In short, regardless of how long it takes to produce an article or volume of research, for which there may be extremely good reasons, there seems little rationale for slower publication after that point. A 'slow humanities' may lead to more rigorous work, greater care, higher levels of feedback and better research. How slow publication, after peer review, might also contribute to any of those aspects is less clear.

This, then, is my argument in favour of preprints, with overlay journals (or even overlay book publishers) as a potential mechanism on top of this: speed of turnaround can have important social and intra-disciplinary ramifications. Making work available for the purposes of communication, rather than purely for certification, could help to foster a better international research community. Likewise, establishing precedence will serve scholars better within a competitive framework. Although it may be desirable to eliminate such a competitive environment, it is worth noting that opposition to faster publication turnaround through preprints does not seem to serve that goal either.

Restoring the editor function

To move now to the second point raised by the concept of overlay journals, it has been my contention throughout this book that prestige, while a pragmatic and useful filtering mechanism for both assessment and readership purposes, comes with many potentially damaging side effects, both social and economic, when applied at the journal/publisher level. The idea of an overlay journal, in which pre-reviewed material is curated by trusted and respected editors into issues that they feel will be of value for their readers, puts editors at the centre of the system and displaces the purely service function of a publisher.

There are, of course, problems with such a setup. For instance, how is such material to be cited without diluting the version of record? As the provisions currently stand, unless material is strictly delineated from original publications in an issue, it might also fall foul of the recommendations of the Committee on Publication Ethics (COPE) on duplicate publication and/or self-plagiarism. One potential solution is to specify precisely on the cover sheet the original venue of publication while placing the 'curated' item in its own separate zone within an issue.

The other advantage of this type of curational setup worth mentioning is that it provides us with a further metric. Pieces that are re-curated by high-profile and respected editors are likely to be valuable themselves. This then gives a mark of prestige to an article based not upon the original publication venue or brand but rather upon the academic 'brand' of the curating editor. This is already the case, at least partially, with respect to special issues and to edited collections. In this way, the value of different publishers will become perspicacious free of a historical legacy of prestige. It may well be, in this new world, that old publishers really do know best. Indeed, the specific reach provided by presses with a history of targeted dissemination and trusted brand is fairly clear, even as critiques are mounted of those systems and companies. However, without a shift away from brand at the publisher level being conflated with academic quality, rather than with the quality of the production and other aspects with which publishers are concerned, it is impossible to evaluate the field fairly. It is, therefore, worth closing this volume with a remark that

seems far from the concept of overlay journals that prompted this discussion but that is, actually, the logical conclusion engendered by such thinking.

It makes no sense for open-access advocates to be 'anti-publishers'. Publishers perform necessary labour that must be compensated and any new system of dissemination, such as open access, will require an entity to perform this labour, even if that labour takes a different form at different levels of compensation. One may be opposed to specific practices of extant publishers, or particular hypothetical publishing enterprises, but it is not possible to desire the functions of filtering, framing and amplification without there being a publisher, even if this publisher is also the author. Nor, I would suggest, should one be opposed to competition between publishers, which can foster benefits for researchers. What seems problematic, instead, are historical genealogies that make it difficult to modify the current system of publication, which was born in an era before the internet. It is also clear, however, from the investigations mounted throughout this book that the economics of scholarly communications are extremely complex and that there are some genuine risks for publishers who dip their toes into these waters, so caution is to be advised and the rhetoric of 'disruption' avoided. However, as the music industry has perhaps best shown us, the internet is not going away and to find new models sooner appears to be a more sensible approach.[19] As the opportunity cost of not venturing into these territories mounts, it becomes incumbent on researchers, librarians, publishers and funders not only to enter into dialogue about suitable transition strategies but also to ensure that our thinking is not bounded by what merely exists. Without a broader horizon of possibility for what our practice might look like, even in the face of pragmatic difficulties, we will not have lived up to McGann's call, the epigraph to this book, to mount a practical self-criticism.