

Canadian Journal of Learning and Technology

Volume 32(3) Fall / automne 2006

Commentary on "A Review of e-Learning in Canada"

Richard A. Schwier

Authors

Richard A. Schwier is a professor and Graduate Program Coordinator of Educational Communications and Technology at the University of Saskatchewan. He can be contacted by e-mail at richard.schwier@usask.ca

Talk about a struggle! This commentary seemed to refuse to be written. I've filled a few virtual wastebaskets with rejected prose, failed approaches and trivial commentary. Finally, I had to ask myself why I was struggling so much with this piece of work, and it came down to a couple of important things:

- The research team (hereafter called Team Concordia) that performed this review did a superb job of it. This, like many of the other products that seem to fly off the desks of Abrami and Bernard, is an impressive piece of research, despite the limited time and resources they had to complete it, and beyond the limitations of using an Argument Catalogue as their primary research methodology. So I can't adopt the role of stern critic and point out all of the failings and oversights in their research with a dismissive air of academic remove; they're too good for that (and besides, it would be un-Canadian to be so rude). I found that the report did what excellent research should do; it caused me to reflect on important issues in e-learning that have been moiling around in my mind for some time.
- If this paper isn't to be about criticizing flaws in their work, what should it be? What literary form should it take? Description? Argument? Conceptual flight of fancy? As a former English teacher, it made some old muscles ache, but I knew I didn't want to write traditional academic discourse.

So I decided this should be a serial essay. With apologies to Ralph Waldo Emerson for using the titling convention of "On...", I will draw a few observations of what I think are key issues in e-learning that come directly out of the report, and that may find their way into a future argument catalogue or meta-analysis.

Some of the things I'll discuss in this essay will be side conversations to the report we reviewed. Much of what concerns me is not included in the review, nor should it be included. The review by its design and intent is aimed at drawing persistent understandings from existing experimental (or sometimes quasi-experimental) studies that meet prescribed criteria. That is their strength and their contribution, and it is a considerable contribution indeed.

But it seems like some of the things that interest me the most are not yet at that level of research maturity. They are topics that haven't yet generated the volume or type of research necessary to allow the kind of compressed scrutiny that this type of review and meta analyses are built to analyze. In fact, much of the research that I'm drawn to these days isn't asking questions that will ultimately provide much if any useful data for generalizing results. But I do think they have important things to say, or at least, they are interesting things to talk about that relate obliquely to this report. Let's take a look at some of the issues that aren't ready for prime time in the world of effect sizes.

On Evidence and Authority

I want to briefly discuss the notion of evidence-based decision making vs. evidence for whether using technology is a good or bad idea. I am glad the authors debunked the Maclean's story on "How computers make Our Kids Stupid". It was an inflammatory, silly charge—in my experience not an unusual feature of cover stories from Maclean's—but we shouldn't be drawn into a useless and divisive debate by trying to answer, "Does technology make our kids smart?" Certainly this team does not fall into the trap; they're far too sophisticated and experienced to allow that to happen. They look for evidence of what practices and procedures work, what we do and don't know with any level of confidence, and what questions need to be addressed. These questions are laudable, reasonable, disciplined and altogether necessary. We can make good use of evidence to support what we do, and a well-designed meta-analysis is an excellent way to distil something useful out of a confusing array of studies.

But educators often do fall into the trap of being defensive, and of trying to provide proof that they are acting professionally when they use technology. We ask ourselves whether technology is a good idea and not should we use it? Should School Boards invest in it? Should teachers be trained in its proper use? All of these questions, in my opinion, miss the mark entirely for a number of reasons:

- **They assume that technology is an entity—an object—that can be plugged into a learning environment, and by its very presence create differences. I don't think so.**
- **The questions assume that technology is one thing, not an array of things, and certainly not a process. We know that soft technologies are more influential than hard technologies, yet we continue to pay homage to hardware as if it has anything to do with anything.**
- **The issue is also one of authority, as Bob Heinich (1995) pointed out some years ago. As long as I've been in this business, I've been asked for evidence that technology is the best way to go about doing things. Should we use technology in this classroom or that learning environment? Can you tell me if it is worth the investment?**

Well, no, I can't answer those questions. I can tell you some ways technology has been shown to work well, and I can share some creative ideas about how to integrate technology into various learning environments, but I

can't prove to anyone's satisfaction that it is the best thing to do in any particular situation.

So, I believe the decision to use technology is based as much on authority as much as evidence of efficacy. Do we have sufficient authority to suggest that particular approaches should be used, and what is the basis of that authority? Is that authority based on evidence that is sufficiently robust to predict outcomes with any certainty? Is that authority based on suasion, experience and professional trust?

I'm conducting a seminar this afternoon with a group of new instructors at our university; we will be discussing learning communities (not necessarily technology-based) and how to promote student engagement in classes. I intend to use a bit of lecture, a smattering of humour, a series of questions and even a few examples of spectacular failures I've experienced over the years. If someone asked me whether the approach I have decided to take is the best possible approach, and to defend my decision from evidence, I would be at a loss. But, of course, nobody will ask. Why? Because I have the authority to make these decisions. There is an element of trust that has been established between me and the people who asked me to do this session, and because I am doing things that people expect. I will seem to be doing the right things because it fits within the range of expectations of the group, not because I can defend my decisions from data. I have the professional authority to make those decisions, and that professional authority is at least partially grounded on the assumption that I have defensible reasons for doing what I do, based on data perhaps, but also based on experience and reputation.

But when we get to innovations, such as some of the emerging ways technology is being used in education, we seem to be challenged to provide data-based evidence for what we do. Why? I think it is because innovative approaches fall outside of the usual expectations of educators, so they are a bit threatening. And I think it is because there is less trust that we know what we're doing when we use technology—it is seen as a silly, showy, expensive and wholly unnecessary add-on, not as an important option for making a difference to learning. Our trust has been eroded by technology enthusiasts whom have over-sold and under-delivered on the promises of technology for decades. As a result, we don't have the same level of authority. Our reasons for using technology may be just as solid and defensible (or just as wishy-washy and indefensible) as our reasons for using a Socratic approach with our students, but we are more likely to receive challenges for using technology.

I'm arguing for two things here:

- We need solid information to help us make decisions about what we are doing, and the data we currently have has been less useful than it might because it is buried in a stew of studies that yield rich but incoherent information. Reviews such as the one done by the Concordia team help us build coherent understandings out of a slice of the studies that meet the necessarily strict criteria to be included in the meta-analyses. They provide a type of authority that is important to evidence-based decision making.
- But we shouldn't throw out or ignore the richness that exists in the excluded studies. In some research circles (not Abrami and Bernard's), there is an attitude of snobbish contempt for studies that don't meet criteria—other people's standards—about what is legitimate, rigorous or worth attention. I think this is foolish because our authority for making decisions about how we perform in the classroom—whether that classroom is distributed, blended or face-to-face is based on more than hard evidence. It is also informed by intuition, creativity, stories of success and failure and just plain good sense, among other things.

On Narrative as a Complementary Approach to Provide Insight and Develop Research Questions

The report calls for increased attention to special needs students, academically advanced students, issues of gender, ethnicity/race/religion, and aboriginal education. In order to get at these thorny issues, I want to suggest that narrative research can offer valuable insights. We need stories—powerful stories to provide understanding, empathy and subjectivity to the research agendas we pursue. I like the clarity that quantitative approaches provide. But I also like the murky places, the swamps where so many interesting and important questions remain half-hidden from view.

Elsewhere, Campbell, Schwier and Kenny (2005a; 2005b; in press) have made a case for narrative as a form of critical inquiry in instructional design practice. The case hinges on the role of language in representing a social system of values, ideas and practices that a community of practice shares as a “commonsense understanding of the social world” (Walmsley, 2004, p. 3). This method asserts that meaning-making is contextual and relational, in other words, the research situation is constantly evolving and the account of it, the research story, is formed in relationship. Both researcher and researched are implicated in the narrative, and from it both learn something about the phenomenon, and about themselves (Murphy & O’Brien, 2006). The unstructured interview design, or collaborative conversation, thus captures the participants’ constructions and interpretations of experience, their families and social cultures, their seminal personal and professional encounters, their moral and ethical beliefs and dilemmas, their development and understanding of their work as instructional designers and how their knowledge is embodied in their relational practice with faculty. Conversations make a space for the researchers’ stories too, and the relational creation and re-creation of the social reality of the participants. The process typically explores four main ideas as methodological issues: reflexivity, voice, strong objectivity, and power/authority.

The problem, of course, is that it doesn't produce much of anything that you can use as evidence for making decisions. It won't likely show up in any argument catalogue or meta-analysis in the future, but it may just provide rich ideas that will feed other types of research programs—research that can inform decisions.

On What We Value as Learning Outcomes

Here was a passage from the report that particularly caught my attention:

Finally, we believe that more emphasis must be placed on implementing longitudinal research, whether qualitative or quantitative (preferably a mixture of the two), and that all development efforts be accompanied by strong evaluation components that focus on learning impact. It is a shame to attempt innovation and not be able to tell why it works or doesn't work. In this sense, the finest laboratories for e-learning research are the institutions in which it is being applied. (p. 4)

I applaud this crisp and clear direction. I am in complete agreement with an emphasis on longitudinal, mixed method research that includes evaluations and a strong focus on learning impact. I hope funding agencies are listening, and I think they are, given some of the noises we have heard from SSHRC in the past couple of years. I would like to add one elaboration to the recommendations, and it has to do with how we think about learning impact. In my opinion, metrics of the learning impact of e-learning are mitigated by institutionalized views of what are legitimate learning outcomes in higher education. If we consider grades, performance on

examinations or other relatively narrow indicators of learning as the primary evidence of learning outcomes, then we are in danger of overlooking the most important possible learning outcomes. I won't belabour the point, because I suspect that the authors had a wide range of learning outcomes in mind, given the list of outcomes they originally coded for before compressing the codes because of high intercorrelations among them. But I'll offer a list of learning outcomes that Stephen Downes (2006) proposed:

How to predict consequences How to read How to distinguish truth from fiction How to empathize How to be creative How to communicate clearly How to learn How to stay healthy How to value yourself How to live meaningfully

This is not meant to be a comprehensive list, nor am I suggesting that they should have been part of this review. After all, the responsibility exists elsewhere for making sure that courses and programs of study incorporate these kinds of value-based outcomes. Presumably, measures of success in courses and programs would include salient outcomes. But I would hope these kinds of outcomes would find their way into future longitudinal research on learning impact, particularly when we look closely at the content of programs of study, although I realize how elusive they are and how problematic it would be to measure them reliably.

On Social Software and the Shifting Responsibility for Constructing Learning Environments

One thing that I noticed about the research that contributed to this report is that the learning environments are presumably formal, bounded and institutionally prescribed instructional contexts. This is natural and appropriate, given the focus of the review, but another area of growing interest could be characterized as "informal learning environments." People use and adapt technology to learn outside of the boundaries of formal learning environments. We see approaches to learning that ask learners to construct their own learning landscapes out of an array of sources, many of their own choosing. A course or program may be part of the landscape, but in many cases people are identifying their own learning needs, identifying a range of resources to address them, and assembling those resources into an aggregate that makes sense to the learners who build them. In other words, learners act as active agents in constructing learning environments that are personally relevant. These types of learning landscapes seem to be tied to social software as tools for networking the human participants in the personally built learning systems. We even see hints in this report concerning formal learning environments: Interestingly, among the Pedagogical Uses of Technology, student applications (i.e., students using technology) and communication applications (both Mean = 0.78) had a higher impact score than instructional or informative uses (Mean = 0.63). This result suggests that the student manipulation of technology in achieving the goals of education is preferable to teacher manipulation of technology (p. 2). The impact of social software is compatible with learning environments that emphasize conversation, collaboration and interpersonal engagement. To my knowledge, there is little information about informal learning environments and how learners marshal resources, and particularly social software, to address their own needs. But it seems to be an issue that is growing in importance given the advances we have seen in recent years in communication technology and social software.

On the Role Played by Instructional Designers as Agents of Social Change

The authors clearly point out that instructional design is often ignored or inadequately implemented in the design of instructional courses and programs. This neglect results in programs that waste money and effort, and that don't produce the intended results. Yes, I agree, and I also agree with the implication that the involvement of instructional designers is important for reasons other than to design courses and programs that meet specific outcomes. In addition to the important role instructional designers play in the design and development of instructional products and programs, they also act in communities of practice as agents in changing the way educational institutions implement their missions. Designers work directly with faculty and clients to help them think more critically about the needs of all learners, issues of access, social and cultural implications of information technologies, alternative learning environments (e.g., workplace learning), and related policy development. As such, through reflexive practice, interpersonal agency and critical practice they are important participants in shaping interpersonal, institutional and societal agendas for change. In essence, instructional design is a social construct and critical pedagogy, in which designers act as agents of social change. A cultural shift has been occurring over the past decade in education—a shift towards environments and approaches based on the ideas of social constructivism. An instructional designer's practice, to which self-reflection is critical, will reflect her or his values and belief structures, understandings, prior experiences, and construction of new knowledge through social interaction and negotiation. Our team, with Richard Kenny and led by Katy Campbell, conducted a four-year program of research to investigate the roles of instructional designers as agents of social change and transformation in higher education. We found that little of the extensive work describing the development of models of instructional design (e.g., Reigeluth, 1999) has been drawn from the lived practice of the instructional designer (Kenny, Zhang, Schwier, & Campbell, 2005). As a result, instructional design theory is not grounded in practice. We argued that it is important to examine the theoretical and experiential backgrounds of these agents of instructional technology, their personal understanding of and values related to learning with technology, and the relation of these to their practice. When asked to tell stories about their values and their practice, instructional designers spoke passionately about their roles as agents of social change and transformation. Would this type of research fit into the kind of analysis we see here? No, of course not; nor should it. The authors of the report seem to support qualitative approaches such as this, however they worry that the preponderance of qualitative research mitigates our opportunity to draw solid recommendations about what should be done. I agree that we need more hard data for making decisions, even though I have primarily been doing qualitative and mixed-method research in recent years. But as part of the balance of approach that Team Concordia calls for, we suggest that narrative approaches offer one way to get at new questions from new directions. We are a long way from finding out whether this direction of research will ultimately result in improved designs or performance; indeed, these questions may never be addressed, and that need not be the criterion against which the value of research is measured.

On a Repository of Failure

The authors observe that the literature in instructional design does not speak about ID failures and successes or the ways that ID is misapplied in e-learning designs. This is an important gap in the literature of instructional design, and one that has been identified by the IDT Futures Group (Bichelmeyer, Boling, & Gibbons, 2006; Hill, et al., 2004). Recently, Tech Trends introduced a special section of the journal called "The ID Portfolio," which is intended to present a critical analysis and description of a specific instructional design project. It is a fledgling attempt to deal with the very thing that the authors of this report noted. Over time, we will develop a significant collection of case studies and analyses of a wide variety of ID projects and products, including their

impact on learning outcomes.

But I'm even more interested in the idea of developing a repository of ID failure, and to my knowledge no such database currently exists. As instructional designers, we know very little about failed designs, and we very likely make the same mistakes in different locations because we did not see what went wrong in other, similar projects. Engineers study their mistakes—the bridges and buildings that fail—as important elements of their initial training and ongoing professional development. Most of us are familiar with why the structural failure of the World Trade Centre buildings happened, and many of us have heard about changes that have been introduced to avoid a similar catastrophe, so one can only imagine how animated the engineering faculties around the world were following that event to study it, understand it, and propose solutions. But what do we know of spectacular failures in instructional design? I'm guessing very little beyond the failures each of us has experienced personally and attempted to hide from view. Our products play out in private locations for the most part. When we have a dramatic problem with a design, we try to fix it, and if we don't succeed then we are, at worst, subjected to scathing internal evaluations and perhaps have a course or program shelved. But generally speaking that is all that comes of it. Our work as instructional designers is not public, even within the boundaries of our own profession, but what a valuable tool we would have if we could study the dramatic failures we have all experienced. Of course, there would be reluctance in many quarters to sharing that kind of information, and there are institutional reasons why people promote only their successes. But it seems like a good idea, and a worthy project for someone wanting to make a career out of the rubble of failed projects in ID.

On a Theoretical Framework of What Makes E-learning Tick

As Team Concordia points out, the existing research on e-learning is diffuse, and this makes it difficult to not only combine and compare results (even without effects sizes), much less figure out what are the key features that drive successful and unsuccessful implementations of e-learning. We don't have a grand theory or framework for situating our growing understanding of e-learning and its impact. My sense is that reviews such as this one and the earlier meta-analyses on specific topics (e.g., Bernard & Abrami, et al., 2004) give us good places to start in developing such a framework.

But developing such a framework is a complex, lengthy and tedious task, even on a narrower topic and smaller scale than "e-learning". At risk of shameless self-promotion, I will use some of our recent work as an example. In recent years, our team at the Virtual Learning Communities Research Laboratory has been working on developing, testing and refining (I won't go so far as to say "validating") a model of formal virtual learning communities (Schwier, in press; Schwier & Daniel, 2007b). We had a growing concern about whether "community" was a useful metaphor for understanding online learning environments, and whether there was any precision in the application of the metaphor. It seems as though the label of learning community is used widely and indiscriminately to describe a variety of online learning environments, from rigid prescribed online classrooms to informal online settings. In addition, while there have been a number of solid and valuable contributions to methods for evaluating online learning environments, they focus very sharply on specific perspectives of community, not a coherent view of what comprises and energizes community.

We felt there was a need for a theoretical framework to understand the phenomenon. To develop our framework, we are employing a variety of approaches to examine methods for determining whether a community exists, and if it does, to isolate and understand interactions among its constituent elements, and ultimately to build a model of formal virtual learning communities. We employed several methods, both

qualitative and quantitative, to answer these seemingly simple questions, including user perceptions of community (Sense of Community Index, Classroom Community Scale), interaction analysis (density, reciprocity,) content analysis (transcript analysis, interviews, focus groups), paired-comparison analysis (Thurstone scaling) and community modeling techniques (Bayesian Belief Network analysis) (Schwier & Daniel, 2007a).

We sensed that these methods could be used in concert to address the questions of whether online communities exist, what their constituent parts are, and how these elements interact. It has been a harrowing journey, and we have a long way yet to go, but we have developed a great deal of respect for mixed approaches and the unique contributions each makes to developing a robust model that can stand up to scrutiny and situate future research.

On Research Questions I Wish I Had Time to Address

My graduate students and I have a running joke. They'll raise a really fascinating question, such as, "If I were an Aboriginal instructional designer, how might my world view influence my process of doing ID?" I usually dodge the issue with, "There's a thesis in that question." With some classes, all I have to do is raise my finger, and they repeat it for me, in chorus. But it's true, isn't it, that there are many more questions than we are able to address. Some are more powerful than others, but most of them deserve attention.

I've heard the complaint, somewhat justifiably I suspect, that in education generally and educational technology specifically we have spent far too much time on trivial questions (insert my 1978 doctoral dissertation here). I don't have a franchise on what the good and worthwhile research questions are, but I do have a few that really burn in my soul.

Some issues are financial and logistical—how does one assemble the technological and personal systems necessary to construct and maintain a successful e-learning environment? But the more important questions center on the design, implementation, pedagogy and effects of e-learning, the socio-educational aspects of learning through distributed and interpersonal means of communication, the role of informal learning to support formal learning and situated learning. A few of the issues that invite investigation are listed below, although many more seem to arise every day.

- How do people select virtual learning communities and how do they make use of them for learning?
- Do voluntary members of informal virtual learning communities differ from those who are assigned to learning communities in formal educational contexts?
- What are the fundamental characteristics of successful and unsuccessful virtual learning communities?
- How do informal virtual learning communities recruit and maintain members?
- Are there rules of engagement or particular protocols for insinuating an individual into the fabric of a virtual learning community, and is the process contextually or culturally bound? Does the process mirror interpersonal group learning contexts?
- How does a new member of a community join discussions of established community members and develop a persona or reputation? Are there power relationships in virtual learning communities, and how do they interact with learning variables?

- In virtual learning communities that permit members to remain anonymous, how does the anonymity of participants influence the tenor of interactions and the satisfaction of the participants?
- What is the nature of learning in virtual contexts, and how do architects, active members and lurkers understand their experiences differently?
- How do political, social, educational and personal agendas interact in the development maintenance and alteration of virtual learning communities?
- Do virtual learning communities exhibit lifecycles, and what significance does this have for their design?

Concluding Remarks

Unless you were raised by wolves, you were already aware of the important research being done at Concordia University on these and related issues. This report is one of several, and ties in nicely with the major meta-analyses on distance education published by Bernard and Abrami teams in recent years. The kind of research that they have taken on is unbelievably difficult to contain and complete, yet they seem to do so regularly and well. Without the kind of synthesis these kinds of analyses provide, we would be at risk of wandering around for decades asking questions that have already been answered or failing to ask the kinds of questions that will lead us systematically to the answers we need next.

Endnote

This paper was supported by a grant from the Social Sciences and Humanities Research Council of Canada. The work described in this paper draws substantially on research programs that included several collaborators over the years, including Katy Campbell, Richard Kenny, Ben Daniel, Gene Kowch and Heather Ross. I gratefully acknowledge their contributions to any of the ideas that exhibit a shred of intelligence, and I assume full responsibility for any ideas that do not.

References

- Bernard, R.M., Abrami, P.C., Lou, Y., Borokhovski, E., Wade, A., Wozney, L., & et. al. (2004). How does distance education compare with classroom instruction? A meta-analysis of the empirical literature. *Review of Educational Research*, 74(3), 379–439.
- Bichelmeyer, B., Boling, E., & Gibbons, A. (2006). Instructional design and technology models: Their impact on research and teaching in instructional design and technology. In, M. Orey, J. McLendon & R. Branch (Eds.), *Educational media and technology yearbook 2006* (pp. 33–50). Westport, CT: Libraries Unlimited.
- Campbell, K., Schwier, R. A., & Kenny, R. (2005a). Agency of the instructional designer: Moral coherence and transformative social practice. *Australasian Journal of Educational Technology*, 21(2), 242–262.
- Campbell, K., Schwier, R. A., & Kenny, R. (2005b). Using narrative inquiry to explore ID as conversation. *Selected proceedings of the Research and Theory Division of the Association for Educational Communications and Technology*, Orlando, Florida.
- Campbell, K., Schwier, R. A., & Kenny, R. (in press). Conversation as inquiry: A conversation with instructional designers. *Journal of Learning Design*.

Downes, S. (2006). Things you really need to learn. Retrieved September 28, 2006 from <http://halfanhour.blogspot.com/2006/08/things-you-really-need-to-learn.html>

Heinich, R. (1995). The proper study of instructional technology. In G. Anglin (Ed.),

Instructional technology: Past, present and future (2nd ed)

(pp. 61–83). Englewood, CO: Libraries Unlimited.

Hill, J. R., Bichelmeyer, B. A., Boling, E., Gibbons, A. S., Grabowski, B. L., Osguthorpe, R. T., Schwier, R. A., & Wager, W. (2004). Perspectives on significant issues facing instructional design and technology. In M. Orey (Ed.),

Educational media and technology yearbook, vol. 29

, 23–43.

Kenny, R. F., Zhang Z., Schwier, R. A., & Campbell, K. (2005). A review of what instructional designers do: Questions answered and questions not asked. *Canadian Journal of Learning and Technology*, 31(1), 9–26.

Murphy, S. A., & O'Brien, A. (2006). Listening above the din” The potential of language in organizational research. *International Journal of Qualitative Methods*, 5(2), Article 10. Retrieved August 22, 2006 from http://www.ualberta.ca/~iiqm/backissues/5_2/pdf/murphy.pdf

Reigeluth, C. M. (1999). What is instructional-design theory and how is it changing? In C. M. Reigeluth (Ed.),

Instructional–design theories and models, Volume II: A new paradigm of instructional theory

(pp. 5–29). Mahwah, NJ: Lawrence Erlbaum Associates.

Schwier, R. A., & Daniel, B. K. (2007a). Did we become a community? Multiple methods for identifying community and its constituent elements in formal online learning environments. In N. Lambropoulos, & P. Zaphiris (Eds.), *User-evaluation and online communities* (pp. 29–53). Hershey, PA: Idea Group Publishing.

Schwier, R. A., & Daniel, B. K. (2007b). Implications of a virtual learning community model for designing distributed communities of practice in higher education. In C. Kimbel & P. Hildreth (Eds.),

Communities of practice: Creating learning environments for educators

. Greenwich, CT: Information Age Publishing.

Schwier, R.A. (in press). A typology of catalysts, emphases and elements of virtual learning communities. In R. Luppigini (Ed.), *Trends in distance education: A focus on communities of learning*. Greenwich, CT: Information Age Publishing.

Walmsley, C. J. (2004). Social representations and the study of professional practice. *International Journal of Qualitative Methods*, 3(4), Article 4. Retrieved August 22, 2006 from http://www.ualberta.ca/~iiqm/backissues/3_4/pdf/walmsley.pdf
