

An Incrementalist View of Proposed Uses of Information Technology in Higher Education

Marvin J. Croy
Department of Philosophy
The University of North Carolina at Charlotte
Charlotte, NC 28223
mjcroy@uncc.edu

Abstract

A number of national educational organizations and individual authors have called for the use of information technology to radically reform higher education. Several projections of how this reformation will unfold are presented here. Three different approaches to critically assessing these projections are considered in this article, two briefly and one in more detail. Brief consideration is given to an approach based on educational values and to an approach based on cost/benefit analysis. After some discussion of the strengths and weaknesses of these approaches, a third approach deriving from a theory of technology control (Incrementalism) is elaborated in more detail and is found to offer helpful criticisms of the called for revolution in higher education. Some recommendations for how these new technologies can be developed in responsible ways are also offered.

This work was supported in part by a grant from The National Science Foundation (SBR 96-17224) and by a grant from the Foundation of The University of North Carolina at Charlotte.

There is presently a growing chorus of calls for swift and radical change in higher education in America. Such calls have been heard in the past, usually from idiosyncratic and seldom acknowledged visionaries. More recently,

however, the calls for radical educational transformation have changed in tone. A number of established national organizations have begun putting forward their views on the future of higher education in America. On all counts, this future is tied closely to the growth of information and computer technology. Years of study and discussion are now being documented by organizations such as EDUCOM, CAUSE, and the American Association of Higher Education. CAUSE (The Association for Managing and Using Information Resources in Higher Education), for example, has produced videotapes to communicate its thinking about the future of American education. In one such videotape (Seeing Higher Education in the Year 2050), Arthur Levine of Harvard's Graduate School of Education and president-elect of Teachers College, Columbia University, predicts what is becoming a not surprising portrait of American universities.

I think that much of higher education is going to disappear. I think that the only institutions that will be left, if we were to look in 2050, might be residential small colleges, and they'll become places in which young people of wealth would find themselves. Bright young people would be sent for a chance for a broader, longer education. And I think we'll be left with research universities. I think the reason for that is that you learn research by apprenticeship and we'll need for people to do that. I can't see any reason that any other sector of higher education would last.¹

Ano:

purpose of envisioning and guiding future developments. One document available from EDUCOM, "Using Information Technology to Enhance Academic Productivity," is authored by William Massey and Robert Zemsky.² Their case begins with two observations. The first is that there will be a huge, probably exponential, growth during the next decade in the demand for courses built upon information technology. The second is that no matter how higher education reacts, information technology will profoundly change teaching and learning, much as the printing press did. On the basis of these assumptions, Massey and Zemsky construct two alternative scenarios for the adoption of

information technology within higher education and explore the consequences of each.

In o

Information technology thus enhances productivity, providing a better ratio of costs to benefits, of inputs to outputs. However, this will require the substitution of technology for human capital and labor, which is to say, for existing faculty, administrators, and their traditional activities. Massey and Zemsky recognize that important questions are raised by this proposal.

The question remains: "What does an institution do with the faculty hours freed up by capital-labor substitution?" The saved hours might relieve shortages elsewhere in the institution, but this outcome becomes less likely if the institution's markets are not expanding. No financial savings accrue if the hours are simply redirected to departmental research as has been traditional in many institutions.... Faculty might take over duties now performed by staff, or regular faculty might displace auxiliary faculty, or the regular faculty might decrease in numbers.³

More

mentors, counselors, curriculum designers, advisors, and evaluators. New technologies will allow students to separate these functions and to select and pay for only those desired. One must be careful to distinguish "contact" from "contact hours." "Some students will continue to want a traditional collegiate education with all of its socialization or "contact" while others will just want the certification, the "contact (or credit) hours."⁴ Education for some will thus become a process of mere "credentialing."

Once productivity is defined as the ratio of inputs to outputs, three paths to improved productivity can be identified. The first ("doing more with more") occurs when increases in benefits outweigh increases in cost. The second ("doing less with less") results whenever modest reductions in benefits can be achieved with significant reductions in cost. The third ("doing more with less") requires that greater benefits are produced while costs remain constant or are

reduced. Massey and Zemsky acknowledge that productivity improvements achieved with information technology thus far are cases of doing more with more, and that, while this reinforces the prevailing faculty culture, this alternative is clearly flawed. Scarcity of resources will unacceptably limit development, and the lack of cost containment will frustrate the enterprise. Not only is "doing more with less" the most promising path to productivity gains, but this path is virtually mandated by rising costs and increasing public scrutiny of education.

Under a second scenario, universities continue in a "business-as-usual" manner, implementing instructional technologies in piecemeal ways that support, rather than change, existing practices. Without that change, the authors forecast a number of ills. For example, the American public refuses to fund higher education at its current level of faculty-centered inefficiency, and ultimately the undergraduate education market is lost to innovative nontraditional providers. These ideas are echoed in Eli Noam's "Electronics and the Dim Future of the University" published in Science. Noam expects that once the university system's control over accreditation is weakened, "we may well have in the future a "McGraw-Hill University" awarding degrees or certificates, just as today, some companies offer in-house degree programs."⁵ In Noam's analysis, the role of economics is clear, and on these grounds it is very doubtful whether the present system can continue.

The American Association of Higher Education's Project Future is also putting forward its view of the coming revolution, particularly in respect to distance education.

In the past two years, American higher education's interest in distance education has exploded. Suddenly, the technology seems to be there; the economics look attractive; we're supposed to serve more students, especially adults, and find new markets and revenue streams. . . . Many roads, it seems, lead to distance education. The new interest in distance

education arouses both unrealistic hopes and unfounded fears. On the hopes side, the claim is that instruction mediated by telecommunications will bring new gains in productivity, that somehow we'll hike access and quality while reducing costs--a claim for which there yet is precious little evidence. Or we hear that technology is the route to new populations of learners in whose wallets there sits a financial bonanza--another unfulfilled hope.⁶

The:

exploring these technologies. In fact, some projects are already taking advantage of these opportunities. The Western Governors Association, led by Governor. Rohmer of Colorado and Governor. Leavitt of Utah, has designed a virtual university which began admitting students in the Summer of 1997. This effort has been supported by a \$150,000 contribution from the Education Management Group (a subsidiary of Simon and Schuster) and has received positive feedback from accreditation agencies. In addition, the Education Network of Maine, the largest in the nation, is already in operation and may provide a model for distance learning programs. Such programs are under development at dozens of American universities and elsewhere around the world.

These views of higher education's future make claims about the aims of education, the role of technology in achieving those aims, and the relevance and significance of educational values. In order to critically assess these answers, three different approaches to issues concerning technology and education will be presented here. The first of these eschews quantitative assessment and emphasizes value concerns while the second advocates quantitative methods of risk and cost/benefit analysis. While each of these orientations possess certain strengths, it will be seen that neither provides a complete response to the kind of optimistic, frequently hyperbolic, claims about education and computer technology witnessed above. For this reason, the Incrementalist approach will be introduced. Incrementalism has emerged as an alternative to quantitative,

synoptic methods of decision making and technology assessment. Once these approaches have been elaborated and analyzed, some conclusions about their usefulness in the present context will be offered.

Educational Values and Concerns Over Technology

Concern over the uses of computers in education accompanied the earliest development. As early as 1970, a study was underway to explore the effects of CAI (computer-assisted instruction) on the classroom behavior of young students.⁷ A few years later, one of the first computer attitude studies was undertaken.⁸ The foci of these studies are indicative of concerns which recognize values within education. These values are generally of two types; either they characterize the educational process (e.g., concern for the individual, the significance of human interaction), or they are expected to be adopted by students (e.g., fairness, honesty, respect).⁹ It is this latter category of values that will be emphasized here, but in either form there have been many studies addressing value concerns both in pre-college settings and higher education.¹⁰ Discussion of educational values is implicitly related to the distinction between the primary and secondary aims of education. Providing students with skills and knowledge relevant to some particular subject may be the primary aim of education, but more than this is expected. It is commonly expected that the educated person will be able to both cooperate with others and to compete fairly, to appreciate and participate in community discourse, to respect the diversity of viewpoints, to resolve disputes even-handedly, to make judgments independent of prejudice, etc. In short, we expect the educated person to have adopted general values such as honesty, responsibility, respect, and fairness. Nevertheless, there is no university course in this "subject." Rather, these dispositions are learned via human interaction and the modelling of exemplary behavior, much of which occurs within classrooms in which other subjects are taught. This may make the "teaching" of these values secondary

in some sense, yet even so they remain an essential component of education. Likewise, the social dynamics of student-teacher interaction may play out primarily within the classroom and secondarily outside of that formal setting, yet that secondary interaction may be invaluable. For instance, Pascarella's review of over thirty studies concluded that "significant positive associations exist between extent and quality of student-faculty informal contact and students' educational aspirations, their attitude toward college, their academic achievement, intellectual and personal development, and their institutional persistence."¹¹

Find

"secondary" outcomes are actually the primary aim of education. In The Ethics of Teaching, Strike and Soltis make exactly this point.

In our view, growth as a moral agent, as someone who cares about others and is willing and able to accept responsibility for one's self, is the compelling matter. Promoting this kind of development is what teachers ought to be fundamentally about, whatever else it is that they are about. We are first and foremost in the business of creating persons. It is our first duty to respect the dignity and value of our students and to help them to achieve their status as free, rational, and feeling moral agents.¹²

Some have expressed concerns about the impact of distance education and other forms of computer technology on such aims, whether construed as being primary or secondary. Cuban's historical overview of technology in education, for instance, stresses that "researchers lack evidence that children exposed to machine interaction over long periods of time develop the full range of values, knowledge and skills expected by parents and the community".¹³ The point here is not that computers can never be used in ways that secure these outcomes, but that these outcomes are linked to educational values and that they may be endangered by technological change in education.

By distinguishing the aims of education, the values perspective encourages an appreciation of both the variety of these aims and the importance of the wholistic, values-centered nature of education in general. Nevertheless, expressing concerns about possible negative effects of technology on the transmission of educational values rarely helps to determine which forms of technology, if any, should be implemented in particular settings. This is because the mere possibility of harm carries very little weight in such pragmatic contexts. Not only is the likelihood of negative consequences unspecified, but no effort is made to quantify the impact of the potential negative consequences. Decisions to adopt particular forms of technology usually depend upon explicit factors, and these are rarely supplied in the arguments of those expressing concerns about the possible impact of educational technologies. Moreover, the vagueness and generality of such arguments makes it difficult to distinguish those which are better supported from those which are less well supported. It is for these reasons that a more formal, quantitative approach to technology implementation is often sought.

Risk Assessment and Cost-Benefit Analysis

Another approach to evaluating claims about the adoption of educational computer technology emerges from cost-benefit analysis and decision theory.¹⁴ This approach, often associated with Bayesian models, finds particular application in the assessment of risky technology and is relevant to the claims put forward by Massey and Zemsky, Noam, and others. Their claims are comprised of, first, predictions of future events given certain conditions, and second, recommendations of what actions should be taken to achieve certain goals. Bayesian theory characterizes decisions as a choice among alternative options while taking into account potential states of the world. Each outcome (combination of option and potential world state) has a certain worth associated with it, often determined by cost/benefit analysis and risk

assessment. When this worth can be properly quantified and the probabilities of potential states of the world are known, a simple calculation can usually identify the best option. Various decision rules, such as maximizing utility or maximizing the minimum, dictate how the best alternative is to be selected. Opponents of this procedure often balk at the prospect of quantifying the value of certain outcomes or of being precise about the probability of certain states of affairs. In this vein, Shrader-Frechette defends risk-cost-benefit analysis, (RCBA), against two categories of attack, one empirical and one normative.¹⁵ On the descriptive front, critics claim that RCBA does not provide an empirically accurate model of how decisions are actually made by experts. Critics offer instead a model of implicit decision making. On the normative front, critics contend that RCBA does not provide an adequate model of how decisions ought to be made, primarily because it shares many of the weaknesses of utilitarianism. Due in part to the inability to quantify moral considerations, RCBA is blind to such factors as rights, obligations, and distributive justice.

In response to these criticisms, Shrader-Frechette maintains that many of the flaws of RCBA can be rectified or diminished. Ethical weighting can be introduced to offset any blindness to such moral concerns as equity of distribution, rights, and obligations. That is, alternatives which serve such considerations can be counted more heavily than alternatives which do not. Given this improvement, ethically weighted RCBA supports democratic decision making through the generation of multiple analyses serving different interests. Moreover, the explicit structure of RCBA contributes to the avoidance of arbitrariness and to the democratic control of technology evaluation. For these reasons, Shrader-Frechette concludes that RCBA is the best procedure available whatever its drawbacks.

The present question is whether ethically weighted RCBA or an educational values perspective can provide a basis for evaluating claims about

the future of education and technology. Each perspective has its own particular strength and weakness. The prospect for increased precision and quantification, of course, is one of the chief motivations for turning to RCBA. Yet it is unclear whether this turn facilitates the assessment of educational technology. The cases in which RCBA is claimed to be useful are typically cases of risky technologies whose dangers are recognized in one of two ways: either through track records which document their costs and accident frequencies, or through estimates of these based on their similarities and connections to other technologies. While Shrader-Frechette's proposals for improving RCBA and its role in the process of evaluating the acceptability of risks are very helpful in the context of risky technologies where the dangers are clear, their application to educational technologies is tenuous. These technologies are not clearly related to other risky technologies nor is much known about the dangers of their use in other contexts. The use of computers is often compared to past failures of radio and TV in education, but no clear evidence demonstrates that those technologies worked against educational values. So, information technology in education does not have an established track record of precluding the realization of important values. But more importantly, few of its relevant consequences and even fewer of their attendant probabilities can be reliably specified. Moreover, while the overall structure of RCBA is open to any values whatsoever, humanists have doubted that all worthwhile outcomes associated with higher education can be empirically measured. At the least, such quantifications will be pragmatically impossible to achieve. Given these reasons, the usefulness of the Bayesian decision and risk assessment model in this context is doubtful. Until the particular risks are empirically documented, it makes sense to make use of a framework which is constructed from non-risk, non-Bayesian concepts. Incrementalism provides such a framework and the version to be introduced here is David Collingridge's theory of technology

control. After elaborating that theory, its relevance to issues concerning information technology in education will be explicated.

Collingridge on the Intelligent Control of Technology

David Collingridge's views on controlling technology are based both on philosophical foundations and on empirical studies of trial and error learning. His early work, which focused on controversies over nuclear power and lead additives in gasoline, was given a philosophical foundation.¹⁶ That foundation derives from the epistemological views of Karl Popper. Popper's epistemology centers around the inevitability of error and the commitment to discovering and correcting error as a prerequisite to progress. Popper denied that the truth of scientific claims could ever be demonstrated. Nevertheless, their limitations could be ascertained, and reasons could be given for preferring one scientific claim over another. Collingridge, following this fallibilist line, makes a similar distinction. Collingridge denies that any decision (preference claim) can ever be justified, but he affirms that reasons can be given for favoring one preference claim over another. In support of his denial that preference claims can be justified, Collingridge attacks the Bayesian account of decision making. In support of his affirmation that preference claims can be rationally compared, Collingridge articulates the role of flexibility and corrigibility in decisions concerning technology control.

The

this model to technology control, for several reasons. For one, the Bayesian model finds application only to simple textbook examples. In the real world many of the model's assumptions cannot be met. The model requires that all relevant states of the world can be identified (forming a mutually exclusive and exhaustive set). Moreover, all options (likewise, mutually exclusive and exhaustive) must be known along with all payoffs for each outcome. Without these desiderata, the most that can be claimed is that some option is the best

of those known at the moment, but this reinforces thinking which is too myopic. The Bayesian model is particularly mistaken in forcing us to think of decisions as occurring in an instant. That is, the model is insensitive to the fact that decisions and the unfolding of their consequences form a process which occurs over some interval of time. During this interval new options unknown at the outset may arise. When these new options are better than any of those previously considered, the decision maker will want to modify the earlier decision. Nevertheless, earlier choices may serve to prevent the adoption of new, better options. The Bayesian model has nothing to say about this important aspect of decision making. Yet it is clear that, when selecting an option from a set of alternatives, one should be careful to choose such that one's future flexibility is not precluded by initial choices. The art of choosing options which maintain flexibility is the cornerstone of Collingridge's theory of decision making and the intelligent control of technology. As evident below, this theory is closely connected with his dim view of our ability to predict the future.

In re

Attempting to control a technology is difficult, and not rarely impossible, because during its early stages, when it can be controlled, not enough can be known about its harmful social consequences to warrant controlling its development; but by the time these consequences are apparent, control has become costly and slow.¹⁷

To escape this predicament, one must grapple with what Collingridge sees as the major horns of the dilemma. This requires either improving our powers of prediction or increasing our ability to change a technology once its flaws are apparent. Most contemporary efforts focus on the first alternative, the prediction horn, and research in Bayesian decision theory provides an example of this. Collingridge claims that these attempts are futile and that the best means of resolving the dilemma is by tackling the control horn. The trick here is to react to (unpredictable) difficulties as they arise but to avoid the failures

caused by the inflexibility of mature technologies. This is accomplished by developing technologies in ways that avoid rigidity and that maintain flexibility. One way of understanding this is through the concept of corrigibility. Decisions to implement particular technologies vary in respect to ease of correction, and preference should be given to decisions whose flaws can be detected quickly and corrected easily. In many cases, ease of correction will be a function of the time required to detect errors. A decision whose flaws go unnoticed for long periods of time will be more difficult and costly to correct, often as a consequence of entrenchment. Technologies become entrenched as they become intertwined such that changing one technology requires changing others as well. This underscores the significance of monitoring ("the continuous scrutiny of a decision's real consequences with the aim of finding error").¹⁸ One must constantly remain on the lookout for signs that a particular decision is mistaken. Decisions which keep one's future options open and which involve systems that are easy to control should also be favored.

Follc

these more recent views are supported by empirical studies of organizational behavior. Collingridge continues to reject the notion that decision makers can identify the one best option in a set of alternatives and to affirm the notion that mistakes are inevitable. The most effective way to minimize the cost of those mistakes is to develop systems in a series of small, slow changes whose effects can be quickly recognized and easily modified. Decentralized decision making and non-hierarchical organizational structures are also important ingredients of incremental development.²⁰ Whether these views can clarify technology issues in higher education is the question to be addressed next.

Applying the Incrementalist View to Educational Computer Technology

The emphasis on unpredictability distinguishes incrementalism from both the values perspective and the risk-cost-benefit perspective. These two

perspectives routinely paint different portraits of the future of technology in education. Technology critics normally adopt the values perspective while technology advocates adhere to a more quantitative approach. The issue then becomes who bears the burden of proof for these projections. Often, that burden is shifted to those who oppose technological innovations. Arguably, those who initiate technological change should bear the burden of proof for predictions of economic and social impact. But from an incrementalist perspective, this controversy is both fruitless and unnecessary. The focus should rather be on developing flexible and corrigible technologies. This development should take place through a series of small, slow changes to existing routines. Given this orientation, educational technology is seen as being young and unpredictable, still malleable, and in need of monitoring. Monitoring is perhaps the single most pertinent concept in this context. As Collingridge and other incrementalists recommend, we should expect to be surprised, and rather than passively waiting for problems to arise we should implement monitoring along with the technology itself. Two points on this issue should be noted immediately. First, the kind of monitoring required here differs from the standardly employed assessment of instructional computing, and second, effective monitoring forces an explicit consideration of value issues. Both of these points deserve elaboration.

Assessments of instructional software and systems standardly focus on performance (and occasionally attitudinal) effects. Even when concerned with more widely defined outcomes, these assessments are carried out via formative and summative evaluations. Formative evaluations aim at early feedback that will guide program improvement and often involve small groups of prospective users. Summative evaluation occurs later and often employs larger, sometimes representative groups of students. Its aim is to document actual program/system achievements prior to wide distribution. It should be clear that

neither of these forms of evaluation fulfill the incrementalist function of monitoring. Monitoring is intended to scrutinize actual implementations of a technology over a period of time in a variety of local settings. In respect to the variety of settings or implementation sites, monitoring does not assume generalizability by random selection. The aim is to assess the impact, perhaps by pre- and post-measures, in natural, possibly idiosyncratic, settings. This aim is much more pragmatic than is normally the case in educational research.

Designing a monitoring regime immediately raises the question of what to monitor for. This question is answered by reference to educational values. Valued outcomes and processes should guide the focus of evaluation. For example, Pascarella's finding of valued outcomes (both attitudinal and behavioral) associated with informal faculty-student contact suggests measures which should be taken when information technologies are implemented in educational settings. Values, both as predispositions to be adopted by students and as desired characteristics of the educational process, should be highly suggestive in directing the monitoring enterprise.

Taking an incrementalist perspective of educational technology highlights certain factors, raises certain questions, and provides an approach for addressing relevant ethical questions. One key question raised by Collingridge's framework is how best to instantiate flexibility in the context of educational technology. This question will take different forms given different uses of different technologies. In the context of distance education, one question will be how quickly and easily course materials can be modified. Massey and Zemsky and other proponents of distance education often emphasize courses which are built around lecturing and which can be easily "canned." But some courses involve learning by doing and forms of instruction other than lecturing, and many courses are in a constant state of evolution as subject matter and technique evolve. Massey and Zemsky themselves claim that "technology

provides more flexibility than traditional teaching methods," but the only grounds offered is that it is easier to reprogram information technology equipment than to retrain professors.²¹ This may be true for simple computer programs, but much more is required for modifying multimedia and other sophisticated forms of information technology. Modification may involve the editing of audio/video components, graphic images, computerized exercises, etc. The constant evolution of computer hardware, operating systems, and development software drives a near continuous process of updating. Changes in course content, the discovery of new knowledge, and innovations in teaching technique will be more recurrent in some courses than in others, but incrementalism is adamant in requiring that these factors be considered in both the design stage and the projection of costs. In addition to determining what forms of technology are more flexible in this context, the sources of inflexibility and entrenchment need to be identified. In the rapidly moving world of computer technology, change is virtually assured, but changing some technologies is bound to cost more than changing others.

The relative role of humans and machines and how different mixes of responsibility will affect flexibility is also a key question. Human judgment is noted for versatility and its ability to cope with unexpected novelty. Automated educational systems are noted for speed, consistency, and reliance on explicit, general rules. What combination of these human and machine capabilities is best suited for the degree of complexity, idiosyncrasy, and changeability inherent in the characteristics of students, subject matter, pedagogical technique, and educational institutions is yet to be determined. This determination cannot be made without a good deal of practical experimentation. It is important not to convince ourselves via a limited number of hypothetical scenarios that the answers to these questions are in hand.

This issue is related to the attempt to provide individualized instruction, often cited as a goal of educational technology. Nevertheless, this goal has gone mostly unachieved in the history of modern education. Educational technology does not create but rather implements individualization. Individualized instruction depends upon empirical discoveries of how student characteristics (in particular, measurable strengths and weaknesses) can be addressed by particular pedagogical techniques. These discoveries, although sometimes supported by data-collecting instructional programs, must pre-date their implementation by instructional technologies. For the most part, these discoveries have not been rich enough to support more than meager forms of individualized instruction.

In respect to trial and error learning, incrementalism recommends gradual changes in well understood routines. Massey and Zemsky's proposals call for much more radical and sudden change. Incrementalists claim that technologies which institute swift, radical change have a history of costly failures.²² If only a few universities adopted their plan, however, this might provide a small change against the background of all educational institutions. Differences in the extent of information technology adoption could provide a healthy diversity among universities, and this diversity itself could provide a form of flexibility. It is important to remember that all of this will be for nothing without effective monitoring. But in any event, it is doubtful whether this piecemeal change would satisfy Massey and Zemsky.

Incrementalism's dim view of the reliability of prediction in this context appears to be justified. Massey and Zemsky begin their arguments with a pair of "observations" (that demand for instruction based on information technology will grow substantially and that information technology will profoundly change teaching and learning no matter what universities do). In fact, these are predictions. Incrementalism urges that they be recognized as such and that

there is little reason to treat them as anything but speculation. By using the term 'observation', it is implied that the truth of these claims is apparent to all. This is not the case. A prediction's truth or falsity rests with the occurrence or non-occurrence of the predicted events. Prior to this, confidence in a prediction is based on the track record of the prognosticator and/or on the regularity of the system whose behavior is being forecast. Neither of those factors appear to carry much weight in this context. Massey and Zemsky's claims about the feasibility of doing more with less are not provided with any empirical support. No examples are provided of information technology that has achieved this goal. Economists are generally doubtful of new technology's ability to provide more productivity for less, and American businesses have had a difficult time documenting productivity gains due to computer technology.²³

In Conclusion

Incrementalism's doubts about the accuracy of technology forecasts calls attention to the slim empirical basis for projections of the impact of information technologies on education. The consequences of these technologies can best be ascertained via actual implementations, but this makes monitoring and careful guidance crucial. Monitoring should be directed, in part, by educational values, concerns, and concepts. The concept of education itself raises questions about the unbundling of educational services and the difference between a college diploma and a certificate of training. Determining what degrees of human and machine interaction and responsibility works better than others is also crucial. This can be learned by slow change, careful scrutiny, and diversity of approach. None of this will yield the swift, radical change called for by technology zealots, but it is our best hope for a system which serves educational needs first and technological innovation second, rather than vice versa.

Notes

1. CAUSE, Seeing Higher Education in the Year 2050, Videotape VID0022, (Boulder, CO: CAUSE, 1994). The CAUSE home page is at <http://cause-www.colorado.edu/>.
2. William Massey and Robert Zemsky, "Using Academic Technology to Enhance Academic Productivity," (Washington, D.C.: EDUCOM, 1995), 11 pages. This and other publications are available through the EDUCOM web site (<http://www.educom.edu/>).
3. Ibid., p.3.
4. Ibid.
5. Eli Noam "Electronics and the Dim Future of the University," Science, 270, 1995, 247-249.
6. Steve Gilbert "Introductory Remarks to 'Why Distance Education'," AAHE Bulletin, 48, 1995. The American Association for Higher Education homepage is at <http://www.aahe.org/>.
7. David Feldman and Paula Sears, "Effects of Computer-Assisted Instruction on Children's Behavior," Educational Technology, 10, 1970, 11-14.
8. Robert Hess and Maria Tenezakis, "Selected Findings from 'The Computer as a Socializing Agent': Some Socio-Affective Outcomes of CAI," AV Communication Review, 21, 1973, 311-325.
9. For a variety of views on this topic see Douglas Sloan (ed.), Education and Values (New York: Teachers College Press, 1980).
10. For works emphasizing pre-college instruction see any of the following:
Larry Cuban, Teachers and Machines: The Classroom Use of Technology Since 1920 (New York: Teachers College Press, 1986); Phillip Jackson, Life in Classrooms (New York: Holt, Rinehart, and Winston, 1968); Douglas Sloan (ed.), The Computer in Education: A Critical Perspective

- New York: Teachers College Press, 1985). For works emphasizing post-secondary education, see the following: Seth Chaiklen and Mathew Lewis, "Will There Be Teachers in the Future? . . . But Then We Don't Think About That," Teachers College Record, 89, 1988, 431-440 (reprinted in Robert McClintock (ed.) Computing and Education (New York: Teachers College Press, 1989), pp. 80-89); Marvin Croy, "Ethical Concerns in Computer-Assisted Instruction," Metaphilosophy, 16, 1986, 338-349; Marvin Croy, "Ethical Issues Concerning Expert Systems Applications in Education," AI and Society, 3, 1989, 209-219; Hubert and Stuart Dreyfus, Mind Over Machine (New York : Free Press, 1986).
11. Ernest Pascarella, "Student Faculty Informal Contact and College Outcomes," Review of Educational Research, 50, 545-595. Questions can be raised about the direction of the causal connections suggested by these associations, or even about the existence of any causal connections whatsoever. But the ultimate point here is that the nature of these associations should be thoroughly investigated before taking any steps that would undercut those potential connections.
 12. Kenneth Strike and Jonas Soltis, The Ethics of Teaching (New York: Teachers College Press, 1985), p. 63.
 13. Cuban, p. 94.
 14. There is a wide literature on risk assessment and management, and there is a long standing debate over the usefulness of cost-benefit analysis. The controversy outlined here is presented in Kristin Shrader-Frechette, Risk and Rationality (Berkeley: University of California Press, 1991).
 15. See Shrader-Frechette, Risk and Rationality, chapter 11.
 16. See David Collingridge, The Social Control of Technology (New York: St. Martin's Press, 1980) and David Collingridge, Critical Decision Making (New York: St. Martin's Press, 1982). Collingridge's basic

- concepts are unfolded sympathetically here, but for more critical treatments see Cassandra Pinnick, "Epistemology of Technology Assessment: Collingridge, Forecasting Methodologies, and Technological Control," Philosophy in the Contemporary World, 3, 1996, and also Marvin Croy, "Collingridge and the Control of Educational Computer Technology," Techne': The Electronic Journal of the Society for Philosophy and Technology, 1, 1996, 1-15 (<http://scholar.lib.vt.edu/ejournals/SPT/spt.html>).
17. Collingridge, The Social Control of Technology, p. 19.
 18. Ibid., p. 32.
 19. See David Collingridge and P. James, "Technology, Organizations, and Incrementalism," Technology Analysis and Strategic Management, 1, 1989, 79-97, and David Collingridge, "Incremental Decision Making in Technological Innovation: What Role for Science?" Science, Technology, and Human Values, 14, 1989, 141-162.
 20. Perhaps Collingridge's best articulation of this view is contained in The Management of Scale: Big Organizations, Big Decisions, Big Mistakes (New York: Routledge, 1992).
 21. Massey and Zemsky, p. 6.
 22. See the chapter length case studies provided in Collingridge, The Management of Scale. See also Joseph Marone and Edward Woodhouse, Averting Catastrophe: Strategies for Regulating Risky Technologies (Berkeley: University of California Press, 1986)
 23. See the Stanford Computer Industry Project's conclusions as reported in the Investor's Business Daily, July 23, 1996, p. A8.