


RAND

Information and Biological Revolutions:

Global Governance Challenges—Summary of a Study Group

Francis Fukuyama

Caroline S. Wagner



DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

Science and Technology Policy Institute

The research described in this report was conducted by RAND's Science and Technology Policy Institute for the Defense Advanced Research Projects Agency under Contract OPA-9215205.

ISBN: 0-8330-2807-3

RAND is a nonprofit institution that helps improve policy and decisionmaking through research and analysis. RAND® is a registered trademark. RAND's publications do not necessarily reflect the opinions or policies of its research sponsors.

© Copyright 2000 RAND

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from RAND.

Published 2000 by RAND

1700 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138

1333 H St., N.W., Washington, D.C. 20005-4707

RAND URL: <http://www.rand.org/>

To order RAND documents or to obtain additional information, contact Distribution Services: Telephone: (310) 451-7002; Fax: (310) 451-6915; Internet: order@rand.org

Preface

This report summarizes the issues that arose and the discussions held during the meetings of a 1998–1999 study group focusing on global governance of information technology and biotechnology. The goal was to bring a policy perspective to bear on a discussion of new technological developments through a series of free-flowing and exploratory presentations and discussions.

An important part of this effort involved bringing together experts from many different fields—journalists, policymakers, scientists, academics, business people—to discuss developments that will affect all of society. By bringing together such a variety of people, the organizers hoped to see whether people from different professions react differently to emerging technological developments. Each study group meeting featured a presentation by a different invited discussion leader, which either explored some aspect of information technology or biotechnology development or examined the capability of human nature or political structure to deal with new technology, followed by a discussion.

This report presents the findings that emerged from these meetings. It addresses a number of issues, with an emphasis on possible U.S. responses on a political or social level to critical technology governance issues. The body of the report summarizes the issues that emerged from the discussion. The appendixes distill the content of the various presentations and discussions.

Francis Fukuyama of George Mason University and Caroline Wagner of RAND's Science and Technology Policy Institute organized this study group and conducted subsequent analysis, with the assistance of Richard Schum and Danilo Pellitiere, graduate students at the George Mason University Institute for Public Policy. Shaun Jones, Defense Advanced Research Projects Agency (DARPA), Department of Defense, and Gerald Epstein, National Security and International Affairs Division, Office of Science and Technology Policy (OSTP), Executive Office of the President, requested this study and provided guidance for this project. However, the conclusions in this report are solely those of the authors and should not be attributed to DARPA or OSTP.

The Science and Technology Policy Institute

Originally created by Congress in 1991 as the Critical Technologies Institute and renamed in 1998, the Science and Technology Policy Institute is a federally funded research and development center sponsored by the National Science Foundation and managed by RAND. The institute's mission is to help improve public policy by conducting objective, independent research and analysis on policy issues that involve science and technology. To this end, the institute

- supports the Office of Science and Technology Policy and other Executive Branch agencies, offices, and councils
- helps science and technology decisionmakers understand the likely consequences of their decisions and choose among alternative policies
- helps improve understanding in both the public and private sectors of the ways in which science and technology can better serve national objectives.

Science and Technology Policy Institute research focuses on problems of science and technology policy that involve multiple agencies. In carrying out its mission, the institute consults broadly with representatives from private industry, institutions of higher education, and other nonprofit institutions.

Inquiries regarding the Science and Technology Policy Institute may be directed to the addresses below.

Bruce Don
Director
Science and Technology Policy Institute

Science and Technology Policy Institute

RAND

1333 H Street NW
Washington, DC 20005

Phone: (202) 296-5000

Web: <http://www.rand.org/centers/stpi>

Email: stip@rand.org

Contents

Preface	iii
Figures	vii
Summary	ix
Acknowledgments	xii
Abbreviations	xv
1. Introduction: Governance Challenges of the Dual Revolutions in Information and Biotechnology	1
2. The Promise and Challenge of Emerging Technologies	2
What Is Technology?	3
Electronic Communications and Computer Intelligence	4
Human Genetic Manipulation and Bioinformatics	8
3. The Problem of Governance	12
What Is Governance?	12
Changing Attitudes Toward Governance	13
Internet Domain Name Regulation	17
Biotechnology	19
Governance of the Dual Revolutions	21
4. Possible Government Actions: Conclusions and Recommendations	23
Broader Decisionmaking Models	24
Citizen Councils	24
Governance by NGO	26
Conclusions	27
Appendixes	
A. Evolution on the Timescale of Thought and Action: Darwinian Approaches to Language, Planning and Consciousness, and Some Lessons from Paleoclimate About How to Speed Up Evolution	29
B. Quantum Theory and Human Consciousness	47
C. Morals, Demonic Males and Evolutionary Psychology	63
D. Implications for Public Policy of the Information and Communications Revolution	79
E. Biology, History, and Social Organization	87
F. Triumph or Tragedy: The Moral Meaning of Genetic Technology	100
G. The Conversion of Genetics and Computing: Implications for Medicine, Society, and Individual Identity	116

Figures

A.1. The Brain's Circuitry	36
A.2. Three Neurons and Their Axons	36
A.3. The Layers of the Cerebral Cortex	37
A.4. Patterning in the Superficial Layers	38
A.5. A Superficial Layer Neuron and Its Synapses	38
A.6. Undecided Cells Receiving Input from Two Areas, in this Case, Representing Apples and Bananas	40
A.7. The Emergence of the Genus Homo	43
A.9. Two Stable States, Cool and Warm	45
B.1. Crystallographic Structure of Microtubules	50
B.2. Schematic View of Two Neurons Connected by Chemical Synapse.	50
B.3. Schematic Graph of Proposed Preconscious Quantum Superposition Emerging Versus Time in Microtubules	58
B.4. Schematic of Quantum Computation	59
B.5. Quantum Superposition Entanglement in Microtubules for Five States Related to Consciousness	60
E.1. Social Organization	94
E.2. Evolution of Brain Size	96
E.3. Social Organization After Transition	97
E.4. History as a Series of Swings of a Pendulum	98

Summary

This report presents the findings of a study group held during 1998 and 1999. A series of eight meetings explored emerging technologies and the governance issues they raise for the scientific and policy communities.

Technology and Governance

In the early part of the 21st century, the technologies emerging from the information technology and biotechnology revolutions will present unprecedented governance challenges to national and international political systems. These technologies are now shifting and will continue to affect the organization of society and the ways in which norms emerge and governance structures operate. How policymakers respond to the challenges these technologies present, including the extent to which developments are supported by public research funds and whether they are regulated, will be of increasing concern among citizens and for governing bodies. New governance mechanisms, particularly on an international level, may be needed to address these emerging issues.

The governance challenges are emerging because of the very nature of these technologies. Information and biological technologies have in common that their control and use are largely in the hands of the individual. The technologies that drove the industrial revolution are systematic and complex, and putting them into use requires collective action, social infrastructure, and technical know-how. Information and biological technologies do not have the same large-scale, systematic nature—making it harder to control their dissemination and use. The governance challenge is no longer democratic control over centralized systems—as it was in the 20th century, with such technologies as nuclear weaponry and energy, telecommunications, pharmaceuticals, medicine, and airlines—but governance over decentralized, distributed systems. The features that make these technologies different from and their potential benefits greater than those of other technologies increase their potential for abuse.

The mechanisms societies use to control, direct, shape, or regulate certain kinds of activities is what we mean by *governance*. Governance is almost always conducted by governmental bodies, although it can be carried out in other ways. Yet, the practical obstacles to governance of these new technologies are

tremendous. Success in governing them requires the cooperation of stakeholders, states, nongovernmental organizations (NGOs), interest organizations, and the average citizen. Within any decisionmaking process, commercial, defense, social, and individual interests will intermingle, and a consensus among many players may be integral to any workable outcome. Accordingly, two central questions seem relevant: Is society likely to call for governance in new technology domains, such as the Internet and biotechnology? What governance issues do these technologies raise?

Changing Attitudes Toward the Need for Governance

Two recent shifts in attitudes strongly influence the issue of governance within technological arenas. The first shift is the decline of conventional top-down governance models and an emphasis on applying privatization; deregulation; downsizing of bureaucracy; and private, market-based solutions to many social problems. This trend is especially evident in telecommunications and information technology (IT). The largely positive and beneficial nature of IT, coupled with the anti-statist attitudes of the late 20th century, have shifted attitudes toward technology more generally and disinclined many from considering regulation as an effective solution to the challenges new technologies present.

The second shift is a changing public attitude toward the conduct of scientific research and the resulting technological innovations that might best be summed up as follows: "Don't leave scientific decisionmaking to the scientists." The most important influence on this perspective may have been our experience with the advent of nuclear weapons. Many greeted this world-changing technological development with great alarm, which has led to the creation of an international regime to prevent the technology from spreading. In the United States, the 1995 Government Performance and Review Act concisely illustrates this trend toward greater societal interest in knowing the outcomes emerging from the scientific enterprise. New reporting requirements are being introduced despite continued protestations from the scientific community that such accountability is not practical and may even be detrimental to innovation.

These trends suggest that the public's perspective about science and technology has become increasingly sophisticated. There seems to be general recognition that regulating new technologies poses substantial challenges and often has unintended consequences that may be as troublesome to society as the problem the regulation was intended to prevent. Recognition also appears to be growing that technological innovation is not always necessarily benign and that some

regulatory actions have served societal objectives effectively. Accommodating both perspectives raises difficult and complex issues for those who would offer governance approaches.

Possible Approaches to Governance

A consensus emerged from the study group that a “top-down” approach to governance of these technologies would not be practical. In the realm of standard-setting, a bottom-up, informal approach could prove workable, given the incentives for participants to converge on a single standard. However, regulation is more challenging. Enforcement across a wide variety of countries is likely to present problems, especially when top-down intergovernmental mechanisms lack force or fail because governments are unwilling to pressure one another. Moreover, the extent of the control of these technologies and their applications that is or will be in the hands of the individual makes regulation particularly difficult. Given that many decisions about use and application will be made on an individual basis, it is hard to image any regulatory structure without wide “buy-in” from the polity.

Accordingly, one approach to regulating technologies like these might be to use a distributed decisionmaking model that would involve a significant number of organizations and users in deciding what technologies to support with research and development funds; what technologies need governance; what the norms of use and application should be; and whether, how, and at what level of formality to regulate technologies.

Another possible approach would be using citizen councils to make recommendations to higher-level, more formal governing bodies. One model might involve aiding the organization of hundreds of citizen councils across the United States (or even around the world) and encouraging them to deliberate the norms of use, regulation, and governance of technology. Using the networking capacities of information technology, such councils could conceivably deliberate and share ideas on a series of governance questions in a way that draws toward a consensus of views on how to manage and govern technologies.

A third model the study group discussed was governance by the actions of NGOs. In numerous recent examples, NGOs, empowered by low-cost electronic communications, have been able to act to achieve outcomes that sovereign nation-states, acting either alone or in concert, could not. However, since NGOs base their authority primarily on the voluntary choices of their members, this can raise issues of legitimacy and may be applicable to only a limited range of problems.

Ultimately, because the technologies emerging from the information and biological revolutions are inherently global, success in governing these technologies is likely to depend on some model that involves all stakeholders—states, NGOs, interest organizations, and citizens—to cooperate in developing governance norms or structures.

Acknowledgments

The authors wish to thank the sponsor of the study group, Shaun Jones, Program Manager, Defense Sciences, DARPA, who provided funding for this effort. Gerald Epstein, Senior Analyst, National Security and International Affairs Division, Office of Science and Technology Policy, Executive Office of the President, also provided guidance and direction for the project.

Particular thanks go to the scholars who gave presentations and led the study group discussions: William Calvin, Richard Wrangham, Bernardo Huberman, Robert Wright, Robin Fox, Roger Penrose, Stuart Hameroff, Leon Kass, and George Poste.

And last, by no means least, we owe a considerable debt of gratitude to the individuals, too numerous to list here, who participated in the study group discussions and who candidly shared ideas and opinions about the effects of the information and biological revolutions on global governance.

Abbreviations

ARPAnet	Advanced Research Projects Agency Network
AT&T	American Telephone and Telegraph
DARPA	Defense Advanced Research Projects Agency
ICANN	Internet Corporation for Assigned Names and Numbers
IEEE	Institute of Electrical and Electronic Engineers
IETF	Internet Engineering Task Force
IT	Information technology
ITU	International Telecommunications Union
NGO	Nongovernmental organization
NSI	Network Solutions, Inc.
OSTP	Office of Science and Technology Policy
PTT	Post, telephone, and telecommunication organization
S&TPI	Science and Technology Policy Institute

1. Introduction: Governance Challenges of the Dual Revolutions in Information and Biotechnology

In the early part of the 21st century, the technologies emerging from the information technology and biotechnology revolutions will present unprecedented governance challenges to national and international political systems. Of particular interest here are electronic communication and computer intelligence, emerging from the information revolution, and human genetic manipulation and bioinformatics, emerging from the biotechnology revolution. These technologies amplify human capabilities so significantly and so profoundly that they stand to alter fundamentally the very notion of what we think of as human. How policymakers respond to the challenges these technologies present, including the extent to which developments are supported with public research funds and whether they are regulated, will be matters of increasing concern among citizens and for governing bodies. New governance mechanisms, particularly on an international level, may be needed to address emerging issues.

The remainder of this report presents the results of a study group that discussed these issues. The report is organized as follows: Section 2 describes the possibilities for technological change that these technologies foreshadow. Section 3 discusses governance questions emerging from these changes. Section 4 presents policy implications. Finally, the appendixes summarize the presentations that the discussion leaders made during the study group meetings and the subsequent discussions held to explore these issues.

2. The Promise and Challenge of Emerging Technologies

The study group set out to examine developments in the information and biological revolutions that present particular challenges for global governance. It has already become evident that the current political and legal infrastructures are inadequate for dealing with global changes in information and biotechnology. The Internet has evolved into a global information network and has developed beyond its original purpose of sharing information into a global commercial trading system. Electronic commerce is straining existing trade regimes, protocols for the protection of intellectual property, and concepts of currency. It is creating problems and jurisdictional issues for taxation and regulation. A number of governance issues have already arisen, such as the fairness of the current system of allocating Internet domain names in an international environment.

In the future, problems related to information security will require a high degree of international cooperation to govern or resolve. These include both the use of the Internet for crime and the misuse of the network by public and private groups in ways that invade personal privacy. Some suggest that contractual relationships will replace regulation and trade protocols. What organization is capable of negotiating and implementing new rules or enforcing net-based contracts? What court of law will adjudicate international contracts agreed to on the global information infrastructure. Can these simply be folded into the World Trade Organization, or do they need a separate institution?

Recent developments in biological sciences, particularly in genetics, raise the question of international organizational and legal governance. Procedures that are judged ethically or medically objectionable in one country may become available elsewhere through market mechanisms, leading to the development of foreign sites where individuals may go to avoid regulations. Can existing organizational structures and laws be adapted to the products of biological science? Will new forms of political organization and law be needed to address these changes?

Similar questions will be raised as biological sciences and computer sciences converge into applications called *bioinformatics*. As science explores creating information technology that can be used as a human prosthetic—either worn on

the body or implanted under the skin—questions about when it is appropriate to use these technologies and under what conditions will arise. Science is also exploring the use of biological materials as information processors in objects, such as “biochips.” Technologists suggest that miniature biological sensors detecting chemical and biological information may soon be available that will be capable of providing instant feedback on individual or group activities and, further, of linking this information into ultrascale networked computing. How can abuses of these technologies, such as surveillance and large-scale information-gathering among the population, be anticipated and regulated or countered?

This section explores the possibilities and challenges that areas of technological change posing particular challenges to global governance may offer.

What Is Technology?

The term *technology*, when examined critically, is often used to refer to tools with which society or individuals are not yet comfortable. The more complex a machine is, and the more that collective knowledge is needed to sustain it, the more time it takes to assimilate this technology into broad social use. The more a technology requires infrastructure, the more likely it is that governance structures will be created to oversee or support the technology.

What does this mean for the technologies we examined during our study? In the case of information technology, the extent of social and technical knowledge needed to sustain the information system is large, as is the infrastructure required, but use of the technology is highly individualistic. Government regulation and private-sector standardization are highly active in these technology areas, although they have trouble keeping up with the pace of change. The rate of adoption has been relatively quick, perhaps because the skill of communicating does not need to be relearned simply to use the new tools.

One could question whether biotechnology should be considered a “technology.” It has anomalous features that make it hard to categorize. Recombinant DNA, for example, while resulting from a complex technique, is a simple product. Moreover, it is not a tool but a biological agent—it acts on its own once it is “programmed” to perform in a certain way. The development of biotechnology products requires extensive social and technical know-how but does not necessarily require a large infrastructure to be deployed. It is not clear what government regulation is required to support or control biotechnology (or even whether it could be controlled), and it appears that private-sector standardization efforts have not yet emerged in any real way. Fetal tissue research, for example

is not allowed to proceed using government funds, but private research groups make their own decisions about its use. The practice of biotechnology falls more nearly into the definition of a technique—the skill for doing something—rather than that of a tool. Tools are used to make biotechnology products, and the outcomes are certainly intended to extend human capacities and make the world more habitable. However, the nature of the product itself and the way it works set it apart from other kinds of technologies.

Information and biological technologies have in common that the individual is more in control of the use and application of these technologies than of many active and reactive machines, in which the systematic nature of the technology often requires collective action to be put into use. The fact that collective action is not required to use these technologies makes them particularly difficult to govern. Many hail the Internet for the “free-wheeling” way in which it operates—having been created from the bottom up and having little governance or control. Indeed, the Internet is often cited as a promoter of “true democracy” because it enables the individual to interact with others directly and in real time. Biotechnology, too, is seen as having special promise because it will tailor treatments and medicines to the individual and place control of certain biological processes in the hands of individuals.

The features that make these technologies different—and that make their effects orders of magnitude greater than those of other technologies that have emerged in the past 50 years (with the possible exception of that of nuclear weapons)—also make the effects of their abuse potentially greater than those of other technologies. Yet, the level of control that is in the hands of the individual makes social governance much more complex than for technologies that require collective action to build, use, or maintain. The problem that emerges is no longer to ensure democratic control over a large and complex centralized system but rather to determine how much governance is necessary for a decentralized, distributed system and how society can accomplish this goal.

The next subsections describe the key features that make these technologies both a promise and a threat.

Electronic Communications and Computer Intelligence

Electronic communication tools have developed over the past 50 years to the point where they are widely assimilated into society. Even so, social organization is still adapting to the use of these technologies. As with earlier

technologies, social adjustment and acceptance of electronic communication is lagging technological development.¹ The speed, ubiquity, and ease of use of electronic technology will continue to grow, and the cost of accessing such communications will continue to drop.

The growth of the use of networked communications has been a phenomenon in itself, as Bernardo Huberman noted in his presentation to the study group:

- In 1994, there were about 1,000 World Wide Web pages in existence in the world; in 1999, there are more than 400 million Web pages.
- In 1993, there were 5 million Internet users; by 1998, there were nearly 100 million.
- In 1992, there were 10 Web servers; in 1999, there are perhaps over 5 million.
- In 1999, traffic on the Internet is said to double every 100 days.

Projections for growth continue to be exponential. Currently, the Internet runs across telephone lines, with the packets of information constituting any particular message "switched" and routed by special computers that Internet service providers manage. Voice messages travel along the same trunk lines as data, but voices are switched and routed by separate switches that telephone companies maintain for this purpose.

In the near future, computers and communication capabilities will converge, with telecommunication companies transmitting far more data than voice communications. In addition, the Internet services, once limited to data, will carry voice. Eventually, one system will emerge, and that system that will eventually expand to carry voice, data, and video. To support these developments, a more robust optical-fiber backbone will be installed, eventually creating links more seamlessly than today with satellite-to-ground and satellite-to-satellite transmissions.² As bandwidth expands, real-time videoconferencing will become available, enabling people to "meet" in cyberspace without having to travel physical distances.

¹As an example, full utilization of technologies in the classroom requires putting educational reform measures in place. Teachers need to change the way they offer instruction and organize the classes. This change will come about slowly and then only if it proves to increase learning in the classroom. See Thomas K. Glennan and Arthur Melmed, "Fostering the Use of Educational Technology: Elements of a National Strategy," Santa Monica, Calif.: RAND, MR-682-OSTP/ED, 1996.

²Changes expected from these technologies are discussed in more detail in Steven Popper, Caroline Wagner, and Eric Larson, "New Forces at Work: Industry Views Critical Technologies," Santa Monica, Calif.: RAND, MR-1008-OSTP, 1998.

As the information revolution evolves, computers will become easier to use and will appear more “intelligent” in a human sense. Industry experts expect that, eventually, networking, computing, and other tools of communication will be miniaturized to the point that the tools themselves (but not their function) nearly disappear. The boxes containing computers will become so small that they are nearly invisible; human-computer interaction will be so facile that computers will require almost no skill to use. Intelligence and memory will be embedded in a host of other machines and environments around us, until most products are “smart” and interconnected.

Within 15 to 20 years, there may be another “computer revolution”—one that uses the rotations found in all atomic particles to take the place of digital switches. New computers based on *quantum computing* will use the rotation of atomic particles in the way the “0s” and “1s” of current digital computers to conduct computations. The result will be computing power that is blindingly fast and utterly ubiquitous.

Joined with emerging developments in three related areas, advanced silicon-based computers (and eventually quantum computers) will almost seamlessly interact with their environment:

- natural-language processing—using computing software that understands vernacular speech
- three-dimensional information processing—using sensors and object-oriented software to enable computers to move about within their environment
- enhanced computer “intelligence” software—having computers that can understand, reason, and react to their environment.

Advanced computing will also be networked so that intelligent control of the environment is available to almost anyone. Homes will be “smart”—they will monitor temperature, repair and clean themselves, and order and prepare food or perform other mundane chores. They will learn patterns of use by the homeowner or user and begin to take over simple tasks. The same intelligent design and ease of use will be true of other machines. Once people adapt to ubiquitous computing, they will cease to be aware of it as separate from them and begin to view it as part of the environment, much as we do with pens, electric lights, and eyeglasses, all of which, in their own time, were once technological breakthroughs.

While ubiquitous, networked computing and electronic communications hold tremendous possibilities for convenience, efficiency, and data capacity, these

capabilities also have significant potential down sides. First, there is an enhanced capacity for surveillance and a corresponding decrease in personal privacy. Considerable discussion on privacy issues has already taken place with regard to the use and marketing of personal information stored in computer databases. Advanced computing and networking raise a host of more specific potential abuses.

Second, there is a potential for increased criminal and terrorist activity. The opportunities for this type of activity are already evident, with money launderers, drug traffickers, hate groups, and pornographers being among the most innovative users of electronic communications. Internet commerce is inherently borderless, which makes the collection of taxes and the regulation of business activities much harder for national governments to carry out. There is also increased possibility for undermining social order by spreading false or misleading information that is difficult to counteract.

Third, cultural communities may find it difficult to limit what their members see and hear. In the United States, this has come up primarily as a matter of controlling pornography, with the U.S. Congress tacking the Communications Decency Act onto the 1996 Telecommunications deregulation bill, a move the courts quickly struck down. While many people in the IT community remain strongly opposed to any abridgments of free speech, the control of pornography and, particularly, of the access children have to pornography remain important issues for a large number of Americans. In other countries, concerns go further, toward the protection of other aspects of cultural heritage: The French, for example, have tried to regulate the use of the French language on French Internet sites.

Fourth, the Internet can influence civil society and social cohesion. As Ithiel de Sola Pool once observed, the shift from "one-to-many" mass communications to "one-to-one" or "many-to-many" conversations of the sort modern computer communications foster had the effect of fragmenting national dialogues by making them much narrower.³ That is, the offerings of the mass media during, say, the middle decades of the 20th century gave all Americans a common set of cultural experiences, whether through watching the Ed Sullivan Show or reading *Life* magazine and the *Saturday Evening Post*. Not just electronic communications but all forms of media have moved to much more-specialized niche markets, with 500-channel cable television replacing the three national networks and with tens of thousands of on-line discussion groups springing up. This shift toward

³Ithiel de Sola Pool, *Technologies of Freedom*, Cambridge, Mass.: Belknap Press, 1983.

narrower groups has had the beneficial effect of improving the quality of communication in many cases but also raises the possibility that the civil society as a whole has become increasingly fragmented.⁴ This means that, in the end, citizens have fewer and fewer common cultural experiences and points of reference, with possible negative implications for their broader ability to associate and work together as a political community.

Finally, there is the simple overload of too much information and too much communication. At some point, the capacity for communication will come up against the capacity of the human body to absorb and use information.

Human Genetic Manipulation and Bioinformatics

Human genetic research, and its application in genetic engineering, will offer three capabilities that are not possible today: It will enable deliberately changing human genes to alter damaged or disease-prone genes; it will enable changing genotypes in a way that is transmissible into succeeding generations; and it will replace or enhance human traits, such as strength or intelligence, beyond what is available in nature.

One milestone toward this end is fast approaching: In the early part of the next decade, participants in the federal government's Human Genome Project, as well as private firms seeking the same end, will complete the initial phase of mapping human genes. A deeper understanding of human genetics, combined with advances in human biology and organic chemistry, seems likely to change the practice of medicine in revolutionary ways.

Sometimes termed *molecular medicine*, the application of genetics to medicine promises to revolutionize that art by applying knowledge of human genetics to the treatment of individual patients. Understanding biological changes at the genetic and molecular levels promises to shift patient care from a standardized set of practices that treat symptoms of disease to individualized treatment for the underlying cause of disease.

As George Poste suggested in his presentation to the study group, "within ten years from now, someone who has a yearly physical should never die of a metastatic malignancy." This will be made possible by understanding the genetic basis of pathologies, joined together with diagnostics. Each person's

⁴For a broader discussion of this issue, see Francis Fukuyama, *The Great Disruption: Human Nature and the Reconstitution of Social Order*, New York: Free Press, 1999, Ch. 4.

health will be tracked according to his or her genetic propensities to develop certain diseases.

Chemical and biological processes are being researched for their ability to detect specific molecular and cellular events accurately and quickly, with the goal of enabling almost immediate diagnosis and treatment. Molecular medicine seeks to enable biological cures—rather than traditional chemical or surgical cures—for biological problems. Costs of diagnosis and treatment are projected go down, and outcomes will improve as diseases are caught and treated at earlier stages, with treatment tailored for individual success.

Currently, for example, certain cancer tumors may have as little as a 20 percent response rate to a specific treatment. Low-probability treatments are combined in hopes that one will work: Standards of care dictate that the treatment is tried as part of a “buckshot” approach to a cure. Molecular medicine will allow diagnosis of the genetic makeup of the tumor, enabling a targeted treatment with a much higher chance of succeeding. Combined with advances in imaging technology and sensors at the molecular level, medical practitioners will be able to diagnose precancerous tissue changes in a way that will enable early treatment and prevention.

As these changes occur in the practice of medicine, the health-care delivery system will need to adjust accordingly. The pharmaceutical and medical-device industries should also begin to shift, conducting more research on genetic treatments and producing fewer mass-marketed drugs.

The information revolution will enable molecular medicine in two ways. First, rapid diagnostics will likely be done using computers and information technologies. Computers will be used for complex diagnoses, and the Internet will be used to conduct remote consultations between physicians. Second, information about the patient’s full medical history will be instantly available through a storage system, either on a chip embedded in the patient’s body or accessible from a central database. These advancements in recording, manipulating, and transmitting biomedical data—made possible by the Information Revolution—will offer improvements in health care over the next 20 years in ways that are only vaguely dreamed of now.

Developments in technology at the convergence of biology and information—such as biosensors and smart materials—are poised to produce perhaps the greatest changes, and challenges, to social and defense systems. Technologists suggest that miniature sensors detecting chemical and biological information may soon be available that will be capable of providing instant feedback on individual or group activities. Joined with large-scale networked computing, this

information could conceivably be available to a broad group of users. Advances in software systems may soon offer the capability to detect previously unknown patterns within massive amounts of unrelated data. On one hand, these developments may increase capabilities in such important public-good functions as environmental monitoring and economic management, but on the other hand, the possibilities of using these technologies to create new forms of terror, conflict, and abuse of power appear to be quite significant.

The actual merging of carbon and silicon, starting with the creation of computer prosthetics and moving toward implantation of artificial devices, is currently being researched. The creation of "smart" prosthetics is being tested as a way to help handicapped persons develop more normal functioning. In its simplest form, widespread use of carbon-silicon technology may come on the form of a subcutaneous chip that will carry all of a person's medical information—particularly useful when an unconscious patient arrives at the medical emergency room. The fact that health care will depend increasingly on information means that access to this information will be essential for good medical care and for the efficiencies it will bring to the health care system. Another possible application for this technology is financial: an implanted ATM or debit card could be created that would carry a person's financial resources on it for instant access.

George Poste suggested that the most intriguing of the silicon-based applications is the ability to implant intelligence or knowledge. Research is under way to explore the limits of biomimetic reverse engineering to understand neurological processes and whether these can be digitized. It appears to be possible to monitor electronic traffic in the optic nerve, for example, because it is transferred as an electrical signal that can be digitized. If it is possible to digitize what is being seen, it would appear to be possible to digitize an image that one person, or a camera, sees and then transfer it to brain cells within a person or to another person or an electronic processor. Under this scenario, a blind person could "see" using a camera that transmits digital signals to the brain.

To the extent that human consciousness has a quantum basis, a possibility that Roger Penrose and Stuart Hameroff have suggested in their presentation to the study group, it may be possible to enhance brain capacities using the emerging principles of quantum computing. This would entail not a convergence of silicon and carbon—with all the attendant problems associated with the body rejecting foreign materials—but of using the brain's matter as the basis for computing. Although many scientists reject the reductionist notion that the brain is a type of computer, Penrose suggested that quantum computation with objective reduction already occurs in the brain and serves as the basis for consciousness,

proffering the possibility that the brain could be programmed to use quantum functioning to increase memory or reasoning power.

While these advances offer significant opportunities for reducing human suffering, they also hold potential for misuse. As Leon Kass pointed out in his study group presentation, the possibilities genetic profiling of individuals presents raise questions both for the individual and for society, in ways that society is not prepared to address. For example, Kass pointed out that "in many cases, practitioners of prenatal diagnosis refuse to do fetal genetic screening in the absence of a prior commitment from the pregnant woman to abort any afflicted fetus." Moreover, economic pressures to contain health-care costs will almost certainly constrain free choice in the face of certain genetic knowledge. Discrimination in insurance may compel certain kinds of genetic intervention that will change the human genotype for future generations. The opportunity for those with funds to access vastly superior health care will likely increase the discrepancies of care between rich and poor.

Finally, the most troublesome of genetic manipulations is the possibility of manipulating human nature itself. Current research suggests the prospect for germ-line manipulations, which could change the characteristics and behavior not just of individuals but of all of that individual's subsequent descendants. The moral questions this technology raises are too numerous and difficult to expand upon in this report, but numerous social misuses of this kind of technology are possible, from widening existing gaps in wealth and life opportunities to producing unintended consequences by altering important but poorly understood aspects of human behavior.

3. The Problem of Governance

One objective of the study group was to examine the public policy implications of technological advances and changing public attitudes toward science and technology, particularly the increased demand for accountability that publicly funded scientific research faces. These implications raise the question of *governance*: whether, to what extent, and how new domains of technologically enabled activity—particularly involving information and biological technology—will elicit public calls for some form of governance. These two domains in particular raise difficult governance issues: the Internet because it is inherently difficult to govern, and biotechnology because it is not only hard to govern, but also, by popular perception, is in need of a higher level of governance than many other kinds of technologies.

This section explores the particular governance challenges emerging from the dual revolutions, with a focus on the issue of domain names on the Internet and biotechnology issues. The section immediately below defines governance and describes why these technologies offer particular governance challenges to the international community.

What Is Governance?

Governance is the effort of human communities to try to control, direct, shape, or regulate certain kinds of activities. At one extreme, these communities can try to ban the activity altogether. Governance is almost always carried out by states, although it can be carried out in other ways as well. In continental Europe, for example, there is a corporatist tradition in which the state works with certain designated representatives of civil society, such as labor unions or employer federations, to achieve public ends. In the United States, the government has tried to off-load many governance functions onto either the private sector or onto nongovernmental organizations (NGOs).¹

For our purposes, there are two distinct forms of governance, which may be labeled *standard setting* and *regulation*. A system like the Internet, for example,

¹For an overview of the role of "third sector" organizations in modern government, see Lester M. Salamon, *Partners in Public Service: Government-Nonprofit Relations in the Modern Welfare State*, Baltimore, Md.: John Hopkins University Press, 1995.

requires standards so that various users and components can interact as smoothly as possible. Some sort of governing function must exist to set those standards. Regulation, on the other hand, is required to constrain activities that produce a private good at the expense of the public. The first form of governance is inherently easier to establish: Once a standard is up and running, others have an interest in complying with it. By contrast, a given regulation might be in place in one country, but other countries might not have such a regulation. For example, if most countries banned fetal tissue research or use of certain genetically modified organisms, some maverick nation might find it profitable to permit these things.

Either form of governance usually constitutes a subset of public policy more broadly. Particularly in support of science and technology, many types of public policy interventions have limited scope, seeking, for example, to stimulate the development of a particular technology, to promote information sharing, or to fix certain specific problems.

It should be clear from the outset that any effort to create governance institutions for either of the two technology areas in question must be international. Modern information technology, in particular, is inherently borderless: The Internet user does not care about the physical location of any given server; so it is possible to defeat an effort by one nation or jurisdiction to control or close down a site by moving it to another nation or jurisdiction. Indeed, the mobility of modern information has led some observers to argue that even a comprehensive international governance regime would be incapable of taxing, regulating, or otherwise controlling the flow of data and information. Biotechnology is less mobile but still presents many of the same challenges: For example, if one country wants to ban cloning or genetic manipulation of offspring, people who want such things can simply obtain them in another country without such regulations. It is useless, therefore, to think about governance except in an international context.

Changing Attitudes Toward Governance

It is clear that attitudes toward the workability and desirability of governance have changed markedly in the last 50 years of the 20th century, particularly in the United States. While the Great Depression and World War II ushered in an era of increasing government intervention in many aspects of economic and social life and the creation of the modern welfare state, the latter half of the century has reversed these trends and emphasized privatization; deregulation; downsizing of bureaucracy; and private, market-based solutions to many social problems.

This is especially evident in the field of telecommunications and IT more broadly. In the first half of the 20th century, virtually all countries that built telecommunication infrastructures did so with state-owned post, telephone, and telecommunication organizations (PTTs) that exercised monopoly control over their national systems. The United States used a private firm, American Telephone and Telegraph (AT&T), but heavily regulated it in the public interest through the Federal Communications Commission.² Traditional PTTs were highly centralized and bureaucratic and exercised dictatorial control over technology development and implementation in their respective national jurisdictions. They saw fit to set prices for services and equipment and to mandate universal service and a system of cross subsidies and were the sole representatives of the users of telecommunication services in any kind of international negotiations.

The international governance mechanisms set up in this period were similarly centralized and top down. The International Telecommunications Union (ITU) in Geneva was simply the supranational regulatory body established by national regulatory bodies through international agreement. Like the PTTs that were its clients, it too sought to set a single, authoritative set of rules regarding frequency allocations, tariffs, protocols, and the like.

It is safe to say that very few advanced countries would set up a similar governance system on either a national or international level if they had to build one from scratch today. Most countries have privatized their PTTs (as in the sales of British Telecom and Deutsche Telekom) or have deregulated their telecommunication sector (as in the cases of the 1984 consent decree that broke up AT&T and of the 1996 telecommunications deregulation act). Countries are much more likely to favor private firms over state-owned monopolies to provide telecommunication and IT services and to accept a higher degree of competition in the provision of these services.

The reasons for this shift are both normative and practical. On normative grounds, the modern, centralized welfare state was seen to be at a dead end by the 1980s, stifling creativity and innovation while wasting taxpayer money through mismanagement and inefficiency. Indeed, the central ideological divide has focused on the role of the state in regulating private markets and providing social welfare. In this regard, the United States has taken the lead, being typically exceptional in pushing the envelope of liberalization.³

²See Eli Noam, *Telecommunications in Europe*, New York: Oxford University Press, 1992.

³On American exceptionalism with regard to anti-statism and market liberalism, see Seymour Martin Lipset, *American Exceptionalism: A Double-Edged Sword*, New York: W. W. Norton, 1995.

There is a further normative reason for wanting to minimize the regulation of IT that is specific to this type of technology. IT is widely perceived as being supportive of the sorts of political values that are emphasized in the West. Cheap, ubiquitous IT—from phones, faxes, and radios to computers, e-mail, and the Internet—has been seen as promoting a deconcentration of power throughout the world and promoting the spread of liberal democracy.⁴ Modern communications was seen as critical in undermining the legitimacy of such communist states as the former East Germany and the Soviet Union,⁵ as well as the right-wing dictatorship of Ferdinand Marcos in the Philippines.⁶ In the future, it is widely believed, IT will help open up closed societies, such as the People's Republic of China and Singapore, whose governments have tried to control Internet use for political reasons.

Americans, moreover, have special reasons for wanting IT to spread around the world with as few constraints as possible. Americans stand to benefit from the IT revolution in power and economic terms. American companies dominate the global IT industry and are on the cutting edge of innovation; American cultural products and media, from CNN to Disney to MTV, will be carried by this equipment; and American values, both political and cultural, will be fostered as the world grows more electronically connected. It is natural that Americans should argue in favor of a minimal IT governance regime, since that favors both their ideology and their self-interest.

On a practical level, it became clear that the evolution of technology itself was making the old governance model inappropriate and even counterproductive. Governments established PTTs in the belief that a telecommunication network constituted a natural monopoly and that private markets could not be relied upon to run such a network in the public interest. By the 1980s, the advance of technology, in particular the advent of long-distance microwave transmission, convinced regulators in the United States that competition in long-distance service was feasible. By the 1990s, other innovations, such as cable, provided alternative routes for obtaining information services at the local level and led the U.S. Congress to permit a much more far-reaching deregulation.⁷ It became clear that PTTs and national regulatory bodies moved much too slowly to provide

⁴On this point, see Peter Huber, *Orwell's Revenge: The 1984 Palimpsest*, New York: Free Press, 1994, and Thomas L. Friedman, *The Lexus and the Olive Tree*, New York: Farrar, Straus, Giroux, 1999.

⁵Scott Shane, *Dismantling Utopia: How Information Ended the Soviet Union*, Chicago: Ivan Dee, 1994.

⁶Gladys D. Ganley, "Power to the People via Personal Electronic Media," *Washington Quarterly*, 1991, pp. 5-22.

⁷Peter Huber and Michael Kellogg, *The Geodesic Network II: 1993 Report on Competition in the Telephone Industry*, Washington, D.C.: The Geodesic Company, 1994.

useful standards and rules in this fast-changing area. This was doubly true at the international level. By the time an organization like the ITU could agree, for example, on a communication protocol, technology had moved on and made the new protocol obsolete. These regulatory institutions were seen as positive hindrances to technological advance.

What was true of telecommunications was even more true of computers, software, networking, and other aspects of IT. Computer technology grew up in an almost completely unregulated environment. IT did not particularly reward scale; in some parts of the industry, at some points in the development cycle, it has been better to be a small and nimble competitor than a large, hierarchical corporation. So, the specter of concentrated power that had driven regulation in other sectors was not so obvious here. The U.S. Justice Department's antitrust suits against IBM and AT&T in the 1970s and 1980s were widely seen as counterproductive, since both corporations were soon upstaged by smaller rivals. The underlying rate of technological innovation and change in the IT business was so rapid that the government was loathe to intervene, although it finally did so when such private firms as Microsoft and Intel developed quasi-monopolies in personal computer software and hardware. The government's pursuit of these companies is, however, controversial, and virtually no one advocates encumbering the IT industry in a regulatory framework like the one that surrounded telecommunications.

Similarly, the ITU has, by and large, given up trying to set standards for the IT industry; the standards that do exist come out of private industry consortia, such as the Institute of Electrical and Electronic Engineers (IEEE) or the Internet Engineering Task Force (IETF). These bodies are less hierarchical than the ITU; participation is voluntary and based on the desire of market participants to have a common set of standards.⁸

The fact that IT regulation is difficult and likely to be counterproductive has led many observers to argue that governance is not desirable and should not be attempted. Attempts to tax electronic commerce or the Internet more generally have met with particular resistance in the United States and with a great deal of special pleading that alternates between arguments that it is illegitimate to tax electronic commerce and arguments that it is technically impossible to do so. During the study group discussions of governance issues, this attitude was strongly represented among the participants.

⁸While it is true that the Internet itself and the use of the telecommunications protocol/Internet protocol was due to an act of government intervention in the form of the U.S. Department of Defense's creation of ARPAnet, once the Internet was privatized, further growth and technological innovation took place rapidly and in a highly decentralized form.

Internet Domain Name Regulation

IT does, however, pose a number of governance issues that call for the creation of new institutions. One clear case is the assignment of Internet domain names. Domain names (that is, Internet addresses ending in .com, .gov, .edu, etc.) have become necessary to electronic commerce and are commonly seen as extensions of corporate trademarks. As such, they constitute a form of property and require a mechanism for protecting property rights. A company, such as McDonald's or Coca-Cola, that has invested millions of advertising dollars in building up a brand name will not take kindly to another firm or individual registering mcdonalds.com or coke.com on the Internet.

In the early days of the Internet, the U.S. government assigned domain names. When ARPAnet was turned over to the National Science Foundation to administer and then privatized in 1990, a private contractor, Network Solutions, Inc. (NSI), was tapped to run the administration center, InterNIC, to register domain names.⁹ NSI simply assigned names on a first-come, first-served basis for a nominal fee. The value of certain domain names quickly created a secondary market for them, however, and new phenomena began to emerge, such as "cybersquatting," in which an individual registers large numbers of domain names in hopes of being able to charge the legitimate trademark owners for them. NSI had to navigate its way through a number of cultural and international issues. For example, while it tried to avoid assigning indecent words as domain names, it could not guarantee that the names it assigned would not be culturally offensive in any of the languages used on the Internet. There is a further issue, moreover, that countries other than the United States require country-specific top-level domains (e.g., .uk, .fr, .de), which some regard as indicating "second-class citizenship."

The international assignment of domain names should in principle be no different from international registration of trademarks, patents, and other forms of intellectual property.¹⁰ Formal governmental institutions and international legal agreements have been established for the registration and protection of these, i.e., a formal system of international governance. If the Internet had somehow been invented 50 years earlier, when government regulation was in greater vogue, it would seem likely that Internet domain name registration

⁹See John A. Zangardi, "Regulation of the Internet Domain Name System," unpublished paper, 1999.

¹⁰With the different that Internet domain names are inherently global. That is, it is quite conceivable that the name McDonald's can be simultaneously registered as a trademark to a U.S. fast food chain in the United States and to an engineering firm in Scotland; by contrast, there can be only one mcdonalds.com.

would have been turned over to a formal, hierarchical national regulatory organization in the United States and to an ITU-like body (if not the ITU itself) on the international level. Given that this issue came up in the 1990s, however, it was solved in a typically less statist form. In July 1997, the Clinton administration ordered the Department of Commerce to privatize domain name registration and created a private, not-for-profit organization called the Internet Corporation for Assigned Names and Numbers (ICANN) to oversee the process. ICANN, in turn, took away NSI's monopoly in domain name registration and proposed to turn the task over to five private companies that had passed its accreditation standard.

ICANN represents a strange and in some ways unprecedented hybrid form of governance. It is neither a formal government organization nor a purely private-sector group, like the IEEE. It seeks to build consensus informally, without enforcement powers and without the legitimacy that would come about had it been established as the result of formal intergovernmental agreement. The hope on the part of its designers was evidently to retain the informality and flexibility of the kinds of governance institutions that characterized the Internet itself in its early days, while meeting the needs of the Internet's increasingly heterogeneous users.

The question, however, is whether an informal body like this can adequately govern domain name regulation. Informal coordination mechanisms require a high degree of consensus over goals and methods. NSI has, naturally, opposed the dilution of its monopoly and has refused to recognize the ICANN's authority. There are broader questions, moreover, within the Internet community about the rules for selecting ICANN board members.¹¹ One can easily imagine these disputes growing more severe in ways that will throw this entire governance system into jeopardy. Why domain name registration should be limited to five companies, what the grounds were for accreditation, and how larger international disputes are to be settled are not clear. Supposing, for example, that a powerful nation, such as China, rejects ICANN's authority or complains about the way that it has assigned names to Taiwanese firms.¹² When such conflicts become too severe, they rapidly outgrow the capability of informal mechanisms and are shifted upward to an intergovernmental level.

The need for more-formal governance institutions may become evident even in fairly technical areas, such as the setting of standards and protocols. The IETF,

¹¹Zangardi (1999), p. 52.

¹²A group in the People's Republic of China has actually registered the domain name "taiwan.com," much to the annoyance of Taipei.

unlike a protocol-setting body within a traditional telephone company, is famous for being decentralized and bottom up. Ideas filter in from all over the Internet community; if they are workable, the IETF simply ratifies them, rather than mandating them out of the blue. Some people have described this as a Darwinian, evolutionary process that fosters rapid innovation. The problem, as in the case of domain name regulation, is that such solutions often do not scale well. The number of participants in IETF meetings tripled between 1987 and 1989 as the number of stakeholders increased.¹³ A body with a more diverse membership and conflicting commercial interests is likely to have much more severe difficulties reaching consensus than the more inbred, homogeneous one that characterized the Internet in its early days.

IT and the development of the Internet have therefore not ended the need for governance but rather have raised new issues (such as domain name regulation) that call for institutional solutions. There is, nonetheless, a presumption that whatever regulatory mechanisms are ultimately established ought to have limited scope, be flexible, and be based on broad consensus rather than hierarchical fiat. This is so for practical technological and economic reasons, as well as a normative one: Very few people perceive the worldwide spread of information technology to be a threat to important social values.

Biotechnology

The same cannot be said, however, for biotechnology, which the public views with much greater alarm throughout the developed world. Genetically altered foods have been in the marketplace for several years, and the U.S. Food and Drug Administration has regulated them from the start because of potential safety concerns. The level of regulation is considerably higher in Europe, which has sought to limit; control; and, in some cases, ban certain kinds of biotechnology outright. European efforts to restrict imports of U.S. biotechnology products, particularly the European Union's ban on beef that has been treated with genetically altered hormones, have led to severe trade disputes.

The reasons for this resistance to biotechnology are complex. In Europe, it is often related to memories of Nazism and the genetic experimentation that went on in the National Socialist period. Over the past 25 years, Europe has developed a very powerful green movement, which, for ideological reasons, does not like tampering with nature. There are also strong concerns about the safety of

¹³Zangardi (1999), p. 17

genetically altered plants and animals, fears that the popular press has stoked with stories about "monster tomatoes" and the like.

The contrast between European and American attitudes toward biotechnology is striking and has led to important misunderstandings on both sides of the Atlantic. Americans tend to be eternally optimistic about technology and its effects on society and regard those raising concerns about biotechnology as Luddites who are frightened of change. Ethical concerns about genetic engineering tend to focus on the narrow issue of safety and whether regulatory and testing procedures are adequate to permit the commercialization of biotechnology products. The commission President Clinton appointed to look into the ethical aspects of human cloning argued for a temporary ban but only on the grounds that the procedure could not be carried out safely. The presumption was that, once the safety questions had been answered, there would be no further objections to individuals cloning themselves. In trade negotiations, Americans tend to focus on safety issues. When the science on an issue is on their side (as it is in the case of beef hormones), they tend to interpret European motives as either dishonest, ignorant, or representing a covert form of protectionism.

As Leon Kass pointed out in his presentation to the study group, however, there are any number of reasons for societies to decide to restrict the use of biotechnology that have nothing to do with safety. Many people have religious or ethical reasons for not wanting to manipulate nature on a genetic level. When biotechnology advocates point out that we already alter gene pools in numerous ways (for example, through simple crossbreeding of plant and animal species), biotechnology skeptics suggest that we are on a slippery slope to ever-greater interventions that may some day alter some of the most basic characteristics of the natural world. The most important of these is human nature itself. As discussed above, so-called "germ-line" research of the future will differ from historical medical technology in its potential to alter human nature by affecting not just the individual to whom it is applied but all of that person's subsequent. Many proponents of biotechnology argue that this kind of problem statement is unduly dramatic and alarmist. The purpose of research in biotechnology is therapeutic: It aims to uncover what are now clearly understood to be the genetic underpinnings of such diseases as breast cancer, Alzheimer's, schizophrenia, and the like and to provide cures for them. Germ-line research, it can be argued, simply takes this form of therapy to its logical conclusion: If the propensity for a disease lies in a genetically heritable characteristic, what is wrong in principle with designing a genetic intervention to eliminate that propensity from present and all future generations that might suffer from it?

The fact that there is no clear answer to the last question suggests, as Leon Kass has pointed out, why biotechnology will be so hard to resist in the future: Any potential negative consequences of genetic manipulation will be intimately connected with positive benefits that will be obvious and measurable. Many people argue that we can draw a line between therapy and enhancement and that we can reserve genetic engineering for the former. But drawing boundaries in gray areas is much easier said than done. There is general consensus that some conditions, such as schizophrenia, are pathological; the problem is that there is no consensus as to what constitutes health. If one can administer growth hormone to a child suffering from dwarfism, why not to one who is in the fifth percentile for height? And if it is legitimate to give it to a child in the fifth percentile, why not to a child in the 50th, who wants to receive the clear-cut benefits of tallness?

All of this suggests that societies will want to control biotechnology much more than they have wanted to control IT but that establishing limits, norms, and principles will be extraordinarily difficult. The problem of implementation arises in biotechnology much the same way as it does with IT: Assuming that societies can agree on limits to the technology (e.g., cloning or designer babies), how can they make these limits stick in an increasingly globalized world?

Governance of the Dual Revolutions

As we have seen, the IT revolution has had beneficial effects in undermining authoritarian hierarchies and distributing power more broadly. In popular imagination, IT is seen as good for democracy, good for the economy, and (if they are Americans) good for the United States as well, since the United States dominates the global IT industry. On the other hand, many nonscientists regard biotechnology with much greater suspicion, despite any beneficial effects.

Efforts to control biotechnology will run into the same practical hurdles as the attempts to control IT. Globalization means that any sovereign state seeking to limit, say, cloning or the genetic manipulation of offspring, will not be able to do so. Couples facing a ban imposed by the U.S. Congress, for example, may be able to go to another country to have a cloned child. Moreover, international competition may induce nations to cast aside their qualms: If one country or region of the world appears to be producing genetically superior individuals through its relaxed rules on biotechnology, there will be pressure for other countries to catch up. The libertarian mindset and the absence of international governance mechanisms, which seem appropriate for the largely benevolent IT revolution, may be less appropriate for what may be seen as a more sinister

biotechnology revolution. But at that point, efforts to close the gate may be unavailing.

The twin revolutions may also influence one another interactively. The belief that technological innovation is necessarily beneficial is not universally shared, even by Americans. The advent of nuclear weapons, after all, was greeted with great alarm and led to the creation of a vast international regime to prevent the technology from spreading. But the largely positive and beneficial nature of IT, coupled with the anti-statist attitudes of the late 20th century, have shifted attitudes toward technology more generally and disinclined many people from even considering regulatory schemes. When a new technology, such as germ-line engineering or cloning, comes along that societies may want to control, we may lack international institutions that are capable of imposing global rules. Existing international institutions, such as the World Trade Organization, are reluctant or unable to regulate technology because such constraints are viewed as nontariff barriers to trades.

Indeed, the practical obstacles to establishing a formal, hierarchical governance system are tremendous. The world has changed enormously since the ITU was established in 1865. The number of actors on the global scene has more than doubled, and the levels of technological sophistication and economic development in states other than the Western powers have vastly increased. Any international regulatory scheme—not just of IT or biotechnology but of child labor, occupational health and safety, or the environment—would have to be hammered out by a group of culturally diverse countries with vastly different levels of development and economic interests. Fifty years ago, for example, no international conference would have regarded the views of China as significant; today, no international agreement can go forward without its approval.

4. Possible Government Actions: Conclusions and Recommendations

Because the technologies emerging from the information and biological revolutions are inherently global, success in governing these technologies depends upon the ability to enlist all stakeholders—states, NGOs, interest organizations, and citizens—to cooperate in developing governance norms or structures. A number of interests intermingle in the applications and governance of these technologies—commercial, defense, social, individual—and these various interests can be made to work to the advantage of cooperation. This section describes some of the policy options available to decisionmakers to aid in governing the technologies the study group examined.

As our RAND colleagues, John Arquilla and David Ronfeldt, point out, in the economic-legal sphere of governing new technologies, the primary concerns are commercial.¹ The incentives for agreeing on standards are relatively straightforward. Furthermore, given the overriding interests of many parties in ensuring safety, efficiency, and security of the flows of information and goods that can be traded as a result of electronic commerce, cooperation in the regulatory sphere may depend upon reaching agreement in making “substantive law”—agreeing on what constitutes criminal activity on the Internet:

Cooperation may also hinge upon acceptance of a body of administrative and legal procedure that would establish jurisdiction and allow enforcement of substantive laws designed to protect property and other assets, both in and out of cyberspace. In the information realm, agreement about such matters as territoriality, extradition, and the notion of “hot pursuit” may form a minimum basis for international cooperation—especially in the area of cyberspace-based territoriality. . . .²

In some cases, law can be decided on a cooperative basis and implemented with a top-down approach. In other cases, norms governing noncriminal use of the Internet may actually emerge from practice: Huberman’s research has suggested that a number of these governing functions will emerge from current practice,

¹John Arquilla and David Ronfeldt, *The Emergence of Noopolitik: Toward an American Information Strategy*, Santa Monica, Calif.: RAND, MR-1033-OSD, 1999, p. 57.

²Arquilla and Ronfeldt (1999), p. 57.

leading to self-regulation of the Internet without having to institute a governing agent.

In the social arena, primary concerns are more numerous and varied than those in the commercial realm, making global cooperation more difficult to achieve. The complexity and speed of technological change in information and biotechnology suggests that top-down approaches to governance will not work. Given the difficulty of achieving global consensus on governance in genetic manipulation, trade in body parts, cloning, implantation of silicon chips, enhanced intelligence and other emerging technologies, it may be worth sorting out which of these issues are worth attempting to address—if not directly with regulation, then in a global dialog to monitor and oversee their use and combat misuse.

Participants in the study group agreed that a top-down or positivist approach to governance of these technologies will not work. In the realm of standard- and norm-setting, a bottom-up, informal approach could prove workable. However, enforcing regulations across a wide variety of countries is likely to present problems, especially when top-down intergovernmental mechanisms lack force or fail because governments are unwilling to pressure one another. Moreover, the individualistic control of these technologies and their applications and use make regulation particularly difficult. Given that many decisions about use and application will be made on an individual basis, it is hard to imagine any regulatory structure without wide “buy-in” from the polity.

Accordingly, one approach to making decisions like these might be to use a distributed decisionmaking model that would involve a significant number of organizations and users in deciding what technologies should be supported with research and development funds; what technologies need governance; what the norms of use and application should be; and if, how, and at what level of formality technologies should be regulated.

Broader Decisionmaking Models

Two specific alternative governance mechanisms were suggested during the course of the study group: citizen councils and NGOs.

Citizen Councils

A possible approach to involving a broad swath of the polity in decisionmaking about technology would be to create citizen councils to provide recommendations to higher-level, more-formal governing bodies. One model

might involve aiding the organization of hundreds of citizen councils across the United States (or even around the world) and encouraging deliberation about norms of use, regulation, and governance of technology. Using the networking capacities of information technology, such councils could conceivably deliberate and share ideas on a series of governance questions in a way that draws toward a consensus on how to manage and govern technologies.

A centrally organized group—in this case, perhaps, the White House Office of Science and Technology Policy, National Science Foundation, or a public-private coalition—could provide the incentive for convening citizens, provide the council adequate information with which to deliberate, and be the repository and clearinghouse for opinions and ideas emerging from the councils.

Citizen councils like this have actually been used quite effectively in Europe. The mechanism of “consensus councils” has a long tradition of settling contentious matters in science. These councils, however, have traditionally been made up of recognized experts and professionals in the field to be considered, leading one observer to liken them to a “synod of Bishops.”³ In the late 1980s, however, the Danish Board of Technology redefined consensus councils as bodies of lay citizens that would be convened to consider the evidence on a particular science or technology issue, participate in public debates, and ultimately provide a consensus report of their findings and policy recommendations. The purpose of the process was not to dictate policy but instead to help the legislature understand where an educated population might stand on an issue before considering specific policies. The consensus council reports have not only shaped policies; industries have also used them to craft research agendas to avoid public opposition that might emerge after they have made significant investments, when it is difficult to change direction.⁴ This success led to the engagement of similar processes in the Netherlands, the United Kingdom, and Australia.⁵ A similar citizens’ panel was convened April 4, 1997, at Tufts University in the United States to consider the topic of “Telecommunications and the Future of Democracy.”⁶

It may be necessary to determine the proper scale for this sort of public discussion and consensus, particularly in a country as large and as populous as

³Peter Skrabeneck, “Nonsense Consensus,” *Lancet* (1990): 1446-1447.

⁴See Richard E. Sclove, “Town Meetings on Technology,” *Technology Review* Vol. 5, 1996, pp. 24-28.

⁵Sclove (1996) and Ian Schibeci, Robin Shaw, and Aidan Davison, “Genetic Medicine: an Experiment in Community-Expert Interaction,” *Journal of Medical Ethics*, Vol. 4, 1999.

⁶Sarah Hackman, “And Now a Word from Your Neighbors,” *Technology Review*, Vol. 6, 1997, pp. 6-7.

the United States. Similarly, it would be difficult for 15 people to approximate the demographics of the entire country. Also, finding a body viewed as nonpartisan may be difficult in the U.S. context. Still, it seems that this sort of considered lay opinion has clear advantages for governance that can complement such mechanisms as opinion polls and expert panels and reports.

Governance by NGO

Another model for future governance came up in James Rosenau's presentation to a similar study group held in 1998–1999⁷ and repeatedly in this study group discussions, was governance by NGO action.⁸ In numerous recent examples, NGOs, empowered by low-cost electronic communications, have been able to bring about outcomes that sovereign nation-states could not easily achieve, either acting alone or in concert with one another. One case of this was the intervention that human rights and other activist groups made on behalf of the Indians in Chiapas, Mexico; the rapid action of the human rights groups forced the hand of the Mexican government in that instance.⁹ More recently, Greenpeace and other international environmentalist groups were able to force Shell Oil to change policy both in Brent Spar in the North Sea and with respect to Nigeria. Other advocacy groups induced the sportswear maker Nike to make promises about the use of child labor. Environmental groups have also been active on a number of biotechnology issues, such as the dispute over beef hormones.

In each of these cases, the NGOs in question were able to change the behavior of a large multinational corporation, or in the Chiapas case, the government of Mexico, in a situation where state action was ineffective. Questions have, of course, been raised about the true effectiveness of this kind of intervention: Who will monitor Nike, for example, for compliance with the promises it made on labor standards? Nonetheless, NGOs have the capability of organizing quickly and transnationally in ways that avoid the bureaucracy and rigidity of conventional international organizations.

This model of governance does, however, have a number of problems. The first concerns legitimacy. One of the reasons that formal government institutions exist is to confer legitimacy on their decisions; in an age when legitimacy results from popular consent, democratic institutions, such as legislatures and

⁷A similar study group was also held in the 1998–1999 year. James Rosenau's presentation was based on his book, *Turbulence in World Politics: A Theory of Change and Continuity* (Princeton: Princeton University Press, 1990).

⁸See Jessica Matthews, "Power Shift," *Foreign Affairs*, Vol. 76, 1997, pp. 50–66.

⁹David Ronfeldt and Cathryn L. Thorup, *North America in the Era of Citizen Networks: State, Society, and Security*, Santa Monica, Calif.: RAND, DRU-459-RC/FF, 1993.

parliaments, are broadly accepted as a means of expressing consensus. Government by NGO, on the other hand, comes about simply as a result of groups being able to organize themselves effectively, with no requirements for being representative or promoting the public interest. Many existing NGOs claim to represent the public interest, of course, and many see themselves as offsetting the inordinate power of multinational corporations or governments in the pay of multinationals. But they are self-appointed tribunes of the public interest, whose actions may or may not correspond to what democratic publics really want. They are also unaccountable; unlike a democratically elected legislature that can be turned out of office in the next electoral cycle, NGOs cannot be removed by popular demand.

A further problem is that the kinds of issues that NGOs can deal with effectively are limited in many ways. Nike and Shell could be pressured by activist organizations only because they had large consumer marketing operations that could be hurt by bad publicity. Companies that sell to other businesses, or which market unique or difficult to obtain products, would be far less vulnerable to pressure, and indeed much less open to public scrutiny.

Conclusions

Today's governance structures are challenged by a unique shift from collective control and hierarchical decisionmaking to individual control and decisionmaking that will mark the technologies emerging from the dual revolutions. The very natures of these technologies make regulating and controlling them particularly challenging. Traditional "top-down" or positivist methods of governance will have little influence over how these technologies are developed, diffused, and assimilated. New governance mechanisms are needed, and they must emerge quickly, be flexible, and have broad buy-in. The alternative methods discussed here—standard-setting bodies, citizen councils, and NGOs—present some options to policymakers considering ways to deal with these challenges.

Appendix

A. Evolution on the Timescale of Thought and Action: Darwinian Approaches to Language, Planning and Consciousness, and Some Lessons from Paleoclimate About How to Speed Up Evolution

Leader: William H. Calvin,¹
Professor, University of Washington
June 2, 1998

Introduction

Questions about consciousness and intelligence arise in the context of both the information and biological revolutions. To understand these phenomena, this meeting began a two-part inquiry into the human brain. This summary was drafted by Richard Schum.

The goal of the study groups on the brain was to understand the its physical possibilities and limitations. In addition, we examined the physical origins of conscious thought and intelligence within the brain. This information is needed to put in context any technological enhancements to the human brain that might be researched and attempted. In addition, this information is key to understanding whether computers or networks of computers can ever achieve a kind of conscious intelligence.

As the information revolution unfolds, understanding the interaction between the biological brain and silicon-based intelligence is becoming an imperative. Electronics research and manufacturing continue to shrink computer chips (and hence computers) and increase computational speed. This suggests a future in which a minuscule computer with blinding speed and considerable memory would be available to nearly everyone on the planet. Recent research using biological and atomic-level processes to conduct computation suggests that the computers of the future may use, or be compatible with, biological material.

¹For additional information on this subject, please visit <http://WilliamCalvin.com>

Some have called these developments “bioinformatics” or biological processing, although the developments are so new that language to describe it is lagging.

Humans already use computers as tools. It is not hard to imagine a day when the technology will enable people to implant a tiny computation machine into the human body or, alternatively, to use human biological processes to do computerlike computations. This raises a number of questions, such as whether such developments might be desirable and to what extent such developments might affect individual and group human behavior. Could such developments be the next step in human evolution?

Although computers are far from “intelligent” in the sense that we think of intelligence today, a number of people are working on creating intelligent systems, such as neural nets. Some would suggest that trying to create a computer modeled on the human brain is misguided and doomed to failure. Even so, this research, as well as research on the quantum basis for consciousness (a subject that will be discussed in the October 22, 1998, study group), has preoccupied a number of computer scientists and brain researchers.

To understand the future of mind, brain, and computing, and to question whether a human carbon-silicon brain is possible or even desirable, we need to understand the nature of the brain, how it evolved, and what challenges lead the brain to develop new capacities. In addition, it is important to consider the capacity of the brain—and possible alterations to this capacity—within the context of challenges to and impacts on human behavior, public policy, and the social order.

William Calvin: Evolution on the Timescale of Thought and Action

To help initiate discussion of mind, brain, and computing, the study group leaders invited Professor William H. Calvin, Ph.D., to present his research into the human brain and evolution. Professor Calvin is a theoretical neurophysiologist at the University of Washington in Seattle. He is the author of nine books including *The Cerebral Code* (MIT Press, 1996); *How Brains Think* (Science Masters 1996); and, with the neurosurgeon George A. Ojemann, *Conversations with Neil's Brain* (Addison-Wesley, 1994).

Calvin's research interests include the recurrent excitatory circuitry of cerebral cortex used for split-second versions of the Darwinian bootstrapping of quality, the fourfold enlargement of the hominid brain during the ice ages, and the brain reorganization for language and planning. He has long been following the

paleoclimate and oceanographic research on the abrupt climate changes of the ice ages, hoping to find a connection to the “big-brain problem.” He has an amateur interest in prehistoric astronomy and the associated archaeology.

He recently returned from a stay at the Rockefeller Foundation’s study center in Bellagio, Italy, collaborating with linguist Derek Bickerton on their forthcoming book about the evolution of syntax, *Lingua ex Machina: Reconciling Darwin and Chomsky with the Human Brain*. His presentation centered on four themes: (1) levels of organization in the human brain, (2) using a Darwinian approach to understanding brain development, (3) the biological basis for a “Darwin Machine,” and (4) how knowledge about the brain might be used to enhance brain functions. This paper summarizes Calvin’s presentation.

Levels of Organization in the Human Brain

Artificial intelligence (AI) approaches to machine intelligence have traditionally been limited to the cost and capability of the computer. Originally, most AI work was analytical. The limited computational abilities of existing computers put a premium on efficiency and limited the ability of researchers to incorporate redundancy into their models. Recent advances in information technology are enabling the introduction of more randomness and redundancy into AI design. In an age of inexpensive computing and networking capabilities, AI is no longer a novelty—it has real applications. As a result, we may be able to model how the brain uses Darwinian processes—the same approach responsible for such biological adaptations as species evolution and immune response—to make decisions on the timescale associated with thought and action.

While computational speed is important, the key factor lies in levels of organization. Consider the following four levels: fleece, yarn, cloth, and clothes. Each is more highly organized and is the product of the one that came before it. For example, fabrics are woven of yarn, and yarn is spun. Each state is transiently stable and reflects a ratchetlike characteristic that prevents backsliding, or disorganization into its former state. Each level is causally decoupled from the next, so one can weave fabric without knowing how to spin yarn or make clothing with it. A whole set of techniques and body of knowledge exist within each level.

The organization of science reflects this kind of approach. For example, chemistry, the study of chemical bonds, can proceed pretty well without understanding anything about atomic spectra or about the Krebs cycle. Nevertheless, it certainly helps culturally for chemists to know something about atomic physics and biochemistry, even though they could function pretty well

without them. Similarly, some of the highest functions of the nervous system, like consciousness, may in fact constitute another level of organization with rules of its own.

Within the neurosciences, there are probably a dozen levels. Some researchers say that memory arises from a synaptic change, while others argue that it is the nerve circuit or even gene expression that changes. Ironically, they are probably all right. This multidimensionality of explanation is what happens when a branch of science spans a number of levels of organization.

Within the nervous system, individual neurons produce impulse trains that can collectively affect the level of spatiotemporal patterns in the brain. The representation of a particular memory is not a result of a single cell firing, but rather the firing pattern within a cell committee (a Hebbian cell assembly, as it is more properly known). Thus, the key to memory functions appears to be pattern recognition and increased cell organization, not the behavior of the individual cells. By way of analogy, consider a computer screen. The behavior of each pixel—whether lit or unlit—has no meaning in and of itself. The meaning is derived from the pattern that is created, not from the individual constituents.

Larger assemblies that go beyond Hebbian-sized groupings also exist and are probably on the order of several square centimeters. These may represent objects, actions, relationships, analogies, or sentences. Composed of individual elements that are about 0.5 mm in size, these “hexagonal mosaics” compete against each other and attract additional members, each adopting the spatiotemporal pattern of its neighbors. This process of quality shaping via ad hoc assemblies continues until a winner emerges.

With various territories competing simultaneously with one another for the limited space on the association cortex, winners and losers are determined in a kind of “playoff.” The winning pattern becomes the conscious focus of the mind, and the “losers” become secondary or subconscious thoughts. A succession of focus occurs when the content of consciousness shifts and a new pattern prevails. This explains how the right birthday present for your spouse might suddenly pop into your head in the middle of a meeting.

A major theme of competitions—whether conscious, subconscious or unconscious—is the search for hidden patterns. In their first four years of life, children go through at least four major stages of discovery in identifying hidden patterns in their environment. During the first year, infants discover some three dozen speech sounds, known as phonemes (for example, ba, da, ca), and create standard categories for variants. They then discover unique patterns in strings of phonemes, known as words, at the rate of about six new words every day.

Between 18 and 36 months of age, toddlers discover structural relationships between words and syntax. The transition can often be very rapid, sometimes within one or two weeks. Finally, children discover Aristotle's rule about narratives—that they properly have a beginning, a middle, and an end—and start demanding appropriate endings to their bedtime stories. Of course, the search for meaning does not end there. As adults, we are constantly trying to make sense of our experiences and discern meaningful patterns in our actions, perceptions, and environment. Indeed, most of the tasks of consciousness are aimed at coping with novel situations, finding suitable patterns amid confusion, and creating new choices. Standard responses to ordinary situations do not require the attention of conscious thought.

Considering these factors, it appears that consciousness operates on different levels. If consciousness is defined as the highest current level of thought, it stands to reason that conscious thought operates at the level of objects and simple actions upon waking up in the morning. Forming relationships, like speaking in sentences, becomes possible only after a sufficient warm-up period or event, like morning coffee. Relations between relationships, like analogies, require even more time acclimation, like a double espresso. Given the ephemeral nature of consciousness, understanding how to improve the stability and duration of these levels is critical to building new ones. Such techniques could then be incorporated into our educational and training programs, enabling students to process information at higher levels of mental operation.

Using a Darwinian Approach to Understanding Brain Development

Given its well-deserved reputation for achieving quality on a timescale of species and antibodies, can the Darwinian process be used to improve the raw material of consciousness repeatedly beyond the incoherence of dreams? In principle, the problem is not whether it *can* be done—it can—it is whether coherence can be achieved quickly enough to be of use to our higher-level mental abilities, on the timescale of conversational replies.

The Darwinian process promotes coherence, or quality, through variation and selection over many generations. However, many aspects of selectionism are referred to as "Darwinian" when they are not truly Darwinian processes. For example, simple selective survival processes, such as leaf culling on trees, which results in a pattern, do not involve a Darwinian process. Neural connections in the brain also engage in the latter behavior by sprouting in abundant amounts, then culling into adult patterns. Take monkeys, for example: The axon count in

the corpus callosum drops 70 percent drop between birth and the age of six months.

Similarly, there are a lot of connections from all parts of the cerebral cortex, even the visual cortex, down to the spinal cord at an early stage of development. By the time an individual reaches adulthood, only the motor strip and the premotor area still have direct connections. All the others have withdrawn. This sort of sprouting and culling is a major feature of the development of a timescale at the individual level. But it is not the recursive bootstrapping of quality that we associate with the reputation of Darwinian processes.

One should also not confuse change with Darwinism. Leaf locations can be modified by rotating a potted plant; climates vary, but so do skirt lengths. Alterations in quality or complexity are not associated with these changes, and successes are neither achieved nor repeated to achieve more success. Darwinian adaptations can be pyramided to achieve new levels; these cannot.

A Darwinian process has six essential characteristics:

- A pattern exists (e.g., genes).
- The pattern is copied or cloned.
- Variant patterns arise because of copying errors and recombinations.
- Populations of some of these variants compete against other populations for area (e.g., bluegrass and crabgrass).
- A multifaceted environment makes some of these variants more common than others (i.e., natural selection).
- The more successful variants serve as the most frequent centers for further variation, and future generations spread out to nearby regions to repeat this process (that is, Darwin's Inheritance Principle).

All six characteristics are required to affect the recursive bootstrapping of quality. Having first five features, without the sixth, results in nothing more than population drift or random jumps from one barrier solution space to another.

So, can Darwinian processes be accelerated so that coherence, or quality, can be achieved quickly enough to be of use to our higher-level mental abilities? Once again, the answer is yes. Four known catalysts help speed the evolutionary process along:

- Systematic recombination: Variation is introduced systematically (for example, bacterial conjugation or sex).

- Fluctuating climates: Climate changes result in more-severe selection and more-frequent culling and, thus, in more-frequent opportunities for variation during expansion.
- Island biogeography, resource scarcity, geographical barriers and climate factors: These can separate a population into isolated subdivisions that discourage migration and promote inbreeding. Repeated separation and unification, or “pumping,” increase diversity and create variants capable of living further out on the habitat’s margins (e.g., the San Juan Islands, which are now surrounded by ocean but were hilltops connected by a broad valley during an ice age).
- Empty niches to fill: These can be due to the extinction of entire subpopulations; pioneers from other subpopulations will rediscover the vacated region and its replenished resources, and a population boom will result; an absence of competition for several generations results, giving rare variants that would otherwise perish a chance to survive. Once established, they may be able to survive future threats to their existence.

Attempts to duplicate the evolutionary process go well beyond the notions of connections and artificial neural networks. The best effort yet is Holland’s genetic algorithm that includes the six essential characteristics of the Darwinian process and one catalyst, systematic recombination. Even more promising, however, is the concept of a Darwinian machine that can incorporate these six characteristics and all four catalysts with stabilizing levels of organization. This concept will likely serve as the basis for intelligent machines of the future.

The Biological Basis for a “Darwin Machine”

While the brain contains the necessary circuitry for a Darwinian process, it has not yet been determined whether one actually occurs, much less how often or in which areas. Such a process would take place in a group of nerve cells in the cerebral cortex. As shown in Figure A.1, each neuron is a treelike structure that contains some 10,000 inputs, called synapses, and about the same number of outputs, called axons, that branch out to connect to other cells.

The pyramidal neurons are the excitatory neurons of cerebral cortex. Figure A.2 shows the axon of each of the three neurons spreading sideways in the superficial layer for a few millimeters in each direction. These cells are arranged in a pattern that is capable of a Darwinian process; those located in the deep layer pyramid do have such a pattern.

The circuitry that the brain uses to do it.




Figure A.1—The Brain's Circuitry

A Darwin Machine in cerebral cortex?

- Pyramidal neurons are the excitatory neurons of cerebral cortex.




Figure A.2—Three Neurons and Their Axons

The cerebral cortex can be thought of as a series of six layers, as shown in Figure A.3. The deep layers tend to act as a sort of outbox, sending axons down to the spinal cord, thalamus, or other destination. The middle layers tend to receive inputs from the thalamus. Finally, the upper layers tend to act like an internal or interoffice mailbox, sending their outputs to other parts of cortex, either immediately to the side or down through the white matter.

Unlike their deep counterparts, the superficial layers exhibit a patterning of their axons, as shown in Figure A.4. The axon tends to go about a 0.5 mm before any output occurs. Figure A.5 is a drawing of a neuron and a cluster of synapses from the superficial layers. These clusters are formed by the overlap of axons terminating near their immediate neighbors, creating a sort of annular ring that surrounds the area at which an input is received. In general, such clusters will occur every 0.5 mm for a distance of 3 or 4 mm in three dimensions, though a local metric may dictate slight variations (for example, 0.65 or 0.85 mm, depending upon the exact location in the cortex). This "express train" arrangement allows outputs to skip intermediate junctions in their path.

Although this may look like a very simple-minded pattern, it is exactly what is required for Darwinian processing. When cells talk to one another, they tend to synchronize with each other. This is true of most excitable systems, and there are




Figure A.3—The Layers of the Cerebral Cortex

In neocortex's superficial layers (I-III), the excitation is spread like an express train, skipping intermediate stops and only delivering excitation in patches about 0.5 mm apart.

But the spread is radial, so the excitation is in a ring, centered on the superficial pyramidal neuron (whose neighbors are also excited).




Figure A.4—Patterning in the Superficial Layers




Figure A.5—A Superficial Layer Neuron and Its Synapses

many examples of this process, known as *entrainment*, in nature. Pendulum clocks on the same shelf synchronize their tics after about 30 minutes or so, and linear relaxation oscillators do it even more quickly. Neural cells are able to synchronize their firings so efficiently because they receive input simultaneously as a result of the overlapping ringlike structures of axons that, taken together, form a triangular array. Cells in this recruitment arrangement tend to fire together.

Many of these triangular arrays exist in the cerebral cortex—one might be sensitive to the color of an object, another might be sensitive to its shape. To avoid redundancy, the largest possible number of arrays is probably limited to a few hundred. Together, these arrays form hexagons in the “mosaic.” The pattern suggests that the spatiotemporal firing within each hexagon is a complete descriptive set, akin to a little musical pattern lasting only a couple of hundred milliseconds.

However, the spatiotemporal firing patterns in each descriptive set are not the only representation of a thought or action in the brain. Synaptic connectivity, the weightings that help maintain these firing patterns, enables the brain to remember these patterns. For example, the spinal cord has the ability to produce a number of different spatiotemporal patterns, called the “gates of locomotion.” This term is used to describe the manner and order in which a leg’s muscles fire to bring it forward. While each type of movement (e.g., walking, trotting, running) has its own spatiotemporal pattern, the same cells are used to do them all—it is just a matter of the initial conditions.

These two levels of representation—a short-term spatiotemporal pattern that is needed to effect thought or action and a long-term spatial pattern that is needed to store it. A good analogy might be a phonograph record whose spatial-only pattern of grooves is able to recreate the spatiotemporal pattern of music and speech. A consequence of this arrangement is that the triangular arrays do not always fire patterns, but rather compete for territory in the cortex. For example, some “undecided” areas (arrays) of the mosaic may receive input from two or more surrounding areas, each with a different firing pattern—say apples and bananas, as in Figure A.6. If the undecided area resonates better (due to a memory imprint) with apples than bananas, it will likely begin firing “apples.” Thus, success in cloning is subject to the extant memories of an environment. This competition provides the mechanism by which decisions are made and ambiguity is resolved. “Winning” spatial temporal patterns are responsible for any motor function output.

With this mechanism in mind, the essential Darwinian characteristics can now be evaluated in the context of the brain:




Figure A.6—Undecided Cells Receiving Input from Two Areas, in this Case, Representing Apples and Bananas

- A 0.5 mm hexagonal spatiotemporal pattern exists.
- The pattern is copied by recurrent excitation and entrainment of triangular arrays.
- Variant patterns arise when these triangular arrays escape conforming neighbors.
- Populations of these variants (hexagonal mosaics) compete for space.
- A multifaceted environment of sensory inputs and memorized resonances makes some variants more common.
- The more successful variants are the most frequent centers for further variation.

The evidence suggests that the neocortex could, in fact, be a Darwinian machine. It has all six essentials and all four optional catalysts, and it produces in advance the spatiotemporal firing patterns needed for converting thought into action. Success and quality are biased by real-time sensory inputs, the environment, and the memorized features of previous environments resulting from synaptic connectivity.

However, there are a lot of ways in which this mechanism can operate. While the brain circuit outlined above is fairly common and appears to be capable of performing all of these actions, it is not known how much time the different parts of the brain spend engaged in this activity. Is it used only during development to lay down the cortical structure or all the time in all areas of the cortex? It is likely somewhere in between. What is known is that these express-train connections exist in all of the common varieties of lab animals with one exception: the rat.

Another thing that is unresolved is whether the brain engages in anything fancier than the Darwinian process. While a more sophisticated process would appear to be unnecessary to perform these kinds of activities, it cannot be determined whether the brain uses this circuit in the exact manner described above until more precise observations from the neocortex are recorded. This technology is around the corner but is not yet available.

“Enhancing” Brain Function

Some glimpses of how we might improve this process emerge from this Darwinian view of how the brain could operate on this timescale.

Understanding how a system works often makes it easier to improve its performance. For example, some higher aspects of intelligence—speed of

learning, speed of operation, number of thoughts that can be held simultaneously in the mind—may be able to be addressed, toward the goal of eventually stabilizing higher, more-abstract levels of consciousness.

The concept of cerebral circuitry and Darwinian processing was first envisioned in the 1930s, but it is the technology of the 1990s that makes it possible to duplicate. Using a high-speed, hybrid computer, it appears feasible to emulate—not reverse engineer—this process. While it can be accomplished as a straight digital simulation, a more-natural fit involves the use of a hybrid digital-analog output device that would digitally copy spatiotemporal codes but record in analog spatial-only resonances. These resonances are what gives the circuitry its interesting properties. In the future, researchers will invent a number of circuits to undertake this task of bootstrapping quality through a series of generations. At that point, the notion of a Darwinian machine that thinks like a human will become a reality.

Finally, there is the issue of how to speed things up. A learning process that takes days to produce results is of little value to most brain functions. In the absence of further information, it is impossible to detail this process. However, some lessons about how to speed up evolutionary processes can be drawn from other Darwinian mechanisms. A primary factor appears to be “windows of opportunity” in behavior—what the French call *avoir l'esprit de l'escalier* [to have the spirit of the staircase]: thinking of the right reply on the stairway after leaving the party. From the perspective of the brain, the timescale associated with an evolutionary process must be a few seconds or less.

To illustrate the catalytic factors at work in this process, consider what happened to our ancestors over the past few million years. About 20 years ago, it was discovered that brain size started increasing some 2.5 million years ago, just around the same time that stone toolmaking became prominent. It turns out that it was during this period that the australopithecine branched off into the *homo* genus, as shown in Figure A.7. The question thus arises as to what was happening back then that could have caused all this to occur and continue? The most likely answer is the onset of the ice ages.

The Role of Climate in Developing the “Big Brain”

During the last 15 years, researchers have concluded that the ice ages were characterized by abrupt climate changes on a number of different timescales, as shown in Figure A.8. What is very obvious are the temperature fluctuations throughout the ages. For example, some 15,000 years ago, at a time when ice sheets covered the northern hemisphere, the temperature abruptly rose to almost

It's been 5 myr
since we shared a
common ancestor
with chimps,
bonobos.

Both brain increase
and toolmaking
started 2.5 myr -- as
did the ice ages.

But how did
climate pump up
brains?




Figure A.7—The Emergence of the Genus *Homo*




Figure A.8—Abrupt Climate Changes During the Last Ice Age

modern levels, despite all the ice. Some 2,000 years later, the temperature abruptly cooled and then warmed back up just as suddenly. One can see the same thing happening even further back on a very compressed timescale. Approximately 122,000 years ago, an abrupt cooling occurred in the middle of the last warm period. It lasted long enough for the sea level to drop some 4 m before warming back up to raise the sea level by about 6 m. About 8,000 years ago, there was also a brief period of moderate cooling.

The explanation for this climatic behavior is as complicated as it is lengthy. What is apparent, however, is that the climate has two stable states—a warm state and a cold state—and it flips between them based on the nature of the ocean currents, as shown in Figure A.9. The consequences of this transition are extreme. It is equivalent to jacking up (or ratcheting down) the entire landscape into a new climatic zone. Contrary to popular myth, however, it is not the magnitude of a cooling that threatens hominids, but its velocity, once the magnitude becomes large enough to effect the mix of plants and prey. That is, the process happens so quickly—within a human generation—that there is not enough time for biological adaptations to take place. The timescale is critical to survival.




Figure A.9—Two Stable States, Cool and Warm

If, for example, the global temperature were to ramp down over a period of 500 years, life would be able to adapt to its new environment. A gradual change in vegetation would occur, emphasizing colder-weather species, like those normally found in higher altitudes, and hominids would likely learn to cope with new challenges. A stepwise cooling over a period of just 10 to 20 years, however, would pose a real threat to the survival of many species, including humans. Reduced rainfall would cause forests to dry up and burn off, leaving grass as the major food resource for at least a couple hundred years. To survive, then, an animal must either be able to eat grass or eat an animal that eats grass until plant secessions allow the ecosystem to advance past this monoculture to a forest more suited for these temperatures. The historical record indicates that our ancestors were subjected to many of these transitions over a period of thousands of years.

Conclusion: The Timescale of Thought and Action

Now transpose these lessons to the timescale of thought and action. The significance of timescale in the evolutionary process indicates that periods of monoculture are important to the neocortex. It stands to reason then that narrowly focused activities, such as concentrating, meditating or sleeping, will

likely result in a downsizing and fragmenting of hexagonal mosaics into regional populations as cortical excitability fluctuates. Because these "climate" fluctuations occur rapidly, they "pump" the other three Darwinian catalysts—systematic recombination, island biogeography, and empty niches for new populations. Thus, it appears that one can control the speed at which change occurs by affecting the noise level in the neocortex.

Since climatic change occurs on various timescales, ranging from the millennia of ice ages to abrupt phenomena, such as el Niño. So, too, would we expect the brain's cortical "climate" to operate on various timescales (if an electroencephalogram is any acceptable measure of the brain's excitability). This process would involve repeatedly reducing and expanding to select the types of cells that are most capable of surviving bottleneck conditions. In addition to these quantum fluctuations, the neocortex is engaged in many parallel processes involving lots of territory, enabling it to maintain independent branches in a "playoff" system of alternatives. To further complicate the matter, different hexagonal arrays represent the different levels of organization and consciousness. As a result, a slow Darwinian process, such as forming a mental agenda, could bias a faster Darwinian process, such as thought and action, thereby skewing the results.

These characteristics are what one can expect from a forthcoming wave of Darwinian technologies. Of course, ethical questions must be considered. If a Darwinian circuit can be replicated in the artificial intelligence of a machine, it ought to be able to do what the Darwinian process is famous for elsewhere: shape up quality. While such a machine would have novel processing and problem-solving capabilities, it would not necessarily be considered "conscious." But as enhancements are made and versatility increases, society will face some very serious issues. For example, what if these technologies are able to work faster than humans? Is it possible to reach a point where all but the most intelligent people can be replaced by these devices? Then what happens when Moore's Law, some 18 months later, makes even those persons obsolete? Theoretically, there is no upper limit on processing speed if enough resources are available. These are the implications and dangers associated with building intelligence into machines.

B. Quantum Theory and Human Consciousness

Leaders: Stuart Hameroff, University of Arizona¹
 Roger Penrose, University of Oxford
 October 22, 1998

On October 22, 1998, the third in a series of study group meetings was held at RAND. The series, sponsored by the Defense Advanced Research Projects Agency (DARPA), focused on social and political governance questions arising from the impacts of the information and biological revolutions. This paper presents a summary of ideas and thoughts presented at the study group meeting.

Questions about consciousness and intelligence arise in the context of both the information and biological revolutions. To understand these phenomena, the session continued the study group's inquiry into the human brain. The goal was to understand the physical limitations of the brain, how it evolved to its current capabilities, and the origins of conscious thought and intelligence. This information is needed for putting any technological enhancements to the human brain in context that might be researched and attempted. In addition, this information is key to understanding whether computers or networks of computers can ever achieve a kind of conscious intelligence.

Presentation by Stuart Hameroff and Roger Penrose

The study group leaders invited Stuart Hameroff and Roger Penrose to present their research on the brain and human consciousness, and discuss the role of the computer in enhancing concepts of human intelligence. Professor Hameroff is a professor in the Departments of Anesthesiology and Psychology at the University of Arizona, and a physician on staff at the University Hospital. Professor Penrose is Rouse Ball Professor of Mathematics at the University of Oxford. He is the recipient of a number of awards, including the 1988 Wolf Prize (which he shared with Stephen Hawking for their research into the understanding of the universe), the Dannie Heinemann Prize, the Royal Society Royal Medal, and the Albert Einstein Prize. His 1989 book, *The Emperor's New*

¹Additional information, as well as more detailed citations, is available on Dr. Hameroff's Web site: <http://www.u.arizona.edu/~hameroff/>.

Mind, was a best-seller and won the 1990 Rhone-Poulenc Science Book Prize. His latest works include *Shadows of the Mind* (1994), *The Nature of Space and Time* (1996) (with Stephen Hawking), and *The Large, the Small and the Human Mind* (1997).

Together, Professors Hameroff and Penrose have developed a theory of consciousness. They propose that quantum theory and a newly proposed physical phenomenon, quantum wave function, are essential for consciousness and occur in cytoskeletal microtubules and other structures within the brain's neurons. Several papers on this theory can be found at Dr. Hameroff's Web site: <http://www.u.arizona.edu/~hameroff/>.

The Problem of Consciousness

Conventional explanations portray consciousness as an emergent property of classical computerlike activities in the brain's neural networks. While there is some disagreement as to the particular point of origin, the prevailing view among scientists in this camp is that (1) patterns of neural activity correlate with mental states; (2) synchronous network oscillations of neuronal circuits in the thalamus and cerebral cortex temporarily binds information; and (3) consciousness emerges as a novel property of computational complexity among neurons.

However, these approaches appear to fall short in fully explaining certain enigmatic features of consciousness, such as

- the nature of subjective experience, or "qualia"—our "inner life" (Chalmers' "hard problem," 1996)
- the binding of spatially distributed brain activities into unitary objects in vision, and a coherent sense of self, or "oneness"
- the transition from preconscious processes to consciousness itself
- noncomputability, or the notion that consciousness involves a factor that is neither random nor algorithmic, and that consciousness cannot be simulated (Penrose, 1989, 1994, 1997)
- free will
- subjective time flow.

Brain imaging technologies have demonstrate the anatomical locations of activities that appear to correlate with consciousness but may not be directly responsible for consciousness.

How do neural firings lead to thought and feelings? The conventionalist (also called functionalist, reductionist, materialist, physicalist, and computationalist) approach argues that neurons and their chemical synapses are the fundamental units of information in the brain and that conscious experience emerges when a critical level of complexity is reached in the brain's neural networks. The basic idea is that the mind is a computer functioning in the brain (brain = mind = computer).

However, in fitting the brain to a computational view, such explanations omit incompatible neurophysiological details:

- widespread apparent randomness at all levels of neural processes (is it noise or underlying levels of complexity?)
- glial cells (which accounts for some 80 percent of the brain)
- dendritic-dendritic processing
- electrotonic gap junctions
- cytoplasmic/cytoskeletal activities
- living state (the brain is alive!).

A further difficulty is the absence of testable hypotheses in emergence theory. No threshold or rationale is specified; rather, consciousness "just happens."

Finally, the complexity of individual neurons and synapses is not accounted for in such arguments. Since many forms of motile single-celled organisms lacking neurons or synapses are able to swim, find food, learn, and multiply through the use of their internal cytoskeleton, can they be considered more advanced than neurons? Are neurons merely simple switches, or are they something more?

Microtubules

Activities within cells ranging from single-celled organisms to the brain's neurons are organized by a dynamic scaffolding called the cytoskeleton. A major component of the cytoskeleton is the *microtubule*, a hollow, crystalline cylinder 25 nm in diameter. Microtubules are, in turn, composed of hexagonal lattices of proteins, known as *tubulin*.

Microtubules are essential to cell shape, function, movement, and division (Figure B.1). In neurons, microtubules self-assemble to extend axons and dendrites and to form synaptic connections, then help to maintain and regulate synaptic activity responsible for learning and cognitive functions (Figure B.2).




Figure B.1—Crystallographic Structure of Microtubules




Figure B.2—Schematic View of Two Neurons Connected by Chemical Synapse. Axon terminal (above) releases neurotransmitter vesicles, which bind receptors on postsynaptic dendritic spine. Cytoskeletal structures microtubules ("MTs"—thicker tubes) are visible within the neurons, as well as actin, synapsin, and others that connect MTs to membranes.

While microtubules have traditionally been considered to be purely structural elements, recent evidence has revealed that mechanical, chemical, and electrical signaling and communication functions also exist, the result of microtubule interaction with membrane structures by linking proteins, ions and "second-messenger" signals, and voltage fields, respectively.

Current models propose that tubulins within microtubules undergo coherent excitation, switching between two or more conformational states in nanoseconds. Dipole couplings among neighboring tubulins in the microtubule lattice form

dynamic patterns, or “automata,” which evolve, interact, and lead to the emergence of new patterns. Research indicates that microtubule automata computation could support classical information processing, transmission, and learning within neurons.

Microtubule automaton switching offers a potentially vast increase in the computational capacity of the brain. While conventional approaches focus on synaptic switching at the neural level, which optimally yields about 10^{18} operations per second in human brains ($\sim 10^{11}$ neurons per brain, with $\sim 10^4$ synapses per neuron, switching at $\sim 10^3 \text{ sec}^{-1}$), microtubule automata switching can explain some 10^{27} operations per second ($\sim 10^{11}$ neurons with $\sim 10^7$ tubulins per neuron, switching at $\sim 10^9 \text{ sec}^{-1}$). Indeed, the fact that all biological cells typically contain approximately 10^7 tubulins could account for the adaptive behaviors of single-celled organisms, which have no nervous system or synapses. Rather than simple switches, then, it seems that neurons are actually complex computers.

Theories of Consciousness: Panexperiential Philosophy Meets Modern Physics

Still, greater computational complexity and ultrareductionism to the level of microtubule automata cannot address the enigmatic features of consciousness—in particular, the nature of the conscious experience. Something more is required. If functional approaches and emergence are incomplete, perhaps the raw components of mental processes (or “qualia”) are fundamental properties of nature (like mass, spin, or charge). This view has long been held by panpsychists throughout the ages.

For example, Buddhists and Eastern philosophers claimed a “universal mind.” Following the ancient Greeks, Spinoza argued in the 17th century that some form of consciousness existed in everything physical. The 19th century mathematician Leibniz proposed that the universe was composed of an infinite number of fundamental units, or “monads,” with each possessing a form of primitive psychological being. In the 20th century, Russell claimed that there was a common entity underlying both mental and physical processes, while Wheeler and Chalmers have maintained that there is an experiential aspect to fundamental information.

Of particular interest is the work of the 20th century philosopher Alfred North Whitehead, whose panexperiential view remains most consistent with modern physics. Whitehead argued that consciousness is a process of events occurring in a wide, basic field of protoconscious experience. These events, or “occasions of

experience," may be comparable to quantum state reductions, or actual events in physical reality (Shimony, 1993). This suggests that consciousness may involve quantum state reductions (a form of quantum computation). But in what medium do such "occasions" occur?

Whether protoconscious experience, or qualia, could exist in the empty space of the universe depends upon how space is defined. Historically, space has been described as either an absolute void or a pattern of fundamental geometry. Democritus and the Michaelson-Morley results argued for "nothingness," while Aristotle ("plenum") and Maxwell ("ether") rejected the notion of emptiness in favor of "something"—a background pattern. Einstein weighed in on both sides of this debate, initially supporting the concept of a void with his theory of special relativity but then reversing himself in his theory of general relativity and its curved space and geometric distortions—the space-time metric. Could protoconscious qualia be properties of this fundamental metric?

Quantum Computing and Consciousness

At extremely small scales, space-time is not smooth, but quantized. Quantum electrodynamics and quantum field theory predict virtual particle-waves (or photons) that pop into and out of existence, creating quantum "foam" in their wake. Lamoreaux verified presence of virtual photons in space-time in 1997. In 1971, Roger Penrose modeled this granularity as a dynamic web of quantum spins. These "spin networks" create an array of geometric volumes and configurations at the Planck scale (10^{-33} cm, 10^{-43} secs), which dynamically evolve and define space-time geometry. If spin networks are the fundamental level of space-time geometry, they could provide the basis for protoconscious experience. Thus, particular configurations of quantum spin geometry would convey particular types of qualia, meaning and aesthetic values. A process at the Planck scale (e.g., quantum scale reductions) could then access and select configurations of experience.

If protoconscious information is embedded at the near-infinitesimal Planck scale, how could it be linked to biology? Penrose's answer is to extend Einstein's theory of general relativity (in which mass equates to curvature in space-time) down to the Planck scale. Specific arrangements of mass are, in reality, then specific configurations of space-time geometry. Events at the very small scale, however, are subject to the seemingly bizarre goings-on of quantum theory. A century of experimental observation of quantum systems has shown that, at least at small scales, particles (mass) can exist in two or more states or locations simultaneously. Penrose views this phenomenon of quantum superposition as

simultaneous space-time curvature in opposite directions—a separation or bubble in underlying reality.

Superposition and subsequent reduction, or collapse, to single, classical states may have profoundly important applications in technology, as well as toward the understanding of consciousness. In the 1980s, Benioff, Feynman, Deutsch, and other physicists proposed that states in a quantum system could interact (via entanglement) and enact computation while in quantum superposition of all possible states (i.e., “quantum computing”). While classical computing processes bits (or conformational states) as 1 or 0, quantum computations involve the processing of superpositioned “qubits” of both 1 and 0 (and other states) simultaneously.

Quantum theory also predicts that two or more particles, if once together, will remain somehow connected, even when separated by great distances. This “entanglement” enables quantum computing to achieve a nearly infinite parallel computational ability. Thus, quantum computers, if they can be constructed, will be able to solve important problems (e.g., factoring large numbers) with efficiency unattainable with classical designs (Shor, 1994).

Results, or solutions in quantum computing, are obtained when, after a period of quantum superposition, the qubits “collapse,” or reduce to classical bit states. As quantum superposition may only occur in isolation from the environment, reduction may be induced by breaching isolation. But what about quantum superpositions that remain isolated—for example, Schrödinger’s mythical cat, which is both dead and alive? This is the famous problem of wave function collapse, or quantum state reduction.

Roger Penrose’s Objective Reduction (OR)

So how or why do quantum superimposed states that avoid environmental interactions become classical and definite in the macroworld? Many physicists now believe that some objective factor disturbs the superposition and causes it to collapse. Penrose proposes that this factor is an intrinsic feature of space-time itself: quantum gravity. According to Penrose’s interpretation of general relativity, quantum superposition—a separation in mass from itself—is equivalent to separation in underlying space-time geometry, or simultaneous space-time curvatures in opposite directions.

Penrose argues that these separations in fundamental reality, or “bubbles,” are unstable—even when isolated from the environment—and will reduce spontaneously and noncomputably to a specific state at a critical threshold of

separation or “decoherence,” thereby avoiding the need for “multiple worlds.” This objective threshold is defined by the indeterminacy principle:

$$E = h/T$$

where E is the gravitational self-energy of the superposed mass separated from itself; h is Planck’s constant divided by 2π , and T is the coherence time until collapse occurs. Thus, the size and energy of a system in superposition, or the degree of space-time separation, are inversely related to the time T until reduction. (E can be calculated from the superposed mass m and the separation distance a . See Hameroff and Penrose, 1996a.)

Assuming isolation, the following masses in superposition would collapse at the designated times, according to Penrose’s objective reduction:

Mass (m)	Time (T)
Nucleon	10^7 years
Beryllium ion	10^6 years
Water speck	
• 10^{-5} cm radius	Hours
• 10^{-4} cm radius	1/20 second
• 10^{-3} cm radius	10^{-3} seconds
Schrödinger’s cat ($m = 1$ kg, $a = 10$ cm)	10^{-37} seconds

If quantum computation with objective reduction occurs in the brain, enigmatic features of consciousness could be explained:

- By occurring as a self-organizing process in what is suggested to be a panexperiential medium of fundamental space-time geometry, objective reductions could account for the nature of subjective experience by accessing and selecting protoconscious qualia.
- By virtue of involvement of unitary, entangled quantum states during preconscious quantum computation and the unity of quantum information selected in each objective reduction, the issue of binding may be resolved.
- Regarding the transitions from preconscious processes to consciousness itself, the preconscious processes may equate to the quantum superposition-computation phase, and consciousness to the actual, instantaneous objective reduction events. Consciousness may then be seen as a sequence of discrete events (e.g., at 40 Hz).

- As Penrose objective reductions are proposed to be noncomputable (reflecting influences from space-time geometry that are neither random nor algorithmic), conscious choices and understanding may be similarly noncomputable.
- Free will may be seen as a combination of deterministic preconscious processes acted on by a noncomputable influence.
- Subjective time flow derives from a sequence of irreversible quantum state reductions.

In what types of brain structures might quantum computation with objective reduction occur? If these events occur in the brain, they would be expected to coincide with known neurophysiological processes with appropriate time scales. For consciousness, then, T should be in range of tens to hundreds of milliseconds.

Event	T (ms)
Buddhist "moment of awareness"	13
"Coherent 40 Hz" oscillations	25
Electroencephalogram alpha rhythm (8 to 12 Hz)	100
Libet's sensory threshold (1979)	500

Objective reduction events in this time frame would require a mass on a nanogram scale. So, what is m ?

Are Proteins Qubits?

Biological life is organized by proteins. By changing their conformational shape, proteins are able to perform a wide variety of functions, including muscle movement, molecular binding, enzyme catalysis, metabolism, and movement. Dynamical protein structure results from a "delicate balance among powerful countervailing forces" (Voet and Voet, 1995). The types of forces acting on proteins include charged interactions (such as covalent, ionic, electrostatic, and hydrogen bonds), hydrophobic interactions, and dipole interactions. The latter group, also known as van der Waals forces, encompasses three types of interactions:

- permanent dipole–permanent dipole
- permanent dipole–induced dipole
- induced dipole–induced dipole.

As charged interactions cancel out, hydrophobic and dipole–dipole forces are left to regulate protein structure. While induced dipole–induced dipole interactions, or London dispersion forces, are the weakest of the forces outlined above, they are also the most numerous and influential. Indeed, they may be critical to protein function. For example, anesthetics are able to bind in hydrophobic “pockets” of neural proteins and ablate consciousness by virtue of these London forces. London force attraction between any two atoms is usually less than a few kilojoules; however, since thousands occur in each protein, they add up to thousands of kilojoules per mole, and cause changes in conformational structure.

If proteins are qubits, assemblies of proteins in some type of organelle or biomolecular structure could act as a quantum computer. So which biological structures are best suited for objective reduction? Ideal structures would

- be abundant
- be capable of information processing and computation
- be functionally important (e.g., regulating synapses)
- be self-organizing
- be tunable by input information (e.g., microtubule-associated protein orchestration)
- be periodic and crystal-like in structure (e.g., dipole lattice)
- be isolated (transiently) from environmental decoherence
- be conformationally coupled to quantum events (e.g., London forces)
- be cylindrical waveguide structures
- have a plasmalike charge-layer coating.

While various structures or organelles have been suggested (e.g., membrane proteins, clathrins, myelin, presynaptic grids, and calcium ions), the most logical candidates are microtubule automata.

The Penrose-Hameroff Orchestrated Objective Reduction Model

The Penrose-Hameroff Orchestrated Objective Reduction (Orch OR) model proposes that quantum superposition–computation occurs in microtubule automata within brain neurons and glia. Tubulin subunits within microtubules act as qubits, switching between states on a nanosecond (10^{-9} sec) scale, governed by London forces in hydrophobic pockets. These oscillations are “tuned” and “orchestrated” by microtubule-associated proteins (MAPs),

providing a feedback loop between the biological system and the quantum state. These qubits interact computationally by nonlocal quantum entanglement, according to the Schrödinger equation, with preconscious processing continuing until the threshold for objective reduction (OR) is reached ($E = h/T$). At that instant, collapse occurs, triggering a “moment of awareness,” or a conscious event—an event that determines particular configurations of Planck-scale experiential geometry and corresponding classical states of microtubule automata that regulate synaptic and other neural functions. A sequence of such events could provide a forward flow of subjective time and “stream” of consciousness. Quantum states in microtubules may link to those in microtubules in other neurons and glia by tunneling through gap junctions, permitting extension of the quantum state throughout significant volumes of the brain.

From $E = h/T$, the size and extension of Orch OR events that correlate with subjective or neurophysiological descriptions of conscious events can be calculated:

Event	T (ms)	E
Buddhist “moment of awareness”	13	4×10^{15} nucleons (4×10^{10} tubulins/cell ~ 40,000 neurons)
“Coherent 40 Hz” oscillations	25	2×10^{15} nucleons (2×10^{10} tubulins/cell ~20,000 neurons)
EEG alpha rhythm (8 to 12 Hz)	100	5×10^{14} nucleons (5×10^9 tubulins/cell ~5,000 neurons)
Libet’s sensory threshold (1979)	500	10^{14} nucleons (10^9 tubulins/cell ~1,000 neurons)

But how could delicate quantum superposition–computation be isolated from environmental decoherence in the brain (generally considered to be a noisy thermal bath), while also communicating with the environment? One possibility is that quantum superposition–computation occurs in an isolation phase that alternates with a communicative phase (Figures B.3 through B.5). One of the most primitive biological functions is the transition of cytoplasm between a liquid, solution (“sol”), phase and a solid, gelatinous (“gel”), phase due to assembly and disassembly of the cytoskeletal protein actin. Actin sol-gel

transitions can occur at 40 Hz or faster and are known to be involved in neuronal synaptic release mechanisms.

Mechanisms for enabling microtubule quantum computation and avoiding decoherence long enough to reach the OR threshold may include

- sol-gel transitions
- plasma phase sleeves (Sackett)
- quantum excitations, ordering of surrounding water (Jibu/Yasue/Hagan)
- hydrophobic pockets
- hollow microtubule cores
- laserlike pumping, including environment (Frohlich/Conrad)
- quantum error correcting codes.




Figure B.3—Schematic Graph of Proposed Preconscious Quantum Superposition (number of tubulins) Emerging Versus Time in Microtubules. Area under curve connects superposed mass energy E with collapse time T in accordance with $E=h/T$. E may be expressed as nt , the number of tubulins whose mass separation for time T will self-collapse. For $T = 25$ msec (e.g. 40 Hz oscillations), $nt = 2 \times 10^{10}$ tubulins.




Figure B.4—Schematic of Quantum Computation. Three tubulins begin (left) in initial classical states, then enter isolated quantum superposition in which all possible states coexist. After reduction, one particular classical outcome state is chosen (right).

Orchestrated Objective Reduction, Cognition, and Free Will

Quantum computation with objective reduction may be associated with cognitive activities. While classical neural-level computation can provide a partial explanation, the Orch OR model allows for far greater information capacity and addresses issues of conscious experience, binding, and noncomputability consistent with free will. Such functions as face recognition and volitional choice may require a series of conscious events arriving at intermediate solutions. Preconscious processing of information occurs in the form of qubits, or superposed states of microtubule automata. As the threshold for objective reduction is reached, these qubits collapse to definite states and become bits, resulting in a conscious experience of recognition or choice.

The problem in understanding free will is that human actions seem neither totally deterministic nor random. In Orch OR, reduction outcomes involve a factor that is “noncomputable.” The microtubule quantum superposition evolves




Figure B.5—Quantum Superposition Entanglement in Microtubules for Five States Related to Consciousness. A. Normal 30 Hz experience. B. Anesthesia: anesthetics bind in hydrophobic pockets and prevent quantum delocalizability and coherent superposition. C. Heightened Experience: increased sensory experience input increases rate of emergence of quantum superposition. Orch OR threshold is reached faster and Orch OR frequency increases. D. Altered State: even greater rate of emergence of quantum superposition due to sensory input and other factors promoting quantum state (e.g. meditation, psychedelic drug). Predisposition to quantum state results in baseline shift and collapse so that conscious experience merges with normally sub-conscious quantum computing mode. E. Dreaming: prolonged sub-threshold quantum superposition time.

linearly (analogous to a quantum computer) but is influenced at the instant of collapse by hidden nonlocal variables (quantum-mathematical logic inherent in fundamental space-time geometry). The possible outcomes are limited (or probabilities are set) by neurobiological feedback (MAPs). The precise outcome (our "chosen" action) is determined by effects of the hidden logic on the quantum system poised at the edge of objective reduction. This could explain why people generally do things in an orderly, deterministic fashion, but occasionally their actions or thoughts are surprising, even to themselves.

Consciousness and Evolution

When in the course of evolution did consciousness first appear? Are all living organisms conscious, or did consciousness emerge more recently (e.g., with language or toolmaking)? The Orch OR model (unlike other models of consciousness) is able to make a prediction as to the onset of consciousness. Using $E = h/T$, the feasibility of consciousness for different organisms can be explored.

A single-celled organism (e.g., a paramecium, with $E = 10^7$ tubulins and $T = 50,000$ ms) would be unlikely to achieve consciousness, whereas a nematode worm (e.g., *C. elegans* with $E = 3 \times 10^9$ tubulins per cell and $T = 133$ ms) might possess the biological complexity to understand what it is like to be a worm. Is it mere coincidence that these organisms were prevalent at the Cambrian “explosion,” a burst of evolution 540 million years ago. Did these creatures possess the first consciousness? Did primitive consciousness (via Orch OR) accelerate evolution?

Would consciousness be advantageous to survival—above and beyond intelligent, complex behavior? The answer appears to be “yes.” Noncomputable behavior (i.e., unpredictability, intuitive actions) is likely to be beneficial in predator-prey relations. The conscious experience of taste may promote the search for food; the experience of pain may promote the avoidance of predators; and the pleasurable qualia of sex may promote reproduction. So, what is it like to be a worm? Absent a sensory apparatus, associative memory, and a complex nervous system, such a primitive consciousness would be a mere glimmer, a disjointed smudge of reality. But qualitatively, at a basic level, it would be akin to ours.

What about future evolution? Will consciousness occur in computers? The advent of quantum computers opens the possibility. However, as presently envisioned, quantum computers will have insufficient mass in superposition (e.g., electrons) to reach the threshold for objective reduction due to environmental decoherence. Still, future generations of quantum computers may be able to realize this goal.

Conclusions

Brain processes relevant to consciousness extend downward within neurons to the level of the cytoskeleton. An explanation of conscious experience requires (in addition to neuroscience and psychology) a modern form of panprotopsychism in which protoconscious qualia are embedded in the basic level of reality, as

defined by modern physics. The Penrose model of objective reduction connects brain structures to fundamental reality, leading to the Penrose-Hameroff model of quantum computation with objective reduction in microtubules. The Orch OR model is consistent with known neurophysiological processes, generates testable predictions, and is the type of fundamental, multilevel, and interdisciplinary theory that may account for the mind's enigmatic features.

C. Morals, Demonic Males and Evolutionary Psychology

Leaders: Robert Wright,¹

Author

September 10, 1998

Richard Wrangham,

Professor of Zoology, Harvard University

November 12, 1998

This appendix summarizes two different, but related, presentations: that of Robert Wright on the Intellectual Foundations for Sociobiology and Its Implications for the Social Sciences, and that of Richard Wrangham, on Human Social Nature, Aggression, and Social Order.

Introduction

For most of the latter half of this century, mentioning the name "Darwin" and the word "policy" in the same breath and in any positive terms implied "Social Darwinism" and was enough to invoke supreme discomfort in most audiences. Since Charles Darwin's *On the Origin of Species* (1859) was first published, many, particularly conservative, readers have attempted to use his ideas to justify existing social hierarchies and even slavery based on a supposedly Darwinian understanding of the natural order of life and "survival of the fittest." Social Darwinians argued that social outcomes and inequality were evolutionarily determined and therefore natural and immutable.

There were also those who reacted strongly to any such notion that evolution had shaped our society or behaviors. The debate on human nature became one of "nature versus nurture." Based both on influential research to the contrary

¹Robert Wright is a science journalist who writes regular columns for *Time* and *New Republic*. His 1994 book, *The Moral Animal*, has won high praise as an overview of the field of evolutionary psychology. While we normally would invite researchers rather than journalists as speakers in this series, we thought that Wright has an especially good grasp of the field as a whole and of its recent intellectual history. Unlike other areas in biology, the epistemological statuses of findings in evolutionary psychology are particularly controversial, with such critics as Stephen Jay Gould contesting many of its central findings. We asked Wright to speak about that debate as an introduction to the broader issue of sociobiology and its implications for the social sciences and public policy.

that showed a significant role for environmental and cultural conditioning and a reaction against the horrific outcomes of social Darwinian policies, such as in Germany under the National Socialists, popular and academic opinion over this century has come to reject nature in favor of nurture. This "cultural determinism" has itself become ingrained in our culture. Our policies are now predicated largely on the assumption that cultural and environmental factors alone shape human behavior.

Yet, even as social Darwinism continues to be rejected as a basis for policy, scientific research in a number of fields, not only biology and genetics but also anthropology and sociology, is revealing an ever greater role for genes and natural selection in shaping the human condition and human behavior. It is now generally accepted that, along with previously identified environmental and behavioral factors, there are genetic factors that influence the occurrence in different individuals or populations of diseases, such as cancer, or even various forms of addiction. On a parallel track, continuing observations of other primates and refinements in the logic of sexual and kin selection have advanced and changed our understanding of how the history of human evolution and the behavior of related species provide clues to our own behavior. This new attempt to understand the influence of evolution in determining human behavior has been called "evolutionary psychology"—a name that serves largely to differentiate this new approach from the earlier theories of social Darwinism.

While this approach appears to provide greater insight into human behavior and is already being applied to the policy arena,² the ideas of evolutionary psychology, if applied without caution or too simplistically, will also be vulnerable to the easy ideological and scientific pitfalls that have befallen other such attempts at social Darwinism. To guide the study group through these politically and scientifically dangerous shoals and to further explore the policy implications of evolutionary psychology, RAND invited Richard Wrangham, a professor of biological anthropology at Harvard University, and Robert Wright, a science and technology writer associated with *New Republic* magazine, to discuss their research in evolutionary psychology with the group. Both have written extensively on this subject.

The first presentation, by Robert Wright on September 10, 1998, treated the intellectual origins of evolutionary psychology and, in general terms, discussed the significance of the field for the policy process. The second presentation, by Richard Wrangham on November 12, 1998, used evolutionary psychology to

²See, for example, Fukuyama, (1998).

discuss the origins of male violence and provided some highly speculative implications such an evolutionary perspective may have for understanding human military behavior.

Nature vs. Nurture?: Evolutionary Psychology and Policy

Both presentations opened with a discussion of the epistemology of the field, to differentiate their research from previous attempts at social Darwinism. Both authors focused on the well-known dichotomy of nature versus nurture.

This dichotomy is now referred to as Galton's error. Francis Galton was a cousin of Darwin's who, upon reading *On the Origin of Species*, felt he had found the key to human nature within the book's pages. To Galton, Darwin's work supported the view that human nature and society had been shaped through the process of natural selection and not by human will or experience: It was our evolutionary past and not parenting, education, or other socialization or cultural influences that determined one's personality or one's fate within human society. Galton plucked the term "nature versus nature" from Shakespeare's *The Tempest* to popularize his arguments.

This simple dichotomy has since been proven an error: There is no either-or relationship between nature and nurture. For example, Wright argues that evolution has prepared us to interact and adopt different strategies—from extremely altruistic to extremely selfish—based on the environment in which we are brought up and the feedback we receive from birth. In a sense, natural selection has favored flexibility (within limits) and the ability to choose from a pallet of behaviors in response to different environmental contexts. Similarly, Wrangham, in his work on male violence, concluded that, for example, evolution and natural selection have determined that chimpanzee groups are ultimately capable of greater violence than bonobo groups but that environmental and historical factors will determine the level and the number of violent incidents the chimpanzees carry out, just as environmental factors may in some cases make bonobos more violent.

Therefore, human nature is not entirely the result of experiences impressed on the tabula rasa of the brain after birth, as perhaps suggested by some cultural determinist literature. Some behaviors, such as male violence, may be more "natural" than we would like to believe. Yet, as both authors point out, this is a long way from saying that human behavior, and therefore human society, is genetically dictated. In Wrangham and Peterson's words, while the

whole logic of evolution would indicate that animals use their intelligence to serve evolutionarily appropriate goals, . . . inherited temperaments in different environments can express all sorts of different behaviors." (Peterson and Wrangham, 1996, p. 177.)

Both authors also stress that their research can offer insight into the evolved behaviors of different species or sexes but not between interspecies groups. Differentiation between human groups has not occurred over an evolutionarily significant time period; therefore, we cannot expect to find any evolved differences between human groups in the way we see such differences between the sexes (separate for most of evolutionary history) or between species.

Beyond this, they also state that just because a behavior has evolved and is therefore "natural" does not mean it is desirable and does not in any way connote an endorsement of the behavior. Speaking of an evolved, natural tendency toward violence or promiscuity does not suggest that either of these behaviors is "good" or beneficial to either the individual or to the group: Pursuing them brings no guarantee of happiness or health, merely that, in the past, these behaviors helped our ancestors pass on their genes.

Thus, both presentations draw four general conclusions about evolutionary psychology:

- 1 Some aspects of species and sex behavior can be understood as adaptations to evolutionary environments. According to both Wright and Wrangham, there is likely an evolutionary explanation for both complex and basic behaviors, such as military behavior or, indeed, morals.
- 2 Rather than supporting Galton's distinction of a dichotomy between nature and nurture, evolutionary psychology refutes it. Instead, it appears that, while certain behaviors are inherited, their expression in interactions is highly dependent on environmental conditions. Thus, evolutionary psychology may help us understand under what environmental conditions behaviors are expressed and by whom (see point 4).
- 3 Furthermore, violence or promiscuity may be natural, but they may also be "globally maladaptive." These are adaptations that facilitate individuals in passing on their genes while they also harm the species generally. The example both authors use is the canine teeth of male baboons. The longer and sharper the teeth, the more successful a male will be in mating but the more likely he is inflict serious wounds on females and other males in his group.
- 4 Evolutionary psychology provides little insight into individual behavior or into differentiation between human groups. Only distinctions that have

occurred within an evolutionarily significant timescale, such as between the sexes or between species, can be analyzed with the tools of the field.

- 5 Finally, while these theories clearly have policy implications, it is too soon to draw any hard and fast conclusions. Despite its close relationship to the field of biology, evolutionary psychology is very much a social science similar to anthropology, psychology, or sociology. Adopting standards of proof from these fields, evolutionary psychology can greatly inform the policymaking process

Richard Wright: Women, Men, and Moral Behavior

Richard Wright sought to answer three questions about evolutionary biology. First, what is it? Second, is it true (the focus of his talk)? Third, does it matter (that is, does it have policy implications)?

Evolutionary psychology, according to Wright, is the study of how our evolutionary environment has shaped human nature. It is not, however, the study of how specific genetic differences yield different behaviors. This is a separate field: behavior genetics. Instead, as with much of evolution theory, evolutionary psychology provides little insight into the differences between individuals but instead is useful in differentiating the behaviors of species or between the sexes and understanding how the behaviors came to be differentiated.

Related to the discussion of nature versus nurture above, one of the key points Wright made clear in his presentation was that differences between most groups of human beings are relatively recent in evolutionary terms. For this reason, it is unlikely that there are evolutionarily determined differences between races or ethnic groups. On the other hand, the difference between men and women is clearly evolutionarily significant, since the differentiation between males and females extends back to the earliest stages of our evolution. Thus, while this may be politically no less controversial, it is scientifically sound to speak of evolutionarily determined sex differences, while the notion of evolutionary determined racial or ethnic differences remains discredited.

Evolutionary psychology is an analytical tool that is a genuine refinement on the work of Darwin. It is based on our more recent understanding of genetics and the work of such authors as Richard Dawkins (1976), William Hamilton (1964), and Robert Trivers (1985). These authors used a genetic perspective (that is, from the perspective of the survival of individual genes rather than of individuals) to explain human behavior and the emergence of such confusing behaviors as altruism. These authors have forwarded convincing theories of how seemingly complex

behaviors and seemingly contradictory outcomes, such as altruism, could arise out of the process of genetic replication.

In his book, Wright argues this further, stating that such a perspective may not just explain obvious animal impulses, such as lust, and more complex behaviors, such as altruism, but also such notions as love. The fact that genetic and evolutionary explanations could be given for animal behaviors had gained new scientific credibility in the 1960s and 1970s and has never really contested since. It was the more "human" emotions and motivations that continued to vex the scientific community. Before discussing this further, he turned to the criticisms that the evolutionary psychology approach has received.

Wright discussed the three criticisms of evolutionary psychology in turn:

1. The field has low standards of evidence.
2. The field's hypotheses are "just-so" stories and are ultimately untestable.
3. Stephen J. Gould has criticized evolutionary psychologists for being too quick to assign observed behaviors to adaptation when they may be by-products of other processes or adaptations.

In response to criticisms about low standards of evidence in the field, Wright suggests that this relates to the field's historical connection to biology (as in E.O. Wilson's sociobiology), when it is in fact more closely related to sociology and psychology. As with these other "soft" sciences, the standards of proof must necessarily be less exact, given the subject matter and the inability to conduct controlled experiments. This is so not only because humans and human societies are involved, but also because it is impossible to speed up or reverse the process of evolution.

This answer is also related to the criticism of evolutionary psychology as a collection of just-so stories that seem to explain why things happen merely because they seem to explain why things are the way they are. Wright answers this criticism saying that, from necessity, evolutionary psychology is a historical science. It is clear that, as with the study of history, we cannot rewind or fast forward evolution. It is therefore necessary to reconstruct the human experience and hypothesize about its implications for our psychology. The study of history proceeds through the application of new evidence and new explanations of the existing historical record. In evolutionary psychology, the ability to look at other species makes it possible to compare experience and genetics. This provides a further layer of evidence beyond the historical record of human societies. Here, we can see patterns such as the sexual bimorphism in animals that is found in most species in which males fight over the scarce sexual resources of females. In the Phalarope species of bird, however, in which males incubate

eggs, females compete more intensely for mates and are larger and more colorful. This sort of evidence helps to understand human male and female behavior.

The final criticism itself needs a little explanation. Gould has criticized the evolutionary psychology approach by using the example of the "spandrels," the area between two adjoining arches seen particularly in medieval architecture. In his paper "The Spandrels of St. Marco and the Panglossian Paradigm," Gould argues that, in many buildings, particularly medieval cathedrals, the area between the arches was used to hang beautiful mosaics. Looking at these structures now, one might assume that arches had been adapted to this form, to hold the mosaic. As we know, the spandrel and the mosaics on them are in fact the byproduct of the arches themselves, which were used to support the heavy walls and ceilings of medieval buildings. According to Wright, Gould claims that evolutionary psychologists, like the social Darwinists and sociobiologists before them, are too quick to assign the emergence of characteristics and traits to adaptation. Their arguments may in fact be similar to arguing that the reason the nose is with us is to support our glasses.

Wright responds that, while the criticism must be well-taken, Gould has never given an example of such a "spandrel" in the literature. In a similar article in the *New York Times*, Gould used literacy as an example of something that cannot be explained by evolutionary psychology. According to Wright, this is quite true, but then again, to his knowledge, no one has been doing research to show that literacy can be so explained.

A question was then asked about whether evolution is over, to which Wright answered that, for practical purposes, evolution is over. That is, since evolution occurs so slowly from a human time perspective, it makes little sense to consider it as a dynamic process. Therefore, the 5,000 or so years since the discovery of agriculture are largely insignificant on an evolutionary scale, and the periods spent in nomadic groups and hunter-gatherer societies for millions of years previously are what must be considered as significant in shaping our species-specific behavior.

How certain is evolutionary psychology? Wright used the example of the Madonna-Whore dichotomy that observers have identified in male attitudes toward women in the vast majority of cultures. Here, men attempt to mate with as many women as possible but respect only those who are less promiscuous. Wright stated that the evolutionary psychology explanation for this behavior has important support but is far from proven, referring to the level of certainty as perhaps only 50-50. He develops this idea extensively in his book. He concludes that while the field does not use the term *spandrels*, it does recognize that there

certainly are many of them. Perhaps that is what differentiates evolutionary psychologists from earlier sociobiologists, such as E. O. Wilson: They are not intent on explaining everything from the evolutionary environment but are willing to assign many attributes to the by-products of the process of evolution and to other factors.

With the small amount of time remaining for his presentation, Wright focused on how the ideas of evolutionary psychology might be applied to policy. Again he stressed that it is too early to draw any clear conclusions, but by determining some of the reasons and the evolutionary patterns of our behavior, we can have a better grasp of what is necessary to influence it.

In the end, even after the discussion, it was difficult to get a grasp of how this world view might be applied to specific policies or even the making of policy. Despite earlier disclaimers, the discussion focused on the possibility for differentiation between human groups.

Robert Wrangham: The Evolution of Male Violence

Robert Wrangham's presentation was less general than Wright's, and hence, perhaps, more easily applicable to the policy arena. The first part of the presentation was based on the ground covered in his 1996 book with Dale Peterson, *Demonic Males*, which seeks to provide a theory explaining why chimpanzees and humans, out of 4,000 species of mammals and 10 million other species, are the only species in which groups of males (and sometimes adolescent females) have been observed hunting and killing males and sometimes females from a rival group. The premise of the book and of Wrangham's presentation is that understanding why only these two species do this may provide a basis for thinking about what natural selection is likely to have done to human psychology, particularly male psychology. In the second part, focusing on the differences rather than the similarities between chimps and humans, Wrangham applied these ideas to policy with a discussion of adaptive explanations for what has been termed human "military incompetence."

Wrangham's thesis, that male violence in humans has been determined by evolution, is clearly controversial. The popular perception is that violence is culturally determined. The recent Seville Statement on Violence declared that violence was not part of our evolutionary legacy and is not in our genes. Humans were able to invent both war and peace,; therefore, violence has nothing to do with evolution, the statement reasoned. Similarly, Gould (1996) continued his criticism of evolutionary psychology and, in this case, Wrangham's thesis in particular, by stating that if both "darkness and light" exist within our capacity

and if both tendencies operate with high frequency throughout our history, we learn nothing by speculating about the Darwinian heritage of these tendencies. While this criticism seems to make a good deal of sense, according to Wrangham, Gould then does go on to allow that at the very least such analysis might show when one might show one behavior rather than the other. As Wrangham is quick to point out, if evolutionary psychology does allow us to discover this, it would be a very useful thing to understand.

The notion that own-species violence is an uniquely human behavior and is therefore culturally determined has been severely weakened as the record of observation of other primates grows. While the group of Gombe chimpanzees that Jane Goodall observed in the 1960s appeared to be peaceful, it became clear by the 1970s that this picture was not complete. By that time, Goodall's original group of chimps had divided into two rival groups. It soon became obvious that, although these two groups had formerly lived together in peace, they now ferociously defended their territories against one another. In 1973, a freshly killed adult female was discovered, and the evidence pointed to other chimpanzees as the culprits. Not long after this, on January 7, 1974, a chimpanzee group of eight was observed to attack a single adult of the other group violently. The group of males and a single adolescent female proceeded to hold the victim down and bite, stomp, and tear at him until he no longer showed signs of life. Although the victim survived the immediate attack, he died shortly thereafter. By 1977, the entire breakaway group of seven adult chimpanzees had either been killed or were females forced to rejoin the original group.

A clear pattern of lethal raiding has since emerged. A male chimp becomes excited, and a border patrol forms. This patrol then sets out for the edge of the group's range. In the process, the group calls and listens to the neighboring range. If they hear a group of chimpanzees calling back, they wait for a while, then return to their own range and their own group. If, however, in their foray into the neighboring range, they spot an individual from the other group alone, their reaction is very different. They cease calling and silently stalk the lone individual. Falling upon him in far greater numbers, they savagely beat and tear at their prey and leave him (it is most often a male) dead or dying.

Although it was originally thought that this behavior may have been due to human interference and to Goodall's provision of food, similar activity has been observed when human interference was minimal and when no food had been provided. This makes it unlikely that human interference is the causal variable in this behavior.

While there have only been nine actual sightings of coalitionary violence or corpses, as evidence that would "stand up in a court of law," a large number of strange disappearances and suspicious bodies have been observed. More important, however, is the ample evidence of border patrol behavior. This pattern of raiding is similar to that of primitive human warfare, and the growing observational record suggests that rates of violent death among adults in chimpanzee groups and in preindustrial societies are similar and high, even by 20th century standards. Wrangham used estimates from 13 primitive societies to show a median male death rate from violence of 20 to 30 percent. Gombe chimpanzee males are estimated to die from violence at a rate of 30 to 40 percent (Goodall, 1986). Lawrence Keeley, using a group of 21 pre-state societies, found that the death rate from raiding is 0.5 percent per year (Wrangham, 1999; Keeley, 1996). Including the world wars, this is well above the rates for modern societies. The evidence of Wrangham and others suggests that the probabilities are similar for chimpanzees, although further evidence and research are needed.

As long as it was thought that intergroup violence, the sort of lethal raiding observed in the Gombe chimpanzees, was a distinctly human activity, there were some attempts to attribute it to a distinctly human trait, such as our larger brain. The fact that chimpanzees also carry out this type of violence suggests that brain size is unlikely to be a causal variable. In this same vein, chimpanzees have culture only in a very limited sense, in comparison to humans; therefore, it appears that attributing coalitionary violence and intergroup killing to human cultural factors alone is similarly insufficient.

So, what motivates chimps to take part in these lethal raids? To answer this question, Wrangham observed that there are specific conditions under which such lethal raiding appears among humans and among chimps. First, there must be persistent rivalry among groups or communities. Second, the costs of aggression must be low—a pattern that Wrangham likens to the Mafia who, it is often said, wait for a 3-to-1 advantage when they attack someone. In chimp and in primitive human raids, the aggressors are rarely injured, and the victim is most often killed or, at the very least, severely injured. Thus, the rivalry suggests that there is a benefit of some sort from reducing the strength of the rival and that the costs of doing so are low.

But why does this behavior arise among chimps and not among other apes? For Wrangham, the answer lies in "Fruit to Party Ratios": Chimpanzees operate in a world of intense scramble competition; food is often scarce and is always unevenly distributed over a wide territory. Groups of chimps move from food source to food source, consuming the food there, and then moving on. The larger the group, the greater the distances that must be traveled and the greater energy

expended. When the food supply is low or in areas with poor food, groups become smaller, and some chimps move as individuals. The patches of food will support large groups, and smaller groups and individuals have a better chance of survival.

Furthermore, males and females move at different velocities when moving between patches of food. Females with children move slower so that they can coordinate their movements with their young. Thus, females often move alone; in large parties, all the females are rarely in the group. Females are also less gregarious and have fewer bonds than their male counterparts, who reach the food sooner and as a group.

This all combines to make chimpanzee groups very unstable internally. All the members of a group are not always present, and particularly, the number of females in a group fluctuates. In particular, female chimps are unable to form coalitionary bonds, and there is no counterweight to male aggression. But equally important, the constantly changing numbers in the group mean that the power relationships between different groups fluctuate often, with the further result that it is very likely that a larger group of males will encounter a smaller group or a lone individual from a rival group. The opportunities for violence are great among chimps.

As a counterpoint, Wrangham uses the example of another closely related ape, the bonobo. Although bonobos live on the southern side of the Congo—quite near the Gombe chips that Wrangham and Goodall studied—they do not exhibit the same sex differences in grouping. Bonobos, particularly the males, are generally less violent than chimps. What explains this difference? According to Wrangham, the bonobos have much better access to food in their territory. On the southern bank of the Congo River, bonobos do not compete with gorillas and therefore have a more diverse diet, eating some of the higher-quality gorilla foods, such as some roots and grasses, as well as chimpanzee foods, such as seeds and fruits.

Wrangham reasons that this explains the greater group stability and a greater female presence in bonobo groups. More food means less scarcity: Groups can remain large, and the differing velocities of males and females are not a factor given the shorter distances between food sources. Female bonobos are as gregarious as males, interact with one another sexually, and form coalitionary bonds of their own. Thus, there is much less instability both between bonobo groups and between males and females within the group. The opportunities for intergroup raiding are therefore very much reduced, because individuals are rarely alone. Similarly, females are able to protect each other within the group, and male violence has fewer potential benefits from a genetic perspective.

Although both bonobos and chimpanzees exhibit similar behavior under the right conditions (that is, chimps can be very peaceful under stable situations, and bonobos can be violent in unstable situations), bonobos have not been observed to reach the extremes of violence of chimps, which appears to suggest an inherited behavioral difference. According to Wrangham, the explanation for the violence of chimpanzee lies not in determining the benefits to the violent chimp but in determining the costs to that chimp. Wrangham's thesis is that the costs of that violence are sufficiently low among chimpanzees (as opposed to bonobos) that, in the course of evolution, it has paid to knock off a rival even if there is no immediate advantage to this behavior. Over time, this behavior has benefited the progeny of successfully violent chimpanzee males.

Before applying these thoughts to military behavior, Wrangham touched on one other notion. Are the similarities between humans and chimps due to their shared descent (synapomorphy) or to some sort of convergence or coincidence? First, humans appear to have faced similar forms of scramble competition for much of their evolution, and it is often assumed that they were organized in similarly unstable social groups (what Wrangham refers to as fission-fusion parties or party-gangs). Then, there is extreme rarity of this sort of lethal raiding in the animal kingdom. Finally, there is the extreme genetic similarity between the two species: Chimps are our closest genetic relative. Taken together, this evidence suggests that the simplest explanation is that lethal raiding arose in an ape prior to the chimpanzee-human split, according to Wrangham. If this is true, lethal raiding has been subject to continuous selection for 5-6 million years. In this time, both humans and chimps have developed violent brains that are capable of premeditated, unprovoked violence.

After laying out this basic theory of the origins of male violence, Wrangham moved into what was at once a more speculative arena and one more interesting for policy. He presented idea from an upcoming article in *Evolution and Human Behavior* titled "Is Military Incompetence Adaptive?"

Simply stated, the basic question here is that, even if raids make sense from evolutionary perspective, since the costs are so low, do battles? In raids, as witnessed in chimpanzees and primitive societies, there is generally an accurate assessment of the victim's strength and that costs to the aggressors will be relatively low. This, however, is generally not the case in battles: Both opponents believe they will win; both engage willingly in many instances; and, therefore, one, if not both, of the participants has failed to make the proper assessment of the costs almost by definition. Furthermore, it is likely that both will incur relatively high costs, no matter who ultimately wins. This clearly describes a completely different logic from that of raiding and has been referred

to in the literature as military incompetence (Dixon, 1976), which Wrangham defines as protagonists losing even when they expect to win.

Battles are uniquely human. Even chimpanzees and other party-gang species do not engage in battles, just raids. Indeed, from the view of animal conflict theory, it would be puzzling if two opponents from the same species of clearly different strength would deliberately fight. The weaker animal should concede and retreat without fighting. Failure to do so would lead to certain injury and perhaps death. Even the chimpanzee victims discussed above, who have little chance of survival, do not stand and fight. Therefore, it must be concluded from the logic of animal conflict theory that conducting battles in which both sides expect to win when only one can is a maladaptive trait, unlikely to be passed on. It might be concluded that commanders who bring their forces to engage in battle must be either emotionally or cognitively incompetent (Dixon, 1976).

Yet, in human history, this military "incompetence" appears almost systemic. It is unlikely that a maladaptive trait would be so pervasive. If both humans and chimps have been selected for raiding, how is it that humans have become so incompetent as to engage in battles regularly? Wrangham argues that perhaps an adaptive explanation is wanting.

In a battle, failures of assessment on the part of one or both opponents are due to positive illusions. Wrangham has two hypotheses for why positive illusions may promote victory and therefore be adaptive. The performance-enhancement hypothesis states that positive illusions can be useful in suppressing negative thoughts, enabling an individual or a group to achieve a goal. The opponent-deception hypothesis states that humans tend to deceive themselves so that they can bluff successfully. Another way to state this is that positive illusions may allow either internal feeling suppression or behavioral leakage suppression, which in turn may lead to greater success in combat.

Positive illusions, however, also reduce the contract zone in which conflicts can be resolved to the mutual satisfaction of both parties. Both parties are selected to overestimate their own abilities and underestimate the abilities or motivation of their opponents. This in turn leads to escalating combat between individuals and increases the likelihood of injury or death.

This type of adaptation is referred to as "globally maladaptive." While it has created the sort of success that allows some males to be more successful than others, it reduces the fitness of the average individual. Wrangham uses the example of the long canine teeth of male baboons, which are due to an evolutionary arms race to possess the most-effective weapons. Although these teeth are advantageous in combats over mates, they also wound females and

other males. Similarly, Wrangham argues that self-deception has been positively selected in military conflicts, because an individual not possessing these traits (or possessing them to a lesser degree) will be ineffective against an opponent who is well possessed of such positive illusions and will be either hesitant or out-bluffed. Thus, although the individual with greater positive illusions is likely to be more successful for the species as a whole, this is still a disadvantage because such illusions lead to more frequent and more severe conflicts.

The level of self-deception through positive illusions might be expected to vary depending on whether the aggressor knows his opponent, his relationships with allies, or his moral ideologies. Variation based on these or similar factors supports the opponent-deception hypothesis, according to Wrangham, apparently because deception will be called for to varying degrees. However, the lack of any apparent relationship between the signals exchanged between opponents and the appearance of a self-serving or self-deception bias supports the performance-enhancement hypothesis.

A question to be answered, à la Gould's Spandrels, is whether this self-deception evolved to meet its current demand, or was co-opted into this role. Wrangham leans toward the hypothesis that self-deception among our ancestors filled a number of uses. There are elements of self-deception in altruism, conflict, and cooperation. It is unlikely that the resolution of violent conflict was of primary significance in this regard. These elements have, however, given some advantage in conflicts and therefore have become significant factors in this behavior.

As discussed above, this flies in the face of animal conflict theory, which assumes that selection should favor accurate assessment, or more specifically, accurate assessors. What it also says is that species with a greater capacity for self-deception would be expected to have longer and more-intense conflicts. Self-deception adds a variable to contests, affecting fighting ability and motivational strength, adding greater variability to the outcomes of these contests. The outcomes of these contests become more unpredictable.

Wrangham also applied the notion of self-deception to the arena Wright covered, the relationship between females and males. He cites research showing a correlation between testosterone and the presence of positive illusions among men and women. In women, self-regard, as measured by the degree to which subjects overranked themselves in a peer-ranking test, was correlated to the presence androgens. This suggests to Wrangham that one way in which testosterone may foster violent behavior is by the promotion of positive illusions about competitive ability.

Thus, an evolutionary history of raiding has left us with a tendency to attack when costs are sufficiently low, and, if Wrangham is correct, we can deceive ourselves into thinking the costs are low or the benefits (often moral benefits) are great. This, suggests Wrangham, is a possible explanation for "military incompetence."

While such authors such as Gould may reject an adaptive analysis of violence because such an analysis may suggest that, if these behaviors are adaptive, they are "good" or biologically hard wired, it appears that such an explanation offers valuable insight. Thus, even if violence is in some way an adapted behavior triggered by factors in our environment, it would be interesting to know how species vary in their capacity for self-deception—not because violence is good but because it may be globally maladaptive: "good" for the genes but not necessarily the species.

Conclusion

The conclusion to be drawn from these two presentations seems to be that there is ample evidence that there are biological bases for our behavior. These have been influenced by our evolutionary past. By studying this past, we can gain insight into our behavior today and thereby be better informed in making policy. What is also clear is that this study and its conclusions must be viewed through the perspective of the historical and social sciences, providing ever more insight but few hard-and-fast answers.

References

- Dawkins, R., *The Selfish Gene*, New York: Oxford University Press, 1976.
- Dixon, N., *On the Psychology of Military Incompetence*, London: Jonathan Cape, 1976.
- Fukuyama, F., "Women and the Evolution of World Politics," *Foreign Affairs*, Vol. 77, No. 5, 1998.
- Goodall, J., *The Chimpanzees of the Gombe: Patterns of Behavior*, Cambridge MA: Harvard University Press, 1986.
- Gould, S.J., "The Diet of Worms and the Defrenstration of Prague," *Natural History*, September 18–76, 1996.
- Hamilton, W.D., "The Evolution of Altruistic Behavior.," *American Naturalist*, 97: 354–56.
- Keeley, L.H., 1996. *War Before Civilization*. New York: Oxford University Press, 1996.

Trivers, R.L. ,1985. *Social Evolution*. Menlo Park, California:
Benjamin/Cummings, 1985.

Wrangham, R. and Peterson, D., *Demonic Males: Apes and the Origins of Human
Violence*. Boston: Houghton Mifflin, 1996.

Wright, R., *The Moral Animal: The New Science of Evolutionary Psychology*, New
York : Pantheon Books, 1994.

D. Implications for Public Policy of the Information and Communications Revolution

Leader: Dr. Bernardo Huberman¹

Research Fellow at the Xerox Palo Alto Research Center and
Consulting Professor of Physics at Stanford University
January 19, 1999

Our increasing use of information and communication technologies has brought about an enormous shift in our economy over this century. To make this point, Huberman quoted Alan Greenspan as saying, "America's output, measured in tons, is barely any heavier now than it was 100 years ago. In the same period, real GDP, by value, has increased 20 times." According to Huberman, this shift has been fueled by two factors: a revolution in communication technologies and a revolution in computing. This summary was written by Danilo Pelletiere.

Today we are able not only to send and receive more information cheaper and faster than before but to do more with it once we receive it. For example, a 3-minute phone call from New York to London in 1930 cost the equivalent in 1999 dollars of \$300. Today, the same phone call costs less than a \$1. Similarly, as Moore's Law continues to hold, computer power today costs only 1/100th of 1 percent of what it cost in 1970.

Even by this standard, however, the growth of the Internet in this decade has been astounding. In 1994, the World Wide Web consisted of about 1,000; in 1999, there are more than 400 million and counting. In 1993 there were 5 million Internet users; by 1998, there were nearly 100 million. In 1992 there were 10 Web servers; today there are perhaps over 5 million. Today, traffic on the Internet is said to double every 100 days.

Initially, beyond its amazing growth, the Internet stood out from other economic and social networks for its relatively egalitarian structure and the new relationships it allowed and encouraged. In its early years as a public network, Internet relationships and exchanges largely were anonymous, ephemeral, and

¹For additional information, visit <http://www.parc.xerox.com/istl/groups/iea/dynamics.shtml>.

dynamic. It was anonymous because there were few mechanisms for ascertaining whether a person was anything other than what or who he said he was. This allowed many individual users and firms to experience and profit from exchanges that would never occur in their own social and geographic environments. It also allowed firms and entrepreneurs to reach new markets with relatively few barriers to entry. The Internet was considered ephemeral because, when the information and personalities that populated it were taken off line, they could disappear without an official trace. There were originally no records or archives recording the traffic for later review, although this has changed of late. Above all, however, the Internet was dynamic, undergoing constant change driven by the largely uncoordinated decisions of millions of individuals and groups dispersed across the globe. These characteristics make the Internet a completely novel economic, social, and political space, promising novel opportunities for both business and governance.

The Internet's novel characteristics, however, are only one side of the coin. There is also great deal of uncertainty in an environment of anonymous, ephemeral, and dynamic information and exchanges. While these attributes provide some activities and users with economic and political opportunities they would not enjoy off line, anonymous, ephemeral, and dynamic relationships are not very conducive to moving mainstream commercial uses on line. As the Internet has grown, it has become apparent that it is subject to congestion and other negative externalities similar to those of other networks.

In his research, Huberman examines the emerging patterns of use in the Internet's dynamic nature. He has found that, as the Internet has grown and as its users have had to confront these problems, many of the emerging patterns of use can best be described and predicted using simple probabilistic models, such as those from economic theory. In one article, he wrote that

[l]ike consumers of natural resources or drivers during rush hour, Internet users, particularly "surfers" are faced with a social dilemma of the type exemplified by the well-known tragedy of the commons. (Huberman and Lukose, 1997.)

What his findings suggest is that a number of the problems that the Internet faces may generate policy and market responses conceptually (though perhaps not technically) similar to those found in the regular economy.

While Huberman's work has focused on observing and modeling emerging patterns of behavior on the Internet, in this presentation, he also offered some ideas about how the Internet is developing social mechanisms, such as branding, document banks, and communities to reduce the complexity and uncertainty of

the on-line environment. The conclusions offer some thoughts on how Huberman's presentation and research apply to our topic of global governance.

Modeling Social Behavior on the Net

With the exceptional growth of the Internet, there has been considerable interest in how people use the in and how their behavior on line is evolving. As more and more users get on line, the statistical patterns of use become more predictable. Thus, while it is perhaps impossible to predict any individual user's needs or actions, regular patterns of aggregate use emerge. Huberman's work describes and models behavior in three areas: social dilemmas and Internet congestion; managing the speed and risk of electronic transactions; and surfing as a real option, i.e., modeling whether an agent decides to surf from one Web page to the next (Lukose and Huberman, 1999). This section reviews the finding from this empirical work. The next section will discuss some of the institutional responses Huberman has observed as the Internet matures.

Congestion and Social Dilemmas

Huberman observes that, since most users are charged a flat rate, bandwidth is a scarce commodity on the Internet, governed by very much the same rules as any common good. Users do not pay according to their use, so they have an incentive to use as much bandwidth as they able to. As individuals, their actions will have little effect on the performance of the network. The well-known catch, the so-called "tragedy of the commons," is that, since most individuals are similarly motivated, the combined impact of this behavior degrades the performance of the network. Behavior that appears rational to the individual makes everyone worse off. The question then becomes, how do users react to this paradox on line?

Huberman and Rajan M. Lukose (1997) used data on the round trip time and route of Internet Control Message Protocol "pings" sent out over the network to study patterns of congestion and use. The time it takes for a ping to be sent, locate the receiving host, and come back is used as a measure of network congestion. The first test they ran simply sent out pings at minute intervals over the course of a day, measuring the travel time of each ping and noting its route. This test confirmed that there were daily periodicities in congestion and that, around these larger fluctuations, smaller, higher-frequency fluctuations also occur. While it seems likely that the wider daily fluctuations are caused primarily by exogenous factors (i.e., usage goes up in the hours when California

gets to work), the smaller fluctuations appear to be more a function of users' reactions to conditions on line.

To test for congestion patterns in these smaller fluctuations, a series of 10,000 pings was sent from a public workstation at Stanford University to a Web server in the United Kingdom over a 45-minute period. The transatlantic route is one of the most congested in the world, and the long distance provided some control for exogenous factors related to locality. The pattern Huberman and Lukose found was consistent with the hypothesis that usage of the Internet is being treated as a public good. It appears that users consume bandwidth greedily, until the whole Internet's performance is degraded, making all users worse off. Differing reactions by individuals to this dilemma lead to the emergence of intermittent congestion spikes with definite statistical properties. Faced with increasing congestion, most users become cooperators, ceasing their activity until congestion eases, while defectors continue to consume bandwidth as it becomes available. The more cooperators there are, the more quickly congestion is reduced, and the greater the overall benefits for the collective network.

The practical implications of this is that a better tools might be developed for accessing remote sites on the Web during periods of high congestion, increasing cooperation or anticipating the behavior of defectors. More generally, however, Huberman concludes that some form of congestion pricing would ease congestion and change this pattern of use. Under current pricing, however, the Internet presents an interesting environment to study the human response to the windfall opportunities and uncompensated costs that network externalities can generate.

Managing Speed and Risk

Congestion adds variability and therefore risk to the on-line environment. Having described the model of how congestion forms and subsides, Huberman moved to the of problem how to execute electronic transactions efficiently over the Internet. Using the time it takes to make a transaction as the metric of efficiency, Huberman proposes a method for quantitatively managing the risk and cost of executing transactions in a distributed network, such as the Internet. The methodology he proposes assumes that congestion will be managed by some form of usage-based pricing and is analogous to the rational investor of modern portfolio theory. He discussed strategies both with and without such pricing.

Related to the congestion discussion above, what response to congestion or expected congestion is likely to generate the best results? For example, if a message is sent and no response is received after a certain time, the user can

continue to wait, can send another message, or can cancel the first (protocols allowing) and send a new message. How long he waits and what strategy he chooses depend on the costs, perceived risk, and expected time of each transaction.

Under the current cost structure of the Internet, cost is not a factor, and expected duration and risk are all that controls the decision to resend or restart a query. As we have seen, this leads to the social dilemma resulting from congestion, since it is in every user's interest to send multiple messages or to try multiple paths of access. Here, the optimal strategy is dynamic, depending on actual on-line conditions and predictions of future use. These factors might be measured and anticipated by artificial agents that would decide on the best strategy.

With a simple pricing mechanism in which charges are proportional to the amount of data users send, the expected time, risk, and monetary cost all influence the decision of whether to resend or restart the transaction. Such a system is currently in place in New Zealand. By adding a simple cost dimension, sending 100 similar messages would now be 100 times more costly than sending one. Therefore, even if sending 100 messages would greatly reduce the expected time and risk involved in completing the transaction, the cost would likely be prohibitive, a disincentive to flooding the network.

Another way to handle congestion pricing is with priority pricing. Under this pricing model, the cost of sending messages could be proportional to the priority the sender placed on it. This would also allow the possibility of sending the same message at different priorities. Using a simple model, Lukose and Huberman (1999) show that a mixed-priority strategy is always preferable, providing a better mix of risk and cost than single-method strategies. Lukose and Huberman compared the calculations that users would make under such a pricing system to those used to choose stocks. Their results are similar to those of modern portfolio analysis, in which distributed cost and risk portfolios can be expected to outperform either the highest returning or least risky assets alone. So, given patterns of use and congestion, their model suggests that, under such a system, it will always be best to use a mixed strategy, combining risk and cost and never choosing the highest-priority or lowest-cost options in transacting business on the Internet.

Surfing as a Real Option

Similar to the case of congestion for data transfer in general, it seems that surfing behavior can be described by a law determining the probability that a random user will surf given Web sites for a given number of links. Huberman and his

colleagues refer to this as the "law of surfing." They again use an analogy from economics and finance to build a model of behavior on the Internet. The analogy here is with the financial notion of a "real option," the flexibility a firm has to invest under uncertainty. Here, the decision of an Internet user to continue surfing to the next Web page is assumed to be a function of the utility of the previous pages accessed and a probability distribution for the number of pages the user is statistically likely to visit before stopping or returning to the original page. When a page is visited that is valuable and provides links to further pages, it is assumed that the reader will be likely to look at the next page assuming that it too will be useful. Future choices are a function of past choices.

They used data collected from a representative sample of (anonymous) America Online Web users on December 5, 1997, to verify their law of surfing, performing detailed measurements of surfing patterns. This empirical test showed that a simple economic model can explain the observed distributions in path lengths and page visits to sites on the Web.

Market and Policy Responses

So, the emerging behavior of users on the Internet can apparently be usefully described using basic economic and financial models. Similarly, the problems that users experience on line are being addressed with institutional solutions analogous to those found in the regular economy.

As discussed above, congestion is not the only obstacle to Internet commerce. Beyond this, it is well-known that the relative anonymity of on-line information and lack of verification of commercial make this information difficult to use. Similarly, the ephemeral quality of the Internet can make it hard to rely on as a steady source of information, goods, or services. Finally its dynamic nature reduces the familiarity of the environment between visits and makes it difficult to know what to expect from subsequent searches and interactions.

These problems have led to a number of policy and market responses. Just as in other market places, branding and third-party verification are playing an increasing role in the value of Internet sites and commerce. Also, established brands from the traditional economy, such as Barnes and Noble, and new Internet brands, such as Amazon.com, are playing an increasing role in the Internet economy. These sites steady and familiar sites for consumers to come to know and return to. The ephemeral nature of the Internet is further being addressed through document banks and electronic databases that hold information from the Web for posterity. And stable communities are emerging from extensive interaction that reduce the uncertainty of what or whom you

might meet. In short, institutions are emerging on line to reduce uncertainty and increase trust.

Branding in Cyberspace

Huberman spent some time discussing the issue of branding in cyberspace. First he provided a simple definition of a brand: a simple encoding of the attributes and reputation of an individual or firm. The purpose of a brand is to solve the problem of having to search a large space or to choose from among many options. The market value of a brand can be established when a firm is purchased or hired; it is the value paid above all the tangible assets. This is what allows brands to be traded, and in theory, they can be in equilibrium.

What are the observed dynamics of brands in cyberspace? In cyberspace, many firms can easily offer the same service to the same markets. These brands interact either as substitutes (competitors) or as complements in what is known as umbrella branding, in which complementary brands are associated in the consumers mind. As the number of brands for a service increases, however, instability occurs as the value of brands (for minimizing search costs) diminishes. Only a few brands will emerge from this instability, and there is greater concentration in the industry.²

Huberman also addressed the issue of privacy and the need to find mechanisms that allow consumers to reveal preferences without revealing private information about themselves. He suggested that, while we could legislate privacy, as they are doing in Europe, a preferable method would be to create technical mechanisms that provide trusted third parties and similar institutions (Huberman, Franklin, and Tagg, 1999).

Governing the Internet

There are two questions that Huberman's presentation might have addressed: (1) What regulation and forms of governance are needed on line? (2) What implications does the Internet have for governance in general? Clearly, his presentation did more to answer the first question.

It seems that, statistically, on-line behavior is more regular than the characteristics of the Internet might suggest. This may mean, for example, that

²This led to a more general discussion by the group on whether branding was easier or more difficult on line than in the regular economy.

simple pricing policies will go a long way toward reducing congestion. As we have seen with highway systems, however, such pricing schemes are likely to meet political and consumer resistance. Huberman's research was published in 1997. To date, there does not appear to be a major move away from flat-fee cost structures. Most commercial Web sites continue to pursue the advertising and commerce models for development and operation. In such an environment, businesses want to encourage use. Therefore, capacity remains the preferred congestion management strategy.

But what is also clear is that many institutions for regulating commerce are migrating from the regular economy to the Internet. Branding, archiving, and community development were all offered as examples. When it comes to the role of the Internet in changing the nature of governance, however, Huberman's presentation offered little insight.

References

- Lukose, Rajan M., "Economic Modelling of the World Wide Web," to appear in *Artificial Societies and Computational Markets*, Cambridge, Mass.: Massachusetts Institute of Technology Press, 1999.
- Huberman, Bernardo A., "Social Dilemmas and Internet Congestion," *Science*, July 25, 1997.
- Huberman, Bernardo A., et al., "Strong Regularities in World Wide Web Surfing," *Science*, April 3, 1998.
- Huberman, Bernardo A., Matt Franklin, and Tadd Hogg, "Enhancing Privacy and Trust in Electronic Communications," *Proceedings of the ACM [Association for Computing Machinery] Annual Meeting*, New York: ACM Press, 1999.
- Huberman, Bernardo A., and Rajan M. Lukose, "Social Dilemmas and Internet Congestion," *Science*, Vol. 227, July 25, 1997.
- _____, "An Economics Approach to Hard Computational Problems," *Science*, Vol. 275, January 3, 1999.

E. Biology, History, and Social Organization

Leader: Robin Fox,
Rutgers University
February 18, 1999

Introduction

On February 18, 1999, the sixth in a series of study group meetings was held at RAND. The series, funded by the Defense Advanced Research Projects Agency (DARPA), focused on the question of whether national and international organizations and current legal systems are adequate to deal with the significant issues of security and governance that emerge from the information and biological revolutions. The potential shortfalls of the current political and legal infrastructure are becoming evident as questions are debated such as research into human cloning or the extent to which privacy on the Internet should be protected. This paper presents a summary of ideas and thoughts presented at the study group meeting.

The study group invited Robin Fox, University Professor of Social Theory at Rutgers University, to present his thoughts on the effects of biology on social structure, organization, and human behavior. A wide range of interests, research, and publications on the human social condition have marked Professor Fox's career. Among his many publications are *Kinship and Marriage*, a widely read anthropological text. In 1970, Dr. Fox and his colleague, Lionel Tiger (a former study group discussion leader), sparked a national debate about biology and behavior with the publication of *The Imperial Animal*, one of the first attempts to introduce ethnological ideas into the social sciences. His works on the evolution of behavior have included *The Red Lamp of Incest: An Enquiry into the Origins of Mind and Society*. Professor Fox was educated at the London School of Economics and Harvard University, and did his postdoctoral work at Stanford University. He has conducted field work in many countries around the world.

Presentation by Robin Fox

How does an anthropologist interested in the story of human evolution view history and social organization? Perhaps the most I can do is to urge a change in perspective.

The Republic was a seminal contribution to our effort to understand morality and life. In it, Plato relates his *Parable of the Cave* in an effort to show how humans bridged the chasm between "unwisdom" and the knowledge of the truth. I have taken the liberty of modifying it a bit to reflect an allegory of the circus. I have entitled it, "One More Hoop for the Tiger." The parable opens with Socrates speaking and the ever obsequious and uncritical sycophant Glaucon listening:

SOCRATES: Next, said I, here is a parable to illustrate the degree in which our nature may be enlightened or unenlightened. Imagine the conditions of men living entirely under the big top of a grand circus. There they've been since childhood, so all they know is what they have seen in front of them.

GLAUCON: Now I see.

SOCRATES: Now in the three rings of the circus, the animals are performing various elaborate tricks. The horses and zebras are standing on their hind legs and walking backwards. The elephants are making a line with each elephant's front feet on the back of the one in front, and they too are walking on their hind legs. The clowns are walking on their hands with hats on their feet. The spectators are amused.

GLAUCON: Naturally so.

SOCRATES: The acrobats meantime are attempting the world's highest human pyramid. And, finally, the tiger is jumping through not one, but two rings of fire.

GLAUCON: It is a strange picture and a strange sort of prison.

SOCRATES: Like ourselves, for we also take for granted all what we see as happening is all that can happen.

GLAUCON: We do indeed.

SOCRATES: You're getting positively irritating. However, to continue with the parable, suppose it was thought that the tiger could only be persuaded or trained to at most jump through two hoops, and this was positively accepted as a fact of nature.

GLAUCON: Well, let us suppose.

SOCRATES: Now suppose that the trainer announced that he intended to prove that the tiger could jump through three or more hoops of fire. Would not the spectators consider this the most significant and marvelous thing in the world?

GLAUCON: They most certainly would.

SOCRATES: Okay, now consider what would happen if they are released from their chains and the healing of their unwisdom should come about this way. Suppose one of them were released and taken into the outside world, where he saw the natural habitats of the animals, and even the human performers. What do you think he would say if someone told him that what he had formerly seen before was totally unnatural and what he now saw was nearer to reality and a truer view of life? Would he not be perplexed?

GLAUCON: Yes, truly perplexed. Sorry, it gets to be a habit.

SOCRATES: What if he were dragged to the top of a hill and shown the panorama of nature with the animals running free on four legs and men on two legs chasing them, and the whole assemblage falling into families and tribes and herds of complex social hierarchies and relationships? Would he not be astonished and disbelieving until his eyes told him that what he saw was truly not an illusion but indeed a reality?

GLAUCON: Yeah, yeah, yeah, most truly astonished, etcetera and so forth.

SOCRATES: There's hope for you yet. But let us ask then, suppose our released prisoner thought back to his fellow prisoners and what passed for wisdom in his former dwelling place. He would surely think himself happy for the change and be sorry for them and not envy them their enjoyment knowing that they mistook the circus for reality. Don't answer.

SOCRATES: Now imagine what would happen if he went back to take his former seat at the circus and told the spectators that he had seen the real world—men and animals—and it was quite unlike the world they judge to be real. What's more, what if he told them what they regarded as dangerous and unacceptable behavior in the animals is just what he had seen outside the tent? They would laugh at him and say that he had gone out only to come back with his sight ruined; that it was not worth one's while to make the attempt. If they could lay hands on the man who was trying to set them free and lead them out, they would kill him.

GLAUCON: Yes, they would and who would blame them?

SOCRATES: G, you are stepping out of character. Every feature of this parable, my dear G, is meant to fit our current analysis. The circus, my sycophant, is history. The spectators are those of us trapped into being unable to see the bounds of historical contingencies. The escaped stranger is one who has seen nature whole and uncontaminated by history.

GLAUCON: When do we get to kill him?

I have great sympathy for G, in this instance. The point is almost childishly simple.

Bearing this in mind, and with respects to Fukuyama's *End of History*, it seems to me that the issue is not really knowing about the end of history, which may or may not happen, but rather its beginning. The end of history we may not live to see; however, the question of where history began is one that can, at least in theory and in principle, be settled.

I remember being taught history in an old English school. The syllabus went strictly by centuries, and you learned the century your history teacher assigned. Mine was the 18th; to me the world began in 1715 and ended in 1815. Nothing important or of consequence happened before or after those two dates. I remember being astonished when my mother gave me a copy of H. G. Wells' *Outline of History* for my sixteenth birthday. I opened it expecting to start at 1715, but found that it started with the origins of life on earth. I guess that was my first introduction to the question of where history began. Of course, we now know Wells' information was wildly inaccurate. Nevertheless, the whole idea shook me. I was brought up distinguishing between history and prehistory. Traditionally, history began wherever prehistory ended. As a result, evolution is not taught in history departments. I am going to question that decision as we go along. My aim is to play the devil's advocate and ask why.

Most history, as it is taught, begins with ancient civilizations. While perhaps a brief nod is given to beginnings in early towns and the like, everything preurban is scarcely counted, by and large. This is the province of prehistory and archeology, not history. Some argue that history begins with written records, which takes us back to no more than 3,000 B.C. and mainly uninteresting lists. Behind this view is the general idea that history is the study of change and progress—a development upward toward something greater. Only after the advent of civilization and writing, according to most historians, did pattern movement emerge, empires rise and fall, and speculation arise as to why these things occurred.

We forget that that was once a new idea and that such noted philosophers as Hegel and Weber (and now Fukuyama) saw history as essentially cyclical. They saw patterns in what they decided was history. Before that there had certainly been societies of some sort, but nothing had happened essentially. This was the accepted wisdom.

The convolution of history with progress—the notion of history reaching some ideal goal—can be traced back to Judeo-Christian beliefs, namely the apocalypse and the end of time. In the 19th century, most people seemed to embrace this perspective. Even H. G. Wells, the noted author and social commentator, spent his life wavering between pessimism about the future and delivering a message

of scientific optimism. Toward the end of his life, he wrote the most extraordinary treatise (which never seems to get read) called *Mind at the End of its Tether*. Having just seen the end of World War II, he argued that the history of *Homo sapiens* had played out—that as far as he was concerned, we had reached the end of history.

The idea of upward progress was seriously challenged in the 20th century. Some have even doubted whether progress occurred at all. Yet despite two world wars, the Holocaust, and the apparent readiness of superior civilizations to decline into savagery, there is general agreement that progress has been made during this period, at least in terms of science and technology. The idea of history as marching forward is extremely powerful.

The only people who have demurred from this progressive view of history, apart from archeologists and anthropologists whose interests are restricted to the artificial realm of prehistory, have been the philosophers. Locke and Rousseau through Rawls focused on the state of nature and the first man. Even among the later theorists of the social contract, this idea of a natural state of man prevailed. However, these are theoretical constructs that represent their idea of what nature would be if one could strip civilization away. They are not derived from the observation of nature, but from the needs of theory.

The important question then is not about the end of history and the last man, but rather the beginning of history and the first man. If we were to abandon all these self-serving definitions, we would be left with what I call a naturalistic view which references only the facts that we know. We have a species, *Homo sapiens*, and a subspecies, *Homo sapiens sapiens*. This species has a duration in time. It came into existence and was gradually differentiated from other related species. From this evolutionary perspective, history is the story of the species, and any questions as to when it started are thus avoided.

The whole scope of history encompasses far more than the beginning of our own species. Another 25 million years of mammalian history needs to be considered. As mammals, part of our essential history is that we are large, bipedal, terrestrial, slow breeding, and land dwelling. Even if we narrow our focus to only the hominids, we still have at least five million years to consider, given that our earliest ancestors broke off from the chimpanzees, our closet relatives with which we share more than 98 percent of our genetic material. Of *Homo* proper, we have at least two million years. And of *sapiens*, our own species, we have more than 300,000 years.

From this perspective, what is called history—even under the broadest definition—accounts for only 10,000 years of an unusually warm interglacial

period. These warm periods coincide with big leaps forward in evolutionary progress. Our current interglacial is little more than a blip at the end of hominid history. In fact, research indicates that we are near its end—most interglacials have not lasted more than 20,000 years. A new ice age takes only the minutest wobble in the earth's axis to occur and can happen with frightening suddenness. Indeed, a very small increase in the polar ice cap would cause Chicago to be under half a mile of ice within 50 years.

So what we may be experiencing are quite unusual and peculiar happenings in the little blip of an interglacial at the very end of the trajectory of human history. What if we were to project ourselves into the future, past another ice age and into an interglacial 50,000 years from now? What would the people of that time think of this one? What they may see as the remains of a brief and unsuccessful episode in history, we now think of as the whole story. We may be prisoners of the proverbial Platonic circus, believing that reality is the circus when it may just be a temporary aberration.

But what good does it do to adopt this point of view? If we have to live and survive in history, what choice do we have but to perform in the circus? Maybe the important thing is that the tiger gets through that third hoop. Perhaps this is reality. I would submit that a change in perspective might benefit us all. A more comprehensive view, not just of one culture or time period, but of the entire thing we call history might make us think differently about what we are doing in the circus. It might make us reconsider what is normal and what is aberrant, what is healthy and what is pathological. We might end up thinking very differently about the way we treat the societies we live in. At least it might stop us from trying to cure things we cannot cure, or change things we cannot change, and enable us to focus our efforts on more useful endeavors. Maybe we can make a better circus. Maybe we can be nicer to the animals.

The time span of the genus, *Homo*, or what we regard as history, is less than 0.5 percent of the total of hominid history. In other words, if we were to make an hour-long film of this period, "history" would pass by in the last few seconds and industrial civilization in the last two frames. We would not even notice it. With this in mind, perhaps we should reconsider the exclusion of evolution from history. Evolution is not prehistory, it is simply history over a very long period. It is a time span so long that significant genetic changes have naturally occurred within it.

From this perspective, the following conclusions can be drawn:

- First, the notion of history should be redefined. History is a brief experiment that may or may not work. For example, the Neanderthals did not work—

they are extinct. Yet they lasted many, many times longer than the period in which we define history. In other words, we have to treat history as problematic, not take it for granted.

- Second, if the progressive attitude—the cavemen-to-computers mindset—is not accurate, history as we know it could be a complete aberration—a colossal evolutionary mistake. We have no reason to suppose that “success” is inevitable or that what we are experiencing is “success.” Some have already begun to pose these very questions in our treatment of the environment. This focus may be one of the first rumblings in realizing that the future is not all onward and upward; that what we are doing may be fouling our own nest rather badly. I find fascinating this talk of taming nature, and I think that is what gave rise to my parable of the circus. Yes, we have tamed the animals, even trained them to dance on their hind legs and jump through hoops of fire. But is this good for the them? And if the animals suddenly refuse to do this, and turn around and start rendering us apart, is this normal or unreasonable behavior on their part? I will pursue this a little further with some examples.
- Third, if progress as we know it is abnormal or a mistake, or at least we can treat it as such, then what is normal? If traveling at 485 miles per hour and realizing nations with populations of more than 800 million people is not, in fact, a sign of progress, how do we decide what is normal? In other words, how do we determine what the animals in the circus should be doing?

In my book, *The Red Lamp of Incest*, I looked at the social organization of all the terrestrial primates. Given the enormous differences that exist between the various species, I tried to determine if there is a set of factors that such societies share. I found that most primate societies broke down into factions.

As shown in Figure E.1, the females and their offspring represent one of the major interest groups. Young females tend to stay with their mothers and the other females, while the young males, when they reach adolescence, move to the periphery, where they set up unstable hierarchies. How these peripheral males make their way into the male hierarchy so that they can breed is central to the organization of primate societies.





Figure E.1—Social Organization

The hierarchical males are organized in ritualized hierarchies by polygamous groups that are more or less egalitarian. The peripheral males try to make their way back into these groups. To do so, they must learn to control their primary impulses through equilibration, or risk attack. If successful, they earn apprenticeships in the hierarchy, and are tolerated by the hierarchical males until they become full members and earn the right to breed. While these "cadets" attempt to engage in sex prior to this time, they are mostly unsuccessful because of the influence exerted by the dominants. Thus, the only significant way of

including their genetic material into the gene pool is to attain full status in the hierarchy of males.

Hierarchical males and females have different agendas and concerns (attributed by sociobiologists to different reproductive strategies), but they do manage to accommodate each other. In addition to protection, grooming, and sex, longitudinal studies of primates have found that kinship is an important link. While paternity is not something that is usually recognized in primates, ranked maternal families form kinship groups consisting of females and their offspring. The dotted line between the hierarchical males and the peripheral males indicates the distance between the groups. It is a wise chimpanzee who knows his own father.

The evolution of brain size is depicted in Figure E.2. For several millions of years, our australopithecine ancestors rumbled along with brains not much bigger than the chimpanzee. Then, about two million years ago, brain size took off in one of the most dramatic examples of growth in the history of organ evolution. Over the next two million years, their brains tripled in size, reflecting the transition from primate to human.

Figure E.3 shows the social organization after the transition, with presumed modifications from the baseline. The females continued to give up boys to the uninitiated males. The uninitiated males continued to form themselves into groups and had to undergo initiation to enter the society of initiated males. If successful, they became part of the initiated male group, first as a junior male and then as a senior member of the hierarchy. The female group underwent some separation due mostly to the presence of postmenopausal females—a phenomenon that does not normally occur among most sexually reproducing species. In general, females who cannot reproduce die off rather quickly. Among hominids, however, menopausal females grew increasingly common due to expanding average life spans.

Changes in kinship caused the formation of another distinct segment, young girls. After the transition, kinship began to take on an allocating function instead of simply linking individuals to one another. Henceforth, young, nubile females were allocated to senior males for the purposes of breeding. In this manner, the link between father and son was forged, and paternity was discovered. Such allocation gave rise to relations between cousins, a normal occurrence throughout the course of human history.




Figure E.2—Evolution of Brain Size





Figure E.3—Social Organization After Transition

This extraordinary period of brain evolution was the distinctive period of human evolutionary adaptation. During this time, the framework remained essentially the same, but the accommodation between the groups changed rather dramatically: We became what we are. Since then, an extraordinary warm

interglacial arrived; the Neolithic revolution occurred; and the domestication of animals and plants began. Then things really took off. All of this has been widely regarded by historians as progress. Compared to a cave, I suppose that a house seems like a step forward. But what if we are all prisoners in this circus, unable to see beyond its three rings? What if one of us were to escape? How might he regard history?

Figure E.4 explores this question a bit further. I took this idea from the work of Daniel Bell about the coming of posthistoric and postindustrial society. He writes about society getting to the point where the pendulum has swung too far. Though he never explains what he means by this, I decided to take a sort of Paleolithic shaman's eye view of the future. Instead of onward and upward development, suppose progress occurred in a series of pendulum swings—away from the Paleolithic norm. So I tried to draw these swings, picking out what I saw as the major leaps forward in human history.




Figure E.4—History as a Series of Swings of a Pendulum

These events seem to come at the peak of the pendulum swings, which got wilder and wilder as the population got larger and larger, and as technological innovations crowded in faster and faster. The major wars, which themselves were responsible for many of these great moves forward, seemed to occur on these major pendulum swings. Of course, I am not sure whether World War III would occur on the next pendulum swing toward a postindustrial society or the swing after that.

My general formula is that we must not mistake what is normal for what has been established by history, for that itself may be abnormal. A great many things we consider to be abnormal or pathological may just be an attempt to get closer to normality, as defined by the basic conditions of the human evolutionary adaptation. They might not be pathologies at all, but indigenous healing processes—what I call immune responses of the body social. Sometimes, the tiger refuses to jump through the third hoop.

F. Triumph or Tragedy: The Moral Meaning of Genetic Technology

Leader: Leon R. Kass

Addie Clark Harding Professor in the College and the Committee on Social Thought, University of Chicago, 1998–1999 William H. Brady, Jr.,
Visiting Scholar at the American Enterprise Institute
March 18, 1999.

As one contemplates the current and projected state of genetic knowledge and technology, one is astonished by how far we have come in the less than 50 years since Watson and Crick first announced the structure of DNA. True, soon after that discovery, scientists began seriously to discuss the futuristic prospects of gene therapy for genetic disease and of genetic engineering more generally. But no one then imagined how rapidly genetic technology would emerge, as the direct consequence of new, utterly unforeseen techniques for DNA recombination. Within a few years, we will see the completion of the Human Genome Project, disclosing the DNA sequences of all the 100,000 human genes. Today, genetic technology companies are thriving, even on incomplete genomic knowledge; the research director for SmithKline Beecham reported at a recent meeting that his company already has enough genetic sequencing data to keep his researchers busy for the next 20 years, developing early detection screening techniques; rationally designed vaccines; genetically engineered changes in malignant tumors leading to enhanced immune response; and, ultimately, precise gene therapy for specific genetic diseases. The age of genetic technology has arrived—and with it, much public anxiety and a growing attention to the some of the attendant ethical issues.

Genetic technology comes into existence as part of the large humanitarian effort to cure disease, prolong life, and alleviate suffering. Attached to the intrinsically humane and morally purposive art of medicine, genetic technology arrives to begin with wrapped in the highly moral mantle of generous and philanthropic humanitarianism. Occupying the moral high ground of compassionate healing, biomedical technology usually receives a royal welcome in our society, even when it raises challenges to other traditional moral norms. To a large extent, the same will be true of much of what genetic technology has to offer in the future. Who would not welcome genetic surgery that corrected the genetic defects that

lead to sickle-cell anemia, Huntington's disease, and breast cancer or that protected against the immune deficiency caused by the AIDS virus?

But genetic technology strikes most people as different from other biomedical technologies. Many people are concerned, anxious, and afraid of "tampering with human genes." Even knowledgeable people, duly impressed by the truly astonishing genetic achievements of the last decade and eager for the benefits, are nonetheless ambivalent. For they sense—I think rightly—that genetic technology, while in some respects continuous with the traditional medical project, is also in decisive respects radically new and, therefore, disquieting. Often hard-pressed to articulate the precise basis of their disquiet, they talk rather in general terms about the dangers of eugenics or the fear of "man playing God."

Enthusiasts for genetic technology, made confident by their expertise and by their growing prestige and power, are often impatient with the public's disquiet. Much of it they attribute to ignorance of science: "If the public only knew what we know, it would see things our way and give up its irrational fears." For the rest, they blame outmoded moral and religious notions, ideas that scientists insist no longer hold water and only serve to obstruct scientific progress. But this sincere yet also self-serving attempt to cast the debate as a battle of beneficial-and-knowledgeable cleverness versus ignorant-and-superstitious anxiety cannot succeed. For the public is right to be ambivalent about genetic technology, and no amount of learning molecular biology and genetics is going to allay its—our—legitimate human concerns. Rightly understood, these worries are, in fact, in touch with the deepest matters of our humanity and dignity, and we ignore them at our peril.

I want this evening to try to articulate some of these concerns, in the hope that we might be less heedless, less arrogant, and more sober as we hurl ourselves forward we know not where. Rather than speak about some ethical questions raised by the use of this or that technique, I want us to consider the moral meaning of the entire enterprise. To do so, we must bear in mind that genetic technology cannot be treated in isolation, but must be seen in connection with other advances in reproductive and developmental biology, in neurobiology, and in the genetics of behavior—indeed, with all the techniques now and soon being marshaled to intervene ever more directly and precisely into the bodies and minds of human beings. I shall proceed by raising a series of questions and comments, the first of which is an attempt to say how genetic technology is different.

What Is so Special About Genetic Technology?

Genetic engineering, when fully developed, will wield two powers not shared by ordinary medical practice. First, medicine treats only existing individuals, and it treats them only remedially, seeking to correct deviations from a more or less stable norm of health. Genetic engineering, in contrast, will deliberately make changes that not only are transmissible to succeeding generations but will even alter in advance specific future individuals (through direct germ-line or embryo interventions). Second, genetic engineering may be able (through so-called genetic enhancement) to create new human capacities and, hence, new norms of health and fitness. True, for the present, genetic technology is being hailed primarily for its ability to improve diagnosis and treatment of disease in *existing* individuals. To the extent that it would and could be confined to such practices, it would raise few questions beyond the usual ones of safety and efficacy. Even intrauterine gene therapy for existing fetuses with diagnosable genetic disease could be seen as an extension of the growing field of fetal medicine. But there is no reason to believe that the use of gene-altering powers can be so confined, either in logic or in practice. For one thing, germ-line gene therapy and manipulation, affecting the unconceived and the unborn, is surely in our future. The practice can be given numerous justifications, beginning with the desire to reverse the dysgenic effects of modern medical success.

Ordinary medicine is not without heritable genetic consequences, though these are not the deliberate or direct goals of therapy, but rather its unintended by-products. Thanks to medicine, individuals who would have died from, say, diabetes now live long enough to transmit their disease-producing genes. Why, it has been argued, should we not reverse these changes by deliberate intervention? More generally, why should we not effect precise genetic alteration in disease-carrying sperm or eggs or early embryos, to prevent in advance the emergence of disease, which otherwise will later require expensive and burdensome treatment, genetic or other? And why should not parents eager to avoid both the birth of afflicted children and the trauma of genetic abortion be able to avail themselves of germ-line alteration? Even before we have had more than trivial experience with somatic gene therapy—none of it successful—sober people are calling for overturning the existing self-imposed taboo on germ-line modification.¹ Never mind the severe ethical impropriety of experimenting upon the unborn (who cannot give their consent) or the countless mishaps that will have to be discarded. The line between somatic and germ-line modification cannot hold.

¹See, for example, Walters (1991).

Neither can we hold or defend the line between therapy and genetic enhancement, despite the naive hopes of many that this will prove possible. Would people reject additions to the human genome that enabled us to produce, internally, vitamins or amino acids we now must get in our diet? Would we oppose the insertion of engineered foreign genes that would be antibiotic to bacteria and parasites or would offer us increased resistance to cancer? Alterations in the immune system that would increase its efficacy or make it impervious to HIV? When genetic profiling becomes able to disclose the genetic contributions to height or memory or intelligence, will we deny prospective parents the right to enhance the potential of their children—by genetic means, among others? Finally, should we discover—as no doubt we will—the genetic switches that control our biological clock and that very likely influence also the maximum human life expectancy, will our life-prolonging culture opt to keep its hands off the process of aging and the upper limit on human life expectancy? Not a chance.

We thus face a paradox. On the one hand, genetic technology really is different, because it can and will go to work directly and deliberately on our basic, heritable, life-shaping capacities, at their biological roots, and it can take us beyond the existing norms of health and healing—perhaps even to alter fundamental features of human nature. On the other hand, we will find its promise familiar and irresistible, precisely because the goals it will serve, at least to begin with, will be continuous with those of modern high-interventionist medicine. This paradox itself contributes to public disquiet: We rightly perceive a powerful difference in genetic technology, but we also sense that we are powerless to use that recognized difference to establish clear limits to the use of genetic power. The genetic genie, first unbottled to treat disease, will, we rightly suspect, go his own way, whether we like it or not.

How Much Genetic Self-Knowledge Is Good for Us?

Quite apart from worries about genetic engineering, gaining genetic knowledge is itself a legitimate cause of anxiety, not least because one of its most touted benefits—genetic profiling of individuals—is guaranteed to increase everyone's anxiety. The deepest problem connected with knowing your own genotype and thus learning your own genetic sins and unhealthy predispositions is neither the threat to confidentiality and privacy nor the risk of so-called genetic discrimination in employment or insurance, important though these practical problems may be. It is rather the various hazards, anxieties, and deformations in living your life that attach to knowing in advance your likely or possible medical future. To be sure, such foreknowledge of predisposition will be welcome in

some cases, if it can lead to easy measures to prevent or treat the impending disorder and if we are talking about genes that predispose to disorders that do not powerfully affect self-image or self-command. But will and should we welcome knowledge that we carry a predisposition to Alzheimer's disease, schizophrenia, or some other personality or behavior disorder? That we definitely carry genes that will surely produce a serious but untreatable disease that will strike us at an unknown future time? Still harder will it be for most people to live easily and wisely with less certain information about predilections and predispositions, say, where multigenic traits are involved or where the predictions are purely statistical, with no clear implication for any particular "predisposed" individual. The recent case of a father who insisted that ovariectomy and mastectomy be performed on his 10-year-old daughter because she carried the BRCA-1 [breast cancer] gene dramatically shows the toxic effect of genetic knowledge.

Less dramatic but more profound is the threat that excessive genetic foreknowledge poses to human freedom and spontaneity, a subject explored 25 years ago by the late philosopher, Hans Jonas, one of our wisest commentators on technology and the human prospect. In a discussion of human cloning, Jonas argued for a novel "right to ignorance," necessary for human freedom and authentic action:

That there can be (and mostly is) too little knowledge has always been realized; that there can be too much of it stands suddenly before us in a blinding light. . . . The ethical command here entering the enlarged stage of our powers is: never to violate the right to that ignorance which is a condition for the possibility of authentic action; or: *to respect the right of each human life to find its own way and be a surprise to itself.* (Jonas, 1974, p. 163. Italics in original.)

To scientists who see only how knowledge of predispositions can lead to rational preventive medicine, Jonas' defense of ignorance will look like obscurantism. But, as Jonas observes,

knowledge of the future, especially one's own, has always been excepted [from the injunction to "Know thyself"] and the attempt to gain it by whatever means (astrology is one) disparaged—as futile superstition by the enlightened, but as sin by theologians; and in the latter case with reasons that are also philosophically sound. (Jonas, 1974, p. 161.)

Everyone remembers that Prometheus was the philanthropic god who gave to human beings fire and the arts, but we forget that he gave them also the greater gift of "blind hopes"—"to cease seeing doom before their eyes" (Aeschylus, lines

250ff)—precisely because he knew that ignorance of one's own future doom was indispensable to any human being's aspiration and achievement. I suspect that many people, taking their bearings from life lived open-endedly rather than from preventive medicine practiced rationally, will prefer ignorance of the future to the scientific astrology of knowing their genetic profile. In a free society, that will be their right. Or will it?

Freedom, Power, and Coercion

Even people who might welcome the growth of genetic knowledge and technology are worried about the power of geneticists, genetic engineers, and any governmental authority armed with genetic technology.² Precisely because we have been taught by these very scientists that genes hold the secret of life and that our genotype is our essence if not quite our destiny, we are made nervous by those whose expert knowledge and technique touch our very being. If, as science has taught us, power over genotype is power over life, not only ours but that of future generations, we have reason to be anxious, even apart from any particular abuses and misuses of that power. C. S. Lewis, friend neither of ignorance nor timidity, put the matter sharply:

It is, of course, a commonplace to complain that men have hitherto used badly, and against their fellows, the powers that science has given them. But . . . I am not speaking of particular corruptions and abuses which an increase of moral virtue would cure: I am considering what the thing called "Man's power over Nature" must always and essentially be. . . .

In reality, of course, if any one age really attains, by eugenics and scientific education, the power to make its descendants what it pleases, all men who live after it are the patients of that power. They are weaker, not stronger: for though we may have put wonderful machines in their hands we have pre-ordained how they are to use them. . . . The real picture is that of one dominant age . . . which resists all previous ages most successfully and dominates all subsequent ages most irresistibly, and thus is the real master of the human species. But even within this master generation (itself an infinitesimal minority of the species) the power will be exercised by a minority smaller still. Man's conquest of Nature, if the dreams of some scientific planners are realized, means the rule of a few hundreds of men over billions upon billions of men. . . . Each new power won *by* man is a power *over* man as well. (Lewis, 1965, p. 69–71. Italics in original.)

²One of the remarkable silences in all discussions of genetic technology has been the naive neglect of its potential usefulness in creating biological weapons, such as, to begin with, antibiotic-resistant plague bacteria or, later, aerosols containing cancer-inducing viral vectors.

Most of our genetic technologists, quite properly, will not recognize themselves in this portrait. Though they concede that abuses or misuses of power may occur, especially in tyrannical regimes, they see themselves not as predestinators but as facilitators, merely providing increased knowledge and technique that people can freely choose to use in making their health or reproductive decisions. Genetic power, they tell us, serves to increase freedom, not to limit it. But as we can see from the existing practices of genetic screening and prenatal diagnosis, this claim is at best self-deceptive, at worst disingenuous. The choice to develop and practice genetic screening and the choices of which genes to target for testing have been made not by the public but by scientists and not on liberty-enhancing but on eugenic (albeit, so far, on *negative* eugenic) grounds. Moreover, in many cases, practitioners of prenatal diagnosis refuse to do fetal genetic screening in the absence of a prior commitment from the pregnant woman to abort any afflicted fetus. And while a small portion of the population may be sufficiently educated to participate knowingly and freely in genetic decisions, most people are now, and no doubt always will be, subject to the (often but not always) benevolent tyranny of expertise. Every expert knows how easy it is to get most people to choose one way over another, simply by the way one raises the questions, describes the prognosis, and presents the options. The genetic preferences of scientists and counselors will always overtly or subtly shape the choices of the counseled.

In addition, economic pressures to contain health-care costs will almost certainly constrain free choice. Discrimination in insurance may eventually work to compel genetic abortion or genetic intervention, through decisions to refuse coverage for this or that genetic disease. State-mandated screening already occurs for PKU. In France, the government has mandated that all citizens will need to carry all their personal information on a "smart card" by the year 2000. The growing tendencies to rationalize health care and to make it more cost-effective may constrain choice precisely as they enhance prospects for prevention and treatment. Moreover, with full-blown genetic screening, there will likely be increasing pressure to limit reproductive freedom, all in the name of the well-being of children. Already, in 1971, geneticist Bentley Glass, in his presidential address to the American Association for the Advancement of Science, enunciated "the right of every child to be born with a sound physical and mental constitution, based on a sound genotype." Looking ahead to the reproductive and genetic technologies that are today rapidly arriving, Glass proclaimed: "No parents will in that future time have a right to burden society with a malformed or a mentally incompetent child." (Glass, 1971, p. 28.) It remains to be seen to what extent such prophecies will be realized, but they surely provide sufficient

and reasonable grounds for being concerned about restrictions on human freedom, even in the absence of overt coercion, even in liberal polities.

Beyond Freedom and Coercion: Questions of Dignity and Dehumanization

Although the public worries about abuses of genetic power and about who will control the controllers, I believe its deepest concerns lie elsewhere. What does and should worry us most can, and probably will, arise even with the free, humane, and so-called enlightened use of these technologies. For, truth to tell, genetic technology, the practices it will engender, and (above all) the scientific teachings about human life on which it rests and which it seems to validate are not, as many would have it, simply morally and humanly neutral. They are pregnant with their own moral meaning, regardless of whether they are practiced humanely or taught humbly. They necessarily bring with themselves changes in our practices, institutions, norms, beliefs, and human self-conception. It is these challenges to our dignity and humanity that most urgently generate the concerns over genetic (and other biomedical and neuropsychological) science and technology. Let me touch on five aspects of this most serious matter.

"Playing God"

Curiously, the worry about dehumanization is sometimes expressed, paradoxically, in the fear of superhumanization, that is, that man, or rather *some* men, will be "playing God." This complaint is too facilely dismissed by scientists and others who are nonbelievers. The concern has meaning, God or no God. By this phrase is meant one or more of the following: (1) Man, or, again, some men, are becoming creators of life, and indeed, of individual living human beings (in vitro fertilization, cloning); (2) they not only create life, but they stand in judgment of each being's worthiness to live or die—not on moral grounds, as is said of God's judgment, but on somatic and genetic ones (genetic screening and abortion); and (3) they also hold out the promise of salvation from our genetic sins and defects (gene therapy and genetic engineering). Man, not God, is a god to man.

Never mind the exaggeration in the conceit and the fact that man, even at his most powerful, is capable only of playing at being God. Consider only that, if scientists are seen in the godlike role of creator-judge-savior, the rest of us must stand in inferior relation to them as creatures-judged-tainted. These worries, despite the hyperbolic speech, are not far-fetched.

One example will suffice. Not long ago, in my own institution, a physician making rounds with medical students stood over the bed of an intelligent, otherwise normal ten-year-old boy with spina bifida. "Were he to have been conceived today," the physician casually informed his entourage, "he would have been aborted." Determining who shall live and who shall die—on the basis of genetic merit—is a godlike power already wielded by genetic medicine. And this power will only grow.

Manufacture and Commodification

But, one will rightly respond, genetic technology holds out the promise of redemption, of the cure for these life-crippling and life-forfeiting disorders. Very well. But to truly practice their salvific power, the genetic technologists will have to greatly increase their manipulations and interventions, well beyond merely screening and weeding out. True, genetic testing and risk management aimed at prevention may in some cases actually cut down on the need for high-tech interventions aimed at cure. But there will be many, many other cases in which increasing scrutiny will necessarily be accompanied by increasing manipulation. And, to produce Bentley Glass's healthy and well-endowed babies, let alone babies with the benefits of genetic enhancement, a new scientific obstetrics will be necessary, one that will come very close to turning human procreation into manufacture. This process has already crudely begun with in vitro fertilization; it will soon take giant steps forward with the ability to screen the in vitro embryos before implantation; with cloning; and, eventually, with precise genetic engineering. Just follow the logic and the aspirations of current practice: The road we are traveling leads all the way to Brave New World—not by dictatorial fiat, but by the march of benevolent humanitarianism, cheered on and enjoyed by the very citizens who, in their ambivalence, also dread becoming simply the latest of man's manmade things.

Make no mistake. The price to be paid for producing optimum, or even only genetically sound, babies is the transfer of procreation from the home to the laboratory and its coincident transformation into manufacture. Increasing control over the product can only be purchased by the increasing depersonalization of the process. More and more, we will give existence to new life not by what we *are* but by what we *intend and design*. As with any product of our making, no matter how excellent, the artificer will stand above it, not as an equal but as a superior, transcending it by his will and creative powers. Such an arrangement will be profoundly dehumanizing, no matter how genetically good or healthy the children. And let us not forget the powerful economic interests

that will surely operate in this area; with their advent, the commodification of nascent human life will be unstoppable.

Standards, Norms, and Goals

Equally troublesome is the matter of standards, norms, and goals. According to Genesis, God, in His creating, looked at His creatures and *saw* that they were *good*—intact, complete, well-working wholes, true to the spoken idea that guided their creation. But what standards will guide the genetic engineers? For the time being, one might answer, the guide would be the norm of health. But even before the genetic enhancers join the party, the standard of health is going to be deconstructed. Are you healthy if you are asymptomatic but carry genes that will definitely produce Huntington's disease or that predispose to diabetes, breast cancer, or coronary artery disease? What if you carry, say, 40 percent of the genetic markers thought to be linked to the appearance of Alzheimer's disease? And what will health and normality mean when we discover genetic propensities to alcoholism, drug abuse, pederasty, or violent behavior? Health will become at once both imperial and vague: Ironically, we will get increased medicalization—via genetic diagnosis—of what have hitherto been mental or moral matters at the same time that we will see the disappearance of any given standard of health, wholeness, or fitness.

Once genetic enhancement comes on the scene, all pretense of standards will go out the window, just when such standards would be most urgently needed. "Enhancement" is, of course, a soft euphemism for improvement, and the idea of improvement necessarily implies a good, a better, and perhaps even a best. But if previously unalterable human nature can no longer function as a standard or norm for what is regarded as good or better, how will anyone truly know what constitutes an improvement? It will not do to say that we can extrapolate from what we like about ourselves and to proclaim that more is better. Because memory is good, can we say how much more memory would be better? If sexual desire is good, how much more would be better? Given that life is good, how much extension of the maximum life expectancy would be good for us? Only simplistic thinking believes it can easily answer such questions. In whose image will the creators of the new and enhanced human beings create them? This is the real problem with positive eugenics: less the threat of coercion, more the presumption of thinking we are wise enough to engineer "improvements" in the human species.

The more modest enhancers, like the more modest genetic therapists and technologists, have no such grandiose goals. They are valetudinarians, not

eugenicists. They pursue, or think they pursue, not some far away positive good, but the positive elimination of evils: diseases, pain and suffering, the likelihood of death. But let us not deceive ourselves. There is in all this avoidance of evil an implicit positive goal: nothing less than a painless; suffering-free; and, finally, immortal existence. What is more, though unstated, this implicit goal is in fact held to be uncontroversial and paramount. Only the presence of such a goal can justify sweeping aside all opposition to the progress of medical science. Only such a goal gives to the principle "cure disease, relieve suffering" its trumping value in nearly all arguments about medical ethics: "Cloning human beings is unethical and dehumanizing, you say? So what: it will help us treat infertility, avoid genetic disease, and provide perfect materials for organ transplantation."³ Never mind whether it means creating and growing human embryos for experimentation, changing the definition of death to facilitate organ transplantation, growing human body parts in the peritoneal cavities of animals, perfusing newly dead bodies as factories for useful biologicals, or reprogramming the human body and mind with genetic or neurobiological engineering: Who can sustain an objection if these practices help us live longer and with less overt suffering?

The Tragedy of Success

That the project is utopian and finally doomed to failure does not slow the enthusiasts. They do not see that we will not eliminate suffering but merely shift it around. They do not remember that contentment means parity between one's desires and one's powers, and they therefore do not appreciate the discontent that we are already seeing as a result of rising desires and expectations in the health-care field.⁴ Worst of all, they do not see the larger human cost of the successes of the humanitarian project. As Aldous Huxley made clear in his prophetic *Brave New World*, the conquest of disease, aggression, pain, anxiety, suffering, and grief unavoidably comes at the price of homogenization, mediocrity, pacification, drug-induced contentment, trivialized human attachments, debasement of taste, and souls without loves or longings—the

³Such was the tenor of *Cloning Human Beings*, the June 1997 report of the National Bioethics Advisory Commission, notwithstanding its call for a temporary ban on human cloning. The only agreed-upon (and temporary) moral objection to human cloning: It "is not safe to use in humans at this time," solely because the technique has yet to be perfected (p. iii). Even this elite ethical body apparently believes that there are no other moral arguments sufficient to cause us to forgo possible health benefits.

⁴A number of recent studies show that, although their actual state of health has improved substantially in recent decades, people's *satisfaction* with their current health status has remained the same or even declined. People seem to be doing better but feeling worse, very likely as a consequence of rising expectations.

inevitable result of making the essence of human nature the final object of the conquest of nature for the relief of man's estate. Like Midas, bioengineered man will be cursed to acquire precisely what he wished for, only to discover—painfully and too late—that what he wished for is not exactly what he wanted. Or, worse than Midas, he may be so dehumanized he will not even recognize that, in aspiring to be perfect and divine, he is no longer even truly human. To paraphrase Bertrand Russell, technological humanitarianism is like a warm bath that heats up so imperceptibly you do not know when to scream.

I am sorry to paint such a gloomy prospect. I surely have no way of knowing whether my worst fears will be realized, but you surely have no way of knowing that they will not. True, Huxley's portrait is science fiction, but what was debunked as mere science fiction not 20 years ago is today genuine biological possibility. But my main point is not the rightness or wrongness of this or that imagined scenario—all this is admittedly highly speculative. It is rather the plausibility, or even the wisdom, of thinking about genetic technology, like the entire technological venture, under the very ancient, profound, yet profoundly un-American idea of tragedy, that poignantly human adventure of living in grand self-contradiction. In tragedy, the failure is embedded in the hero's success, the defeats in his victories, the miseries in his glory. The technological way of approaching both the world and human life, a way deeply rooted in the human soul and spurred on by the utopian promises of modern thought and its scientific crusaders, seems to be inevitable, heroic, and doomed.

Science, the Soul, and Shrunk Self-Understanding

To say that technology as a way of life is doomed, left to itself, does not yet mean that modern life—*our* life—*must* be tragic. Everything depends on whether the technological disposition is allowed to proceed to its self-augmenting limits, or whether it can be restricted and brought under intellectual, spiritual, moral, and political rule. And here, I regret to say, the news is not encouraging. For the relevant intellectual, spiritual, and moral resources of our society, the legacy of civilizing traditions painfully acquired and long preserved, are taking a beating, not least because they are being called into question by the findings of modern science itself and by biology's most public and prophetic voices. The technologies present troublesome ethical dilemmas, but the underlying scientific notions call into question the very foundations of our ethics.

The challenge goes much further than the notorious case of evolution versus biblical religion. Is there *any* elevated view of human life and goodness that is proof against the belief that man is just a collection of molecules, an accident on

the stage of evolution, a freakish speck of mind in a mindless universe, fundamentally no different from other living—or even nonliving—things? What chance have our treasured ideas of freedom and dignity against the teachings of biological determinism in behavior, the reductive notion of “the selfish gene” (or, for that matter, of “genes for altruism”), the belief that DNA is the essence of life, and the credo that survival and reproductive success are the only natural concerns of living beings—or, rather, of their genes?

As sociologist Howard Kaye notes:

For over forty years, we have been living in the midst of a biological and cultural revolution of which innovations such as AID, in vitro fertilization, surrogacy, genetic manipulation, and cloning are merely technological offshoots. In both aim and impact, the end of this revolution is a fundamental transformation in how we conceive of ourselves as human beings and how we understand the nature and purpose of human life rightly lived. . . . Encouraged by bio-prophets like Francis Crick, Jacques Monod, E.O. Wilson, and Richard Dawkins, as well as by humanists and social scientists trumpeting the essential claims of race, gender, and ethnicity, we are in the process of redefining ourselves as biological, rather than cultural and moral beings. Bombarded with white-coated claims that “Genes-R-Us,” grateful for the absolution which such claims offer for our shortcomings and sins, and attracted to the promise of using efficient, technological means to fulfill our aspirations, rather than the notoriously unreliable moral or political ones, the idea that we are essentially self-replicating machines, built by the evolutionary process, designed for survival and reproduction, and run by our genes continues to gain. But still the public’s ambivalence persists, experienced in the form of anxiety at what such a transformation would mean. (Kaye, 1997b; see also Kaye, 1997a.)

These transformations are, in fact, welcomed by many of our leading scientists and intellectuals. Last year the luminaries of the International Academy of Humanism—including biologists Crick, Dawkins, and Wilson and humanists Isaiah Berlin, W. V. Quine, and Kurt Vonnegut—issued a statement in defense of cloning research in higher mammals and humans beings. Their reasons are revealing:

What moral issues would human cloning raise? Some world religions teach that human beings are fundamentally different from other mammals Human nature is held to be unique and sacred. Scientific advances which pose a perceived risk of altering this “nature” are angrily opposed. . . . [But] [a]s far as the scientific enterprise can determine . . . [h]uman capabilities appear to differ in degree, not in kind, from those found among the higher animals. Humanity’s rich repertoire of thoughts, feelings, aspirations, and

hopes seems to arise from electrochemical brain processes, not from an immaterial soul that operates in ways no instrument can discover. . . . Views of human nature rooted in humanity's tribal past ought not to be our primary criterion for making moral decisions about cloning. . . . The potential benefits of cloning may be so immense that it would be a tragedy if ancient theological scruples should lead to a Luddite rejection of cloning. (International Academy of Humanism, 1997.)

To justify ongoing research, these intellectuals are willing to shed not only traditional religious views, but all views of human distinctiveness and special dignity, their own included. They are seemingly unaware that the scientific view of man they celebrate does more than insult our vanity. It undermines our self-conception as free, thoughtful, and responsible beings, worthy of respect because we alone among the animals have minds, hearts, and aspirations that aim far higher than mere life and the perpetuation of our genes. It undermines the beliefs that hold up our mores, practices, and institutions, not excluding science itself. Why, on these intellectuals' understanding of "the rich repertoire" of human thought, should anyone choose to accept as true the results of *their* "electrochemical brain processes" rather than adhere to those of his own?

The problem may lie not so much with the scientific findings themselves but with the shallow philosophy that recognizes no other truths but these and with the arrogant pronouncements of the bioprophets. In a recent letter to the editor complaining about a review of his book, *How the Mind Works*, evolutionary psychologist and popularizer Stephen Pinker rails against any appeal to the human soul:

Unfortunately for that theory, brain science has shown that the mind is what the brain does. The supposedly immaterial soul can be bisected with a knife, altered by chemicals, turned on or off by electricity, and extinguished by a sharp blow or a lack of oxygen. Centuries ago it was unwise to ground morality on the dogma that the earth sat at the center of the universe. It is just as unwise today to ground it on dogmas about souls endowed by God. (Pinker, 1998.)

One hardly knows whether to be more impressed with the height of Pinker's arrogance or with the depth of his shallowness. But he speaks with the authority of science, and who can dispute him on his own ground?⁵

⁵For an attempt to dispute such reductionist claims and to point the way to a more adequate account of living nature (on philosophical, not religious grounds), see Kass (1985) and Kass (2nd ed., 1999). See also Jonas (1982).

There is, in fact, nothing novel about reductionism, materialism, and determinism; these are doctrines with which Socrates contended. What is new is that these philosophies now seem to be vindicated by scientific advance. Here, in consequence, is perhaps the most pernicious result of our technological progress—more dehumanizing than any actual manipulation or technique, present or future. We are witnessing the erosion, perhaps the final erosion, of the idea of man as something noble, dignified, precious, or godlike, and its replacement with a view that sees man, no less than nature, simply as more raw material for manipulation and homogenization.

Hence, our peculiar moral crisis. We are in turbulent seas without a landmark precisely because we adhere more and more to a view of human life that both gives us enormous power and, at the same time, denies all possibility of nonarbitrary standards to guide its use. Although well-equipped, we know not who we are or where we are going. We triumph over nature's unpredictabilities only to subject ourselves, tragically, to the still greater unpredictability of our capricious wills and our fickle opinions. Engineering the engineer, as well as the engine, we race our train we know not where.

This, I submit, is the truest moral meaning of all of today's wonderful biomedical technology and of the scientific view it reflects and fosters. It is only our infatuation with scientific progress and our naive faith in the sufficiency of our benevolently humanitarian impulses that prevent us from recognizing it.

Does this mean, therefore, that I am in favor of ignorance, suffering, and death? Am I in favor of killing the goose of genetic technology even before she lays her golden eggs? Surely not. But I do insist on the importance of seeing the *full* human meaning of this new enterprise in biogenetic technology and engineering. Important though it is to set a moral boundary here or devise a regulation there, hoping to decrease the damage caused by this or that little rivulet in the belief that one is avoiding the torrent, it is even more important to be sober about the true nature and meaning of the flood itself. The new biologists and their technological minions do not know all they think they know, and they never will. For all their ingenuity, they do not even seek the wisdom that just might yield the kind of knowledge that keeps human life human. If, unlikely though it seems, they could be persuaded to face squarely the full import of the project they are launching, they might proceed with less heedless exuberance and greater humility. And if the rest of us become clearly aware of the dangers—not just to privacy or insurability but to our very humanity—we might be better equipped to defend the increasingly beleaguered pockets and principles of human dignity, even as we continue to reap the considerable benefits genetic technology will inevitably provide.

References

Aeschylus, *Prometheus Bound*, lines 250ff.

Glass, Bentley, "Science: Endless Horizons or Golden Age?" *Science*, Vol. 171, 1971, pp. 23-29.

International Academy of Humanism, "Statement in Defense of Cloning and the Integrity of Scientific Research," May 16, 1997.

Jonas, Hans, *Philosophical Essays: From Ancient Creed to Technological Man*, Englewood Cliffs, NJ: Prentice Hall, 1974.

_____, *The Phenomenon of Life: Toward a Philosophical Biology*, Chicago: University of Chicago Press, 1982.

Kass, Leon R., *Toward a More Natural Science: Biology and Human Affairs*, New York: The Free Press, 1985.

_____, *The Hungry Soul: Eating and the Perfecting of Our Nature*, 2nd. ed., Chicago: University of Chicago Press, 1999.

Kaye, Howard, *The Social Meaning of Modern Biology*, 2nd ed., New Brunswick, NJ: Transaction Publishers, 1997a.

_____, "Anxiety and Genetic Manipulation: A Sociological View," paper presented at the Hastings Center, December 1997b.

Lewis, C. S., *The Abolition of Man*, New York: Macmillan, 1965.

National Bioethics Advisory Commission, *Cloning Human Beings*, June 1997

Pinker, Steven, "A Matter of Soul," *The Weekly Standard*, Correspondence Section, February 2, 1998, p. 6.

Walters, LeRoy, "Human Gene Therapy: Ethics and Public Policy," *Human Gene Therapy*, Vol. 2, 1991, pp. 115-122.

G. The Conversion of Genetics and Computing: Implications for Medicine, Society, and Individual Identity

Leader: Dr. George Poste
Vice President Smith Kline Beechum
Chief Scientist—R&D
April 19, 1999

Summary by Danilo Pelletiere, George Mason University

In his talk to the group, George Poste discussed what he anticipates will be the implications for human medicine from the "two formidable technological imperatives" of advances in genetics and progress in computing. Poste also spent some time discussing the economic and policy implications of the growing linkages between the two.

The theme of his presentation was that advances in these technologies hold the potential to move medicine away from being a descriptive discipline to becoming a mechanistic discipline: Used in tandem, advances in computing and genetics will allow medicine to move further beyond the mere description of diseases to elucidation of the fundamental basis for pathology and will allow this information to become part everyday medical practice. Understanding the genetic basis for pathology and having the information technologies to hold, transfer, and analyze that information should in turn lead to more rational mechanisms for the diagnosis and, particularly, the treatment of disease. The ultimate goal of this progress, according to Poste, will be the prediction and prevention of disease.

The increase in genetic knowledge and our ability to handle and analyze it is also changing the way we treat disease. Today, researchers are finding a bewildering degree of heterogeneity and quite distinct genetic pathologies underlying diseases that are symptomatically quite similar and that, until only recently, were thought to be one and the same. Following from this, it seems likely that the ineffectiveness of our therapeutic interventions in some cases may be due to our failure to understand such variations in the underlying pathology of the disease. Each of these distinct variations may require a different treatment modality.

While discovering these underlying differences will be a great step forward for therapy, such discoveries will also likely have a great effect on the economics of the pharmaceutical industry. With the current average cost to bring a new drug to market at around \$500 to 600 million, the question quickly arises, how far can you fragment a major disease set that is currently treated by a single drug or set of drugs and still find it compatible with the economics of drug discovery and development? The implicit costs of the biological and information technology were another theme in his talk. What Poste made clear throughout was that the potential of the current biological and information technology revolutions in health care is subject to real constraints and uncertainty when considered through the industry's economic lens. In the case of fragmenting diseases and therefore therapies, the benefits to firms are likely to be predicated on the emergence of a parallel process that radically reduces discovery and development costs. Without such a parallel process, Poste argued, costs will rise as the market size for each new drug—designed to target only a subset of disease—diminishes. Later in his presentation, Poste suggested that the ability of physicians and the health industry to use the huge genetic databases that are being compiled today depends on massive investment in information technology infrastructure, which he does not see occurring today. This may ultimately lead to the merger of pharmaceutical and telecommunications firms in the long run, and is likely to be a major obstacle to the application of new knowledge in the short run.

The other likely development that Poste highlighted is that not only will diseases with similar symptoms be treated as increasingly distinct ailments, but patients too will no longer be treated as uniform. It will become possible for physicians and patients to be much more informed about individual reactions to therapy based on an individuals' genetics and even their environments.

The real issue, however, that lies at the heart of all these advances will be the testing of individuals for predisposition to disease. Poste was quick to state that such technology would only become generally available in 25 years or so because of technical and economic obstacles yet to be overcome. The primary obstacle is that, while some diseases can be classified as unigenic, most are multigenic diseases. For example, 17 separate genes have thus far been implicated in the development of diabetes. As he discussed later, given the influence of further variation within genes, the combinatorics alone associated with sorting out the genetic causes of the disease and whether it will occur in any single individual remain technically daunting. Furthermore, many of the necessary data and much of the infrastructure are still missing. For the complex multigenic diseases, such as cancers, cardiovascular, and neurodegenerative diseases—all of the

diseases that impose the greatest economic toll on society—we are still technically a long way from individual profiling. He stressed, however, that, although the industry is not as far along in this respect as the headline writers might suggest, it is still a valid theme that deserves both technical and ethical consideration today.

After laying out the speculative nature of his general themes, Poste went on to highlight two “rude and blunt” challenges that are almost certain to test current medical systems severely in the next five to ten years. These emerging situations will painfully reflect the shortcomings of current medicine delivery, and can not be ignored in any discussion of the future of human medicine. First, there is now ample evidence that large-scale antibiotic-resistant epidemics are likely in the near future. Poste estimated that, between 2004 and 2010, western societies will experience a vulnerability to disease similar to that before 1945. The implications of this emerging situation are compounded by what Poste referred to as a population increasingly “cocooned” from most of the risk and dangers of disease. Furthermore, according to Poste, governments are not giving the issue adequate public attention. It is important to recognize that the last new class of vaccines was discovered in the 1970s.

The second situation, of which we are more aware, is an aging population that represents an ever-growing demand for what will likely be increasingly finite resources. We began this century with an average life span of 40, and we are ending it now almost twice that, thanks to public sanitation, improved nutrition, pharmaceuticals, and vaccines. This issue is receiving some attention, but according to Poste, the true extent of its effects is largely being ducked by politicians, whether in a private system, as in the United States, or a public system, as in Great Britain.

Given these two further and more immediate imperatives, it is necessary to look at how the parallel revolutions in genetics and information technology will be able to improve the provision of antibiotics, vaccines, and health care in general. In this regard, Poste suggested that the greatest and most immediate benefit could be gained from using these new technologies in the treatment of existing disease. With increased availability and analysis of genetic information, misdiagnosis could be radically reduced. Not only would this reduce injurious and even fatal mistakes in the prescription of pharmaceutical therapies, but it would also reduce the use of new and existing antibiotics, perhaps increasing their effective life.

Better understanding of the genetic basis for pathologies is likely to lead to a dramatic increase in specifically targeted genetic therapies for most of the major

diseases. The real power of genetics in the short term, however, will be in diagnostics. Here, Poste was careful to draw a distinction between "genetic testing" for predisposition to disease and testing to diagnose existing disease. These are technically and ethically very different. In the near term, the second is more likely and more important. The improved diagnosis of existing disease will improve the way we treat disease and our ability to define disease earlier. Increasingly, it is becoming clear that distinct diseases with similar symptoms are being misdiagnosed and ineffective therapies are the result. Poste used the example of acute lymphocytic leukemia. Five years ago it was considered one disease, and we would have we would have treated patients having it by assaulting them with various cytotoxic "napalms," as Poste referred to the chemicals used in the treatment of the disease. While some patients would respond, most would not. Five years later, we now know there are 14 different types of acute lymphocytic leukemia based on highly reproducible genopathologies that the potential victims carry. Therefore, it is now understood that effective treatment of the disease requires using different therapies according to the distinct genetic type of disease present in the patient. As discoveries like this are made, the effectiveness (and economic efficiency) of treatment should greatly improve.

In this regard and in the short term, the most dramatic gains will therefore be in oncology and infectious disease. At a SmithKline Beechum division in California, Poste set the goal that, within ten years, a person who has a yearly physical should never die of a metastatic malignancy. The four biggest malignancies for humans are breast, lung, prostate, and colon. By diagnosing the disease before metastaticization has begun and by starting treatment, it should be possible to save lives from cancer without finding the cure to cancer. Poste feels that doing this within ten years is a realistic goal. Again returning to the profound economic and social changes that these technological revolutions portend, however, he referred to the implications for the Social Security system and the costs of health care of this treatment and the survival of more patients. The trade-offs here are as always uncertain.

In keeping with the theme of costs, he also suggested that, now that we understand the heterogeneity of disease and the economic and health costs (discussed further below) of misdiagnosis, we will no longer have the luxury of waiting a few days to find out what exactly is affecting people. Therefore, we must develop highly automated detection systems, not only to identify the disease in terms of its family of origin but also to profile fully its antibiotic resistance and susceptibility spectrum. To be able to provide a real benefit, this should occur within 30 minutes, since most patients will have to leave the

doctor's office with some sort of treatment in hand. The economic cost of unnecessary therapy is a great burden for today's health-care system. In the future, decision-support software will not allow the physician to prescribe expensive therapy unless the patient has a validated subset of the disease associated with that treatment. Thus, the physician's prescription will be verified against the individual's and the disease's genetic and environmental information, to check that an optimal and safe treatment is being prescribed. In theory, this should both save money and increase the useful life of antibiotics. Many HMOs are already beginning to adopt this type of approach.

Moving on to predispositional diagnostics, Poste made the claim that this is the most powerful technology that could ever be unleashed. It would allow to use a genetic profile to assess a person's risk of major diseases. This is a capability we do not have in any real sense today. For this to occur, however, a number of technical and economic components need to be in place first.

One of these components is the recently announced Single Nucleotide Polymorphism Consortium (SNIP). Ten major pharmaceutical companies are working together to examine the single nucleotide polymorphisms in genes and the implications of these polymorphisms for an individual's susceptibility or resistance to disease. It is not just the presence of particular genes that causes disease but also distinct variations and mutations within genes.

Although it now seems feasible to create such a SNIP map, undertaking predispositional analyses will still be a huge logistical problem. Even if, three to five years from now, we have the SNIP map of man and if it can be superimposed on the full genomic code of man (the Human Genome Project), the challenges from the combinatorial associations alone will make the task of mapping risk to major diseases quite formidable. As discussed above, 17 distinct genes that have been implicated to date in the development of diabetes alone. Each of these genes, however, can be heterozygous, which means that 3^{17} , or 129 million, genetic combinations may (or may not) be implicated in the disease. These combinations are then subject to additional variation due to polymorphisms (the SNIP map), which must also be considered.

This brings us to an important third component that must be in place for predispositional testing to take place. Once these vast databases are in place, the computing power must be available to trawl, transport, and analyze the data and results. The future of medicine is data heavy and highly dependent on continued advances in computing power and information technology networks.

Therefore, three technical achievements must be in place before the medical potential of these current advances can be realized:

1. the full genome of man
2. the full SNIP man
3. cost-effective genotyping software and hardware.

The last condition suggests that the industry will need the equivalent of a Moore's law to drive down the unit cost of a processing unit of genetic information. Today, using current technologies, even if we had the full SNIP and human genome maps, even rudimentary analysis would cost several tens of thousands of dollars per patient. It is not immediately apparent that there will be an exponential improvement in the development of the necessary software (that is, algorithms to analyze genetic data) and hardware similar to what has been occurring in information technology. Still, Poste expects these technical domains to be in place within the next ten to fifteen years. This will begin to put in place the building blocks for predispositional risk profiling, and then (if not sooner) we will have to face the dauntingly complex bioethical issues.

The primary bioethical issue is discrimination. The dangers, according to Poste, are that probabilistic outcomes will receive the wrong interpretational overlay from policymakers and the public. This is the problem of interpreting nature versus nurture. In the 1960s, one could hardly mention nature; now, the pendulum has almost swung back to where genetic "nature" is all. This ignores the complexities of gene-environment and gene-gene interactions that are little understood and will continue to be little understood for quite some time.

Beyond the problems of discrimination and the robustness of the evidentiary standards used to define risk, this new technology will most certainly inflame both the abortion and eugenics debates. If a fetus is shown to be carrying "genes for"—a term that Poste views as a dangerous misrepresentation of the level of determinism implicit in such a finding—the opportunity and pressure for eugenics based on economic or other considerations will appear. The possibility of defining neurogenetics—a genetic basis for behaviors—further suggests that there may be pressure either to treat pharmacologically or to weed out behaviors that are judged somehow to be aberrant. The implications of this for society and for the genetics of man are not well understood.

To show how quickly these issues will emerge, however, Poste used an example from a few years ago, "Dutch Violence Cohort Gene." This grew out of a study of a family in the Netherlands whose members showed a tendency to become suddenly and explosively violent. The genetic analysis of this pedigree showed a correlation with one particular variant of a gene. This was corroborated by a study of the Old Order Amish in Lancaster County, PA. The same variant of the

gene was found in the family members who were prone to explosive outbursts after periods of relative calm. Within a week after the papers reporting this were published, SmithKline Beechum received an invitation from a state correctional system asking whether the company was interested in typing their inmates. Poste reported the company refused.¹ For him, this was an example of how simplistic interpretation of the results of this sort of predispositional risk profiling can be adopted by society and lead to improper or misunderstood policy pressure.

Even if were possible to assign highly robust correlations between genes and phenotypic fate, the problem arises of determining whether we want to know and when we want to know what diseases we are likely to have. In this case, a patient's decision would likely be influenced by whether the risk was mitigatable and whether therapy is available.

Poste returned to his discussion of predispositional testing in the field of pharmacogenetics—understanding how an individual's genetic makeup influences his or her responses to current and future drugs—and its likely contribution to health care. Drug action is evaluated by two criteria, efficacy and safety. With the current broad drug classes available today, between 50 and 70 percent of a population is likely to show an efficacy response roughly along the lines the therapy provider expects. Somewhere between 25 and 30 percent of the population will be completely refractory and will show no or even adverse effects. With new pharmacogenetic technologies, you might have an *a priori* classification of the molecular foundation of responsiveness versus nonresponsiveness to pharmaceutical therapies. That will add an element of rationality and safety to medicine. Today, 3.1 billion prescriptions are issued in the U.S. every year, of which 2.1 million result in an adverse reaction. Of the prescriptions with adverse reactions, 1 million will result in hospitalization and 106,000 in death. This makes adverse drug reactions the fourth biggest killer in the United States, and this is due to genetic variation among individuals in drug metabolism enzymes. A number of companies, including SmithKline Beechum, are moving aggressively to create diagnostic assemblies that use all human drug metabolism enzymes and all known genetic variance to screen drugs based on the genetics of an individual.

The potential benefits of such a screening system are vast, but as discussed above, this technology is unlikely to appear in the majority of doctors' offices

¹Later, in the question period, Leon Kass, an ethicist at the University of Chicago (and a previous speaker) asked the reason for this refusal. Was it based on the nature of the evidence or the ethical dilemmas of such a program? Poste responded that it was the lack of a robust evidentiary standard of correlation.

unless there is parallel progress in building the information-science capability to handle large-scale human genetic databases cheaply. The need for economic investment in this area will necessitate entirely new business arrangements. The evolution of large-scale health information systems may well require the merger of health-care, telecommunications, and software companies.

Once again, however, even if the necessary infrastructure and the software capability to handle and analyze the emerging genetic information are in place, the policy and ethics questions will remain formidable challenges. What of informed consent? What about "retrospective trawling" of existing genetic databases? Today, samples are tested for many things through "multiplex testing" that cannot be analyzed with our current technology and understanding. Current laws in Europe, the European Data Directive in particular, that protect this information are likely to become an obstacle to the transfer and use of these data for research and treatment. The European Data Directive rejected safe-haven exemption for medical data in the name of privacy. Therefore, the ethical issues are unavoidable. Encryption constraints and cultural differences are likely to define the environment in which this process proceeds.

The direction of the genetic revolution and its interface with information technology depend further on two factors. The first of these is the direction that research on human embryonic stem cells takes. In theory, this research opens up the prospect of regenerative medicine through the apparent potential of these cells to differentiate into any type of tissue. In the short term, this is going to be a flash point (as evidenced by a recently released ethics report cited in *The Washington Post*, May 23, 1999, p. A1.) for research in the United States because of the current necessity of using embryos to carry out this research. Linked to this is the second factor, reprogramming of the human nucleus, "Dolly-type" technology. We already knew that an adult cell nucleus has every single gene necessary to code for a complete adult organism. What Dolly taught us was that some set of chemical triggers in the egg are capable of totally reprogramming the genetic repertoire of an adult nucleus. In this way, a completely new organism can be created.

If regenerative tissue production is to become routine, it will become necessary to link these two technologies. We need a nucleus from the patient if we are going to create organs that the patient will not reject. The only potential escape from the embryonic complexity is if, at some point, we actually learn to understand the signals that confer plasticity and reprogrammability. This will allow to use the patient's own cells. In the short term, however, we will have to use embryos if this research is to continue.

This vision of the future of medicine suggests that, sometime in the next 25 years, the genetic manipulation of humans will occur. Whether this will take the form of simple somatic lineage or of direct manipulation of the gametes and the germ line to create heritable changes will vary by society and culture. There will be therapeutic strategies that are no longer connected to exogenous chemicals but instead to the ability to switch genes on and off selectively for therapeutic purposes.

But the really fascinating thing will be the actual merging of carbon and silicone, the implantation of artificial devices, as we have never seen before. In the simplest form, it seems possible to implant a subcutaneous chip that will carry all of a person's medical and personal information. Poste made the point that, in France, citizens are already required to carry a smart card with this information. As health care becomes increasingly information dependent, having complex medical information on (or in) the patient or the codes to access this information will be essential for good medical care and, particularly, for the economics of the system. This information will allow health-care providers to avoid duplicative tests and to provide immediate attention. This is clearly dependent on seamless data flows and raises the specter of unequal access to this technology and the inability of the system to treat those without such technology.

But beyond the simple implantation of information to be accessed by external sources, there is the issue of true cyborgian technology. Of particular interest is the ability to implant intelligence or knowledge externally. Here, the dichotomy to consider in the merging of carbon and silicone systems is between (1) augmenting intelligence, that is, how we use computational systems to create knowledge set of use to us in an increasingly facile way, and (2) imbedding intelligence. A number of advances have been made in implanting microdevices linked to neurosets, creating direct computational interfaces that can, for example, move a cursor on a computer screen.

What are the limits of biomimetic reverse engineering as neurological and genetic information continues to be mapped? Poste was recently exposed to work that used miniaturized electrode technology to look at spatial firing patterns to determine how rats find their way in a maze. What the researchers found was that these patterns operate exactly the same as visual pathways. This leads to the question of whether it is possible to transfer that intelligence to another individual without its having the actual experience. Will it be possible to arrive in a new city and know exactly where you are? It begins to look as though it will be possible.

Similarly, it is possible to monitor traffic in the optic nerve because visual images are transferred as electric, and therefore digital, pulses. If it is possible to digitize what is being seen, it seems possible to digitize an image one person sees and transfer it to those who have never seen it. Such cyborgian interventions are based on processes having similarities between the cognate signaling of biological systems and the digital signaling of totally artificial systems.

Poste summed up by stating that these emerging technologies place a whole new set of dichotomies, polarizations, antagonisms, and confrontations in the public-policy arena. The primary policy question is likely to continue to be access to care within this country and across the globe. But even beyond access to care, these new technologies will offer us remarkable alternatives that will force continual reassessment of our notions of nature and nurture. There is a real danger that the information and capabilities unleashed by new medical technologies may be interpreted in a deterministic vein and that free will become only the latest wiring diagram that offers the greatest flexibility. We will also continually and increasingly have to confront society's belief in the sanctity of the embryo and the fetus. And the eternal confrontation between beneficial and malignant manipulation of these new technologies will have to be addressed. The threats biological warfare and terrorism present to our security will only increase. The dangers and the benefits of technology have been in conflict throughout human history, and Poste concluded by warning that we are beginning to approach an entirely new domain of complexity and intensity.