CHAPTER EIGHT

Media Content Labeling Systems: Informational Advisories or Judgmental Restrictions?*

Donald F. Roberts

The past several years in the United States have witnessed a remarkable debate over whether and how to control media content. The discussion has included most of the media—film, television, popular music recordings, computer games and video games, and, of course, the Internet and the World Wide Web (traditional print media have been largely ignored)—and has ranged from arguments about whether controls are needed at all, to what kinds of controls best fit U.S. political and social needs. One recent upshot of this debate, although hardly the end of the discussion, has been Federal legislation mandating that a V-chip be installed in virtually every new television set sold in the U.S., the industry announcement of a companion TV rating system in January 1997, and a remarkable outpouring of public and government dissatisfaction with that system, leading to its modification less than a year later.

This paper considers why the content rating issue has gained such momentum, briefly reviews empirical research on current portrayals of violence on television and on consequences of exposure to such portrayals, and discusses what the V-chip is and how it works. It proceeds to argue that an informational content labeling system is preferable to a judgmental and restrictive rating system such as the one recently adopted, at least for the time being, by the U.S. television industry, and closes with a description of such an informational advisory system.

PROTECTING CHILDREN

Children are presumed, quite justifiably, to be different from adults—to be more vulnerable, less able to apply critical judgmental standards, more at risk (cf. Roberts, 1993, 1997). As

a consequence, attempts to do anything about media content, whether to label it, to restrict access to it, or to censor it totally, are generally justified in terms of keeping children from harm.

Such arguments are not new. Consider these comments by a psychiatrist, Dr. Edward Podolsky, to a U.S. Senate Subcommittee on Juvenile Delinquency. He spoke following the committee's viewing of excerpts from several televised crime shows.

> Seeing constant brutality, viciousness and unsocial acts results in hardness, intense selfishness, even in mercilessness, proportionate to the amount of exposure and its play on the native temperament of the child. Some cease to show resentment to insults, to indignities, and even cruelty toward helpless old people, to women and other children. (in Starker, 1989, p. 137)

I selected that particular quote because it implicates several of the consequences I will be discussing in this essay, and because of its date—1954. I suspect the programs that committee viewed forty-four years ago would elicit smiles, or yawns, if they were held up as examples of television violence today.

Here is another statement about children and the mass media:

> The tendency of children to imitate the daring deeds seen upon the screen has been illustrated in nearly every court in the land. Train wrecks, robberies, murders, thefts, runaways, and other forms of juvenile delinquency have been traced to some particular film. The imitation is not confined to young boys and girls, but extends even through adolescents and to adults. (in Starker, 1989, p. 8)

That is taken from a now defunct periodical entitled *Education*, commenting on the new mass medium—film . . . in 1919.

I could continue moving back through history in hundred year chunks, reading similar expressions of concern about media content referring to each and every new medium, including print. But let me end with one final quote:

> Then shall we simply allow our children to listen to any story anyone happens to make up, and so receive into their minds ideas often the very opposite of those we shall think they ought to have when they are grown up?

The classicists may recognize that this is Plato, giving his justification for censorship as a necessary condition for building the ideal citizen to inhabit the Republic. My point is simply that fear of what the media may do to children is nothing new. Humans have always wrestled with the issue of what kinds of media content might be inappropriate for children—and what should be done about it.

Calls for Content Labeling

Responses to the question seem always to have ranged from "do nothing" at one extreme to "burn the books" (films/games/records—authors!) at the other. A middle

ground, in the U.S. at least, has taken the form of calls for implementation of some kind of content labeling or rating system—that is, some means to identify the appropriateness of media content for children, and then to use that system either to empower parents, to control children's access, or some combination of the two. Most of us are familiar with motion picture ratings. In the U.S., they have been around since at least 1931, when the Hayes Production Code went into effect, and have been continued since 1968 in the form of the Motion Picture Association of America's (MPAA) movie classification and rating system (see Federman, 1996). So why the recent upsurge in concern and debate?

Why have ratings become such a social and political issue in the 1990s? There are probably many reasons that public concern with "doing something" about media content has reached such a crescendo in the past few years. In the case of the U.S., two of the more important factors are that several negative social trends began to peak at the same time that advances in communication technology enabled popular media to present content in new and more disturbing ways than ever before. Just when our society was experiencing dramatic (and unconscionable) increases in teenage violence and crime, in teenage pregnancies and venereal disease (Hechinger, 1992), and in just plain incivility, the media also began to portray violence, sex, and incivility in what seemed to be greater proportions (actually, levels of television violence have remained remarkably constant for over twenty years; see Gerbner & Signorelli, 1990; National Television Violence Study, 1997) and—more important—more graphically than ever before. (I suspect that increased graphicness feeds the perception of increases in amount of violence portrayed.) Film and television have now developed techniques to make bodies explode and blood spray right before—if not into—audiences' eyes; video games now reward kids for the number of on-screen enemies they can decapitate, with bonus points for extra blood and gore; some popular music lyrics, Web sites, and premium cable channel films make available—indeed, make almost commonplace—sexual content that, in the U.S. at least, once resided almost exclusively within "brown paper wrappers."

Given that adults have always worried that the messages media bring from "outside" may exert undue influence on children (Roberts, 1997), it is not surprising that the co-occurrence of these two trends led to a perception that the mass media are "obviously" having a negative impact on society, and, therefore, that controls or restrictions are needed. Given the complexity of devising regulation that satisfies the First Amendment guarantee of freedom of expression, one of the few viable options for exercising some kind of control seems to lie with a content labeling or rating system—*so long as it is not implemented by the government.* A *New York Times* poll published in July 1995, found that over 80% of all adult Americans and 91% of all parents favored the establishment of a rating system for television; 80% of parents believed that music recordings should be rated; 86% of parents thought videotapes and video games need ratings (Sex and power in popular culture, 1995).

RESEARCH ON MEDIA VIOLENCE

Before considering the kinds of rating systems that have been proposed and implemented in the U.S., I want to define some boundaries, make clear a premise or two, and briefly

consider several of the issues that have been discussed in the ongoing debate about ratings, issues that I believe are central to understanding what a good content labeling system will look like. For the most part (albeit not exclusively), I will focus on television and television content—especially violent content—because that has been the most consistent subject of relevant scientific research. Nevertheless, and this is a basic premise that should be explicit: *a screen is a screen.* Viewers, especially children, do not respond differently to movie screens, television screens, and computer screens; what holds for one probably holds for the others. In other words, in terms of how members of the audience are affected, the issue is the nature of the content, not the channel by which the content is delivered. My second premise is that most of the psychological principles that guide human responses to screen portrayals of violence also guide responses to portrayals of any other kind of behavior, from sexual to altruistic to how to kick a football. The same kinds of things that increase the likelihood a child will learn a violent act from television also increase learning of an altruistic act, or any other kind of act. Obviously, there are some differences across media and across types of content, but on the whole the evidence indicates that the similarities are far more important than the differences.

It is also important to note at the outset that I distinguish between content labeling systems and content rating systems. The two terms are not interchangeable; they refer to quite different approaches to content advisories. Indeed, the distinction between the two is at the heart of this paper's argument. For me, the fundamental difference is one of providing information about content and allowing consumers to make decisions (good or bad) versus imposing restrictions or prohibitions on potential consumers based on someone else's evaluation of the information and judgment about the capabilities and/or vulnerabilities of potential consumers.

Given those caveats and assumptions, let us look at some of the issues in the debate over whether ratings are needed, and, if so, what form they should take. First, we need to spend a few minutes looking at what the research tells us about the impact of screen violence on children's beliefs, attitudes, and behavior, and about what is currently portrayed on U.S. television.

Most research on the consequences of exposure to media violence has focused on viewers' learning of aggressive behavior or attitudes through exposure to entertainment violence. Several exhaustive reviews of the hundreds of scientific studies conducted during the past forty years lead to the unequivocal conclusion *that exposure to mass media portrayals of violence contribute to aggressive attitudes and behavior in children, adolescents, and adults* (see, for example, Comstock with Paik, 1991; Huston, Donnerstein, Fairchild, Feshbach, et al., 1992; Paik & Comstock, 1994). Obviously, media violence is not the only cause of violent social behavior, but few social scientists any longer debate that it plays a contributory role. As long ago as 1982, a National Institute of Mental Health report on television and behavior concluded: "In magnitude, television violence is as strongly correlated with aggressive behavior as any other behavior variable that has been measured" (National Institute of Mental Health, 1982). Studies conducted in the intervening fifteen years have not altered that judgment (Comstock & Paik, 1991). More often than not, those who continue to claim that there is no evidence for such a causal connection

tend to be associated with the media industry—and/or simply have not read (or choose to ignore) the scientific literature.

What the past decade and a half of research has added, however, is evidence that exposure to media violence can have negative consequences beyond increasing the likelihood of viewers' aggressive behavior. We now know that prolonged violence viewing can also lead to emotional *desensitization,* engendering callous attitudes toward real-world violence and decreasing the likelihood of helping real victims. In addition, a third consequence of violence viewing is increased *fear* of becoming a victim, which in turn leads to such things as mistrust in others and to increases in self-protective behavior. In short, research evidence confirms that excessive exposure to media violence can lead to learning aggressive behavior, to desensitization, and to fear, and that several of these outcomes might occur simultaneously (for reviews see Comstock & Paik, 1991; Wilson, Kunkel, Linz, Potter, Donnerstein, Smith, Blumenthal & Gray, 1996).

Given that such consequences of viewing violence are well documented, the more interesting research questions (particularly when faced with developing a content labeling system) concern identification of the contextual factors within media content that seem to make a difference. In other words, what are ways of portraying violence that increase or decrease the likelihood of a negative effect? Both intuitively and on the basis of scientific research, we know that some violent programs are more problematic than others, that some ways of displaying violence are likely to increase learning, fear, or desensitization, but that other depictions are quite likely to decrease these outcomes. It does not take a scientific background to sense that the consequences to viewers of the violence in a film like *Schindler's List* (in which a man saves numerous Jews from the Nazi concentration camps during World War II) are probably quite different than the consequences of the violence in a film like *Natural Born Killers* (in which young adults blast a bloody swath across the U.S.). Both films portray brutal violence, both show a number of killings, both are relatively graphic—yet one is generally thought of as an antiviolence statement while the other has been accused of celebrating violence. The interesting question is: Why? What are the differences in how each portrays violence that make the two films so different? If we are to design a content rating system that will differentiate between two such different media portrayals, such questions are critical. A simple body count will not do the job.

Fortunately, as part of a massive content analysis of violence on U.S. television (The National Television Violence Study, 1996, 1997), Barbara Wilson and her colleagues (1996) reviewed the experimental research on media violence with an eye to identifying contextual factors that make a significant difference in how viewers respond to violent content. Nine factors emerged from the experimental research literature: (1) the nature or qualities of the perpetrator; (2) the nature or qualities of the target or victim; (3) the reason for the violence—whether it is justified or unjustified; (4) the presence of weapons; (5) the extent and/or graphicness of the violence; (6) the degree of realism of the violence; (7) whether the violence is rewarded or punished; (8) the consequences of the violence as indicated by harm or pain cues; (9) whether humor is involved. Although the amount of research on each individual factor varies (i.e., we know a great deal about the role of rewards and punishments but not a great deal about the role of humor), Wilson and her associates contend that there is adequate evidence to safely conclude that each

identified factor either *increases or decreases* the probability that a violent portrayal poses a risk to viewers on at least one of the three outcomes: learning, desensitization, or fear. When these contextual elements are mapped over the three outcomes, the matrix shown in table 1 results. The arrowheads show what experimental research says about how each contextual factor affects each outcome. Thus, for example, when violence is rewarded we expect an increase in both learning and fear; when violence is portrayed as unjustified, we expect a decrease in learning but an increase in fear; humor should increase both learning and desensitization; and so on. The spaces where no arrowhead occurs indicate a lack of adequate evidence concerning how that particular contextual factor affects that particular outcome. For example, no research examining how harm and pain cues affect either fear or desensitization was located.

Research on Television Content: The National Television Violence Study

The matrix in table 1 served to guide the content analysis component of the National Television Violence Study (NTVS), an ongoing, three-year study of violence and U.S. television. Although the overall study includes several different components, description of the work is limited here to its examination of the nature of violent television content.

Each year the NTVS researchers sample and analyze the content in a *representative week of U.S. entertainment television.* I include the italics to emphasize the magnitude of the task. For example, for the 1994–95 season, they sampled 23 channels of television available in the Los Angeles area, including broadcast networks, independent channels, public television, basic cable and premium cable channels. For each channel they randomly selected two daily, half-hour time slots between 6:00 A.M. and 11:00 P.M. over a period of 20 weeks, ultimately taping a total of 3,185 programs. After eliminating news programs, game shows, religious programs, sports, instructional programs and "infomercials" (none of which fell within their contracted definition of entertainment programming),[1] they were left with a sample of 2,693 programs—2,737 hours of programming. This resulted in a representative 7-day composite week of programming for each of the 23 channels, the largest and most representative sample of entertainment television content ever collected.

The coding scheme in this study is equally detailed and comprehensive. Violence is defined as:

> . . . any overt depiction of a credible threat of physical force or the actual use of such force intended to physically harm an animate being or a group of animate beings. Violence also includes certain depictions of physically harmful consequences against an animate being that occur as a result of unseen violent means. Thus there are three primary types of violent depictions: credible threats, behavioral acts and harmful consequences. (National Television Violence Study, Content Analysis Codebook, 1994–1995, p. 3)

But more important than the definition of violence per se, precise definitions have been developed for all of the contextual factors listed in table 1. That is, coding instructions were created to enable coders reliably to identify such content elements as harm, pain,

TABLE 1

Predicted Impact of Contextual Factors on Three Outcomes of Exposure to Media Violence

	OUTCOMES OF MEDIA VIOLENCE		
CONTEXTUAL FACTORS	LEARNING AGGRESSION	FEAR	DESENSITIZATION
Attractive Perpetrator	↑		
Attractive Target		↑	
Justified Violence	↑		
Unjustified Violence	↓	↑	
Presence of Weapons	↑		
Extensive/Graphic Violence	↑	↑	↑
Realistic Violence	↑	↑	
Rewards	↑	↑	
Punishments	↓	↓	
Pain/Harm Cues	↓		
Humor	↑		↑

From the National Television Violence Study, 1996. Predicted effects are based on review of social science research on contextual features of violence. Spaces are used to indicate that there is inadequate research to make a prediction. Reprinted with permission.

↑ = likely to increase the outcome

↓ = likely to decrease the outcome

TABLE 2

Selected Findings from National Television Violence Study
Attributes of Television Content, 1994–95

58% OF ALL ENTERTAINMENT PROGRAMS CONTAINED VIOLENCE

Violent Programs

- 33% contained 9 or more violent interactions
- 51% portrayed violence in realistic settings
- 16% showed long-term consequences of violence
- 4% had an antiviolence theme

Scenes within Violent Programs

- 15% of violent scenes portray blood and gore
- 39% of violent scenes use humor
- 73% of violent scenes portray violence as unpunished

Interactions within Violent Scenes

- 25% of interactions employed a gun
- 35% of interactions depicted harm unrealistically
- 44% of interactions showed violence as justified
- 58% of interactions **did not** depict pain

Adapted from the National Television Violence Study, Executive Summary, 1996. Reprinted with permission.

humor, justification for violence, attractiveness of the target of violence, and so forth. Thus, rather than simply counting how often violence occurs in current entertainment television programming, the NTVS analysis provides a detailed picture of the contextual features associated with portrayals of violence. Finally, the coding scheme operates at three distinct levels—that of the overall program, of the scene, and of the individual violent act (i.e., violent interaction), enabling independent inferences about the nature and context of violent acts, violent scenes, and violent programs. Such a multilevel approach is necessary if one is to be able to differentiate between programs that glorify violence and those that condemn violence. For example, although a program with an antiviolence theme may depict as many violent scenes as a program that glorifies violence, the more global antiviolence message may emerge only at the program level. If analysis were to be limited to individual acts or scenes, this point might be lost.

Obviously, a study of this magnitude produces results far too extensive to detail here. Nevertheless, a brief summary of a few of the findings will help form the foundation of my argument about what kind of content labeling system will best serve the television audience.

The first conclusion from the NTVS study is not surprising—there is a great deal of violence on U.S. entertainment television. Indeed, in 1994–95 more than half of all entertainment programs—58%—contained violence (the comparable number for 1995–96 was 61%). More interesting than the total amount of violence, however, is its nature—that is, the context in which it is portrayed and the attributes with which it is associated. Table 2 summarizes a few of the contextual results.

For the most part, the table speaks for itself. Although a substantial proportion of U.S. entertainment television contains no violent content whatsoever, over half of the programs do portray violence. Moreover, when violence does occur it is often portrayed in ways that are more likely than not to increase the chances of some kind of negative effect on viewers. Violence often goes unpunished, seldom results in either immediate pain or negative long-term consequences, and is often portrayed as something to laugh about. Over a third of violent interactions depict harm unrealistically, almost 45% portray violent acts as justified, and well over half fail to depict any associated pain. These are all factors that have been shown to increase the likelihood of viewers learning to be more aggressive, or becoming more fearful, or becoming more desensitized. In other words, the contextual factors characteristic of much (not all, but much) U.S. television programming are just those that *increase* the likelihood of negative consequences among youthful viewers. Indeed, the results of the NTVS content analysis read like a primer on how *not* to produce programming for children.

THE V-CHIP

In February 1996, President Bill Clinton signed into law the Telecommunications Act of 1996, a far-reaching piece of legislation that is destined to change the face of the U.S. telecommunications. One small part of that legislation was intended to empower parents by providing a way for them to control the television content to which their children can have access. This was accomplished by mandating that within two years of the signing of the bill (February 1998), all new television sets sold in the U.S. must contain a V-chip, and that within one year (February 1997) the television industry must have developed a system to implement V-chip capabilities (otherwise, the Federal Communications Commission would appoint an independent committee to do it for them). Now what does this mean? What is the V-chip, and what is the system designed to implement it?

Briefly, a V-chip is simply a piece of hardware, a very tiny piece of hardware, that will be included in the electronics of new TV sets (or added to existing TV sets). It allows consumers to block programs depending on how their content is labeled or rated. The chip reads a signal that is not visible to viewers (it is embedded in the vertical blanking interval, the portion of the television signal that currently carries closed caption services for the hearing impaired). That signal, which is to be included within every television program, will carry information about the content of the program. The consumer can

program the chip to recognize and respond to any particular rating or level of intensity or other kind of information embedded in the signal. Programs that fail to meet the selected criteria, whatever they might be, are blocked. Thus, for example, a show might be labeled somewhere between V-0 (for no violence) to V-4 (for a great deal of violence). Or, using the television rating system employed from December 1996 through September 1997, it might be rated anywhere from TV-G (General Audience), through TV-PG (Parental Guidance Suggested) or TV-14 (Parents Strongly Cautioned), to TV-M (Mature Audience Only). On the basis of such a label or rating, parents decide what kinds of shows are to be allowed in their home and set the V-chip to block anything in excess of, or not conforming to, the selected criteria. Once that selection is made, the chip automatically decodes the signal embedded in each program and acts in accordance with parental (or other consumer) decisions. If the program exceeds the rating, the V-chip picks up the signal and the screen simply goes blank. In short, the chip is simply a device that enables consumers to decide what kinds of television content they want to allow into their homes at any given time, and to block out any content that does not meet their standards.

Several other things about this technology are important to note: (1) the V-chip is capable of accommodating any one of a number of different labeling or rating systems; (2) the chip can accommodate several different systems simultaneously (there is no requirement to settle on a single approach); (3) a single program can have independent ratings for different kinds of content; that is, there can be one rating for violence, another for sex, and another for language, all pertaining to the same program; and (4) the chip can be turned on or off, or reprogrammed, at any time. For all these reasons, I believe the chip has been misnamed. Initially the "V" was appended to indicate violence chip, but since it can do much more than respond to violence levels, I think a more appropriate name would be C-chip, standing for "Choice Chip." To make it a real choice chip, however, requires giving consumers the necessary information to make reasoned choices—what I call an informational content advisory.

Informational vs. Judgmental Systems

An informational content labeling system posits that information contained in an advisory helps consumers direct their behavior by telling them what is in "the package"—that is, what is contained in the program, film, or game they are considering. The usefulness of the information depends on how clear, specific, and relevant it is to a given consumer. For example, assume one wishes to avoid—or select—content depicting violent or sexual behavior. In this case, a label explicitly describing the kind and amount of either behavior is more helpful than content-free proscriptions that simply warn the content may be problematic but do not state why (e.g., TV-14). In other words, informational systems assume that the primary function of content advisories is to inform viewers about what to expect, and that the more fully they do this, the better. An informational system leaves open both the question of appropriateness and the selection decision.

Judgmental approaches—for example, the MPAA Film Classification System—generally do not provide much descriptive information. Rather, they make judgments about what is or is not judged to be appropriate for particular audiences—specifically, for

different age groups of children. Thus, a TV-14 rating tells consumers that somebody has made a judgment that something about the content is inappropriate for children younger than fourteen, but says little or nothing about what that content is (e.g., violence, sex, inappropriate language, etc.). In the most extreme cases, such judgments become proscriptions. For example, in the U. S., youngsters under seventeen years old are prohibited from attending an R-rated film unless accompanied by an adult. In other words, judgmental approaches hand over to someone other than the consumer the question of what is appropriate, and in some cases, the selection decision. Usually the judgment is made by some relatively anonymous ratings board (Federman, 1996).

Typically, two rationales are offered for adopting a judgmental as opposed to an informational approach. First, it is argued that given the thousands of hours of media content produced each year, there is no way to develop a descriptive system complex enough to identify the kinds of content differences that proponents of informational systems would like to describe, but still simple enough to be employed by whomever is charged with the task of labeling. Second, even if an informational system could be developed, proponents of judgmental systems say that it would be far too complex for most consumers to use. Rather, they argue, parents are more likely to use a system that only requires them to make a single, simple, age-based choice.

By now it should be clear that I favor informational content labeling systems over judgmental systems. There is, of course, the possibility of combining the two approaches—of both telling the consumer what is in the package *and* providing judgments about its age-appropriateness. But even that, I think, is a mistake. Not only do judgmental systems take fundamental decision-making power away from parents, but they also increase the risk of attracting children to the very kinds of content from which we would like to protect them. Even though content advisories are intended to help parents monitor and guide their children's media consumption, we cannot lose sight of the fact that youngsters also see and respond to these ratings. Nor can we ignore that content decisions are under control of at least some children most of the time and of most children at least some of the time. It follows that how content advisories affect children also warrants careful consideration.

BOOMERANG EFFECTS

Unfortunately, there is mounting evidence that advisory labels can boomerang, attracting youngsters to inappropriate content—a kind of "forbidden fruit" effect (Christenson, 1992). And although both informational and judgmental advisories have been found to boomerang sometimes, the effect is more general and more consistent with judgmental ratings (see Bushman & Stack, 1996; Cantor & Harrison, 1996; Cantor, Harrison & Nathanson, 1997; Christenson, 1997; Morkes, Chen & Roberts, 1997).

To the extent that informational systems attract children to "forbidden fruit," they do so because they identify content that youngsters seek because of some need or interest independent of the labeled content. For example, youngsters interested in sex or violence for whatever reason will read the advisory to determine whether a given program can satisfy their interest, and act accordingly. Children not interested in these topics may either

ignore or actively avoid the program, also depending on the information in the advisory. It is no different than when people who are interested in gardening read program listings to locate programs about gardening. For better or worse, information is used to guide choices; not approving of the choice a person might make on the basis of information is not a legitimate reason to withhold information.

The boomerang effect associated with judgmental ratings, on the other hand, is not primarily a function of fulfilling a child's information needs; judgmental ratings provide little information. Rather, when children are drawn to content rated as inappropriate by a judgmental system (e.g., when a twelve-year-old chooses a TV-14 program), it is primarily because they are reacting against what they perceive to be someone attempting to control their media choices; such a reaction is quite independent of what the content may be. Reactance theory (Brehm, 1972) posits that a perceived threat to individual freedom motivates humans to restore freedom by actively seeking to engage in the proscribed behavior. Thus, to the extent that children perceive content advisories as attempts by some "authority" to limit their access to content or otherwise impose control or censorship, the theory predicts that they will strive to consume the proscribed material *regardless of the nature of the content.* Several studies indicate that youngsters perceive labels proscribing content on the basis of age or appearing to put control in the hands of others—particularly parents (e.g., an advisory such as "Parental Discretion Advised")—as highly restrictive, and that they react strongly against them. At least three experiments have shown that the MPAA Film Classification system is particularly likely to engender reactance and a boomerang effect among children (Cantor & Harrison, 1996; Cantor, Harrison & Nathanson, 1997; Morkes, Chen & Roberts, 1997; also see Bushman & Stack, 1996).

Of course, the informational and the judgmental models are not entirely independent. Simply the fact that any rating is assigned—whether a single letter icon (e.g., "R") or a descriptive phrase (e.g., "Humans killed; blood and gore")—indicates that someone hopes to control at least some consumers' access to the content, thus creating some potential for reactance. Similarly, even the most "content free" label usually elicits consumer inferences about the nature of the proscribed material. For example, when asked what an MPAA "R" stands for, most young adolescents in the U.S. will refer to sex and/or violence. Moreover, both mechanisms may operate simultaneously. A twelve-year-old boy might seek out an R-rated film both because he has been told he can't see it (reactance) and because he believes it may portray activity about which he is curious (information seeking). Whether a particular rating or advisory provides information or elicits reactance, then, is a matter of degree. When concerned with children's responses to ratings from a practical point of view, the question is better phrased in terms of which systems are less likely to cause reactance and more likely to provide useful information.

THE RSAC CONTENT LABELING SYSTEM

Let me turn, then, to the issue of whether it is possible to design an informational system complex enough to give relatively fine-grained information about a program but still simple enough for both labelers and parents to use. I shall describe a content advisory system I

helped devise a few years ago and start by describing some questions we addressed from the beginning of the project.

First, consider how satisfied a parent would be with a system that rated programs with either a simple G (good for children) or NG (not good for children). Prior to viewing, that's all one would know about the program—either its G or NG. Would the situation be better if there were four or five ratings levels—say from 0 to 5, or from TV-G through TV-PG, and TV-14 to TV-M (the TV Parental Guidelines)? My experience has been that most parents would prefer either of the latter options to the simple G or NG. But is that enough? Wouldn't the "parent" want to know what the ratings mean by "children?" That is, would it help to know if G - NG referred to seven-year-olds, ten-year-olds, or fourteen-year-olds? Would your answer to the questions change depending on *who* gave the rating, on whether, for example, the G (for "good") or NG (for "no good") was assigned by a leader in your public educational system or by a youthful college dropout marking time as a television content rater while awaiting a "real" job? What if the rating was always assigned by one of two educators, one of whom was obsessed with keeping violence off television while the other made keeping children safe from nudity his life's work—but you never knew which gave a particular rating? Would it make a difference which one assigned the rating? Are you more concerned with portrayals of violence than with portrayals of sex? Do you have different feelings about depictions of nudity? Of vulgar language? And even if you decide which kind of content concerns you most, are you certain about how any particular portrayal should be rated? Perhaps what you see as brutal violence someone else will judge to be little more than a friendly tussle. Indeed, consider the wide range of answers the preceding questions are likely to elicit from a large, diverse group of parents.

The importance of such questions began to emerge for me when I was asked by the U.S. Software Publishers Association to help develop a parental advisory system for computer games.[2] In 1994, in response to the release of several particularly violent and bloody video games, some members of the U.S. Congress brought pressure to bear on both the video game and computer game industries to develop some kind of parental advisory label to be placed on game packages (see Federman, 1996). At minimum, the argument went, parents should have some indication of what is in a game before they purchase it for their children. I won't detail the history of that particular debate, except to note that in order to preclude threatened government action, each of the two industries (video games and computer games) developed its own system, and the one produced for computer games took an informational approach. Ultimately, when the computer game content labeling system was completed, it was turned over to a nonprofit advisory board independent of the computer game industry, the Recreational Software Advisory Council (RSAC), and became known as the RSAC system. Over the past two years a slightly revised version of the system, called the RSACi system ("i" for Internet) has gone into effect on the World Wide Web, and over 60,000 Web sites currently use it to label content.

Several factors influenced the shape of the RSAC content labeling system. Most important was the issue of whether an advisory should be judgmental or informational—(evaluative or descriptive). That is, should a content advisory make an evaluative judgment about what a child should see, or should it provide descriptive information about

what is in the game, allowing parents to make the evaluative judgments appropriate to their personal beliefs and value systems? Is it better to label a program as "inappropriate for children under thirteen years," or to say "this game depicts violence that goes unpunished and that results in injury to humans," asking parents to decide whether their children should play? When I talked to parents about the Motion Picture Film Classification System, I found that many objected to age-based, judgmental ratings because they believed that often such ratings were not appropriate for their own children. Some felt their ten-year-olds were perfectly capable of handling some kinds of content likely to get a PG-13 rating, but not other kinds; some felt that their fourteen-year-olds should not see a PG-13 movie, but had a great deal of trouble defending their position in the face of such "expert" ratings; most complained that the simple lettering system simply did not tell them enough to enable them to exercise informed judgment. They believe that a PG-13 rating can be assigned to a film on the basis of violence, or sex, or language, but they often are uncertain about what particular kind of content is at issue in any given film—sometime even after they have seen the film. Many parents indicate that all they really know when faced with a PG-13 rating is that someone has made an evaluative judgment that the content is "inappropriate" for younger children.

Of course, to the extent that an advisory provides both descriptive information and an age-based judgment, parents can make a decision based on descriptive information combined with the additional knowledge of someone else's evaluative judgment about appropriate age levels. That evaluative judgment, however, is or is not valuable to the parent depending on who that someone else is and what the criteria underlying the judgment were—information that is not currently available for the MPAA system (Federman, 1996). Moreover, in some circumstances that judgment can override a parental decision; that is, parents cannot decide to have their sixteen-year-old attend an R-rated film absent the company of an adult.

In any case, given: (a) parents' expressed desire for more information, (b) game developers' antipathy toward others making evaluative judgments about their products, and (c) reactance theory's prediction that age-based content restrictions are likely to boomerang, we opted for informational content labeling as opposed to judgmental ratings. Ultimately, we took as our model the U.S. food labeling system, which requires food packagers to list the ingredients in the package. Consumers are not told what they should or should not eat; rather, they are given adequate information and the consumption decision is left to them. The RSAC content labeling system used the same principles.

Another important factor that shaped the final form of the RSAC system was logistical. The nature of computer games makes it very difficult to require that they be screened by independent raters. Unlike films or videotapes, which can be viewed in ninety minutes or so, it can take upwards of one hundred hours to review a computer game (make that two hundred hours if you are over forty years old). Given the hundreds of games that need to be labeled each year, it would be extremely expensive and impractical to require independent coders to describe or rate each game. We decided, therefore, to develop a self-rating system—that is, a system whereby the game developers themselves rate their own games. This, of course, created a new problem. To ask a game developer to label his or her own game, particularly when many developers tend to believe that labels or ratings

indicating higher levels of violence (or sex, or vulgar language) may decrease sales, is like asking the fox to guard the henhouse. It would seem to invite the developers to bend the rules. Thus, we had to find a way to keep game developers accurate and honest as they labeled their own games, and equally important, a way that would also assure the public that such self-administered ratings are, in fact, accurate and honest. Ultimately, the solution turned out to be quite simple. We took the norms and canons of science and moved them into the public arena. That is, we developed a content labeling system that is *reliable* and *public*, and those two attributes largely solved the problem.

A reliable system means that any two individuals using the coding procedures correctly will describe or rate a game identically. This requires very concrete, very detailed definitions of everything to be described, and a set of questions about the content based on those definitions that ask for nothing other than yes/no responses. The idea is that no matter how different the individuals, if they use the same objective definitions correctly, and answer the questions honestly, they cannot help but assign the same label or rating to a game.

A public system means open to public oversight; that is, anyone and everyone has access to the system, its definitions, and its procedures. To the extent that open access to a reliable system is guaranteed, then anyone should be able to check the label or rating given to any game at any time. The idea underlying this requirement is that if it is easy for anyone in the public to raise questions or objections in those instances when they do not agree on the rating (using, of course, the same rating system), the threat of such checks keeps game developers honest. If the game developers misuse the system, they face loss of their rating (which can cost them access to retail outlets) and heavy fines. (A public system can also provide increased flexibility in that, over time, public input can be used to sharpen or modify questions and/or definitions, keeping the system in step with cultural norms.)

Finally, there remained the question of what to label and how to label it. Both public opinion and prodding from Congress dictated advisories addressing each of four content dimensions—violence, sex, nudity, and language (ultimately, the labels combine sex and nudity, but the two kinds of content are still rated separately). Here, however, I will focus on violence.

We decided that there would be five levels of intensity for each content area—that is, from 0 for no violence to 4 for the most potentially harmful portrayals of violence. We reviewed the research literature on the effects of media violence, identifying content dimensions—what Wilson et al. (1996) called contextual factors—that were most likely to increase negative effects and that seemed most appropriate to the *content of games*. (Since games do not have the kinds of story lines found in dramatic narratives, their content labels focus on slightly different dimensions than might be the case for television programs.) By this procedure we settled on five primary features that would make a difference in the level of the advisory:

1. the nature of the target (victim)—i.e., is the target humanlike, nonhuman, or an object?;
2. the stance of the target (victim)—i.e., is the target threatening or nonthreatening?;

3. consequences to the target (victim)—i.e., death v. injury v. disappearance v. no consequences;
4. depiction of blood and gore;
5. consequences to the player—i.e., is the player rewarded or not rewarded for aggressive behavior?

To the extent that one or more of these attributes occur within a computer game, the advisory of the level of violence increases. For example, if the game portrays a threatening human attacked but not injured, the game gets a 1 for violence; if the threatening human is injured, it gets a 2, and so on. Combinations of these various dimensions result in the logic chart shown in table 3, which illustrates the various attributes in a computer game that result in different violence advisories.

Of course, the person judging the content is not required to make his or her way through that chart. Rather, there is a set of highly concrete, highly objective definitions for every term used in the chart, and a parallel set of yes/no questions employing those definitions. For example, one question asks: "Does the software title depict blood and gore of sentient beings?" That question results in a straightforward "yes" or "no" response from the person doing the labeling because the terms "depict," "blood and gore," and "sentient beings" are each explicitly and extensively defined. (Terms such as "sentient beings" would never be used in a descriptive label, but are included in the definitions in order to cover the wide array of creatures inhabiting the world of computer games—from realistic humans, to animated space aliens, to killer frogs.) Here, for example, is part of the definition of "blood and gore":

> Blood & Gore: Visual Depiction of a great quantity of a Sentient Being's blood or what a reasonable person would consider as vital body fluids, OR a visual Depiction of innards, and/or dismembered body parts showing tendons, veins, bones, muscles, etc., and/or organs, and/or detailed insides, and/or fractured bones and skulls.
>
> The depiction of blood or vital body fluids must be shown as what a reasonable person would classify as flowing, spurting, flying, collecting or having collected in large amounts or pools, or the results of what a reasonable person would consider as a large loss of the fluid such as a body covered in blood or a floor smeared with the fluid . . . etc.

There are literally dozens of pages of such definitions and associated examples, one for every important term in each of the questions.

The questions are arranged in a branching format and are typically administered on a computer. Depending on the response to any given question, the system either gives an appropriate content label or determines what the next question should be. Depending on the amount and nature of violence in a given game, whoever does the labeling may respond to as few as two or as many as fifteen questions. The same procedure is followed for sex/nudity and for language. Finally, depending on how the questions have been answered, the program determines what the advisory icon should be (see figure 1) and

TABLE 3

RSAC Methodology Logic Chart

	All	1	2	3	4
Maximum Violence					
Rape					X
Wanton and Gratuitous Violence					X
Blood/Gore				X	
Human Threatening Victims					
No Apparent Damage/No Death		X			
Damage with or without Death			X		
Death/No Damage			X		
Human Non-Threatening Victims					
Damage/No Death					
Player Not Rewarded (unintentional act)			X		
Player Rewarded				X	
Death With or Without Damage					
Player Not Rewarded (unintentional act)			X		
Player Rewarded (gratuitous violence)					X
Non-Human Threatening Victims					
No Apparent Damage/No Death		X			
Damage With or Without Death		X			
Death/No Damage		X			
Non-Human Non-Threatening Victims					
Damage/No Death					
Player Not Rewarded (accidental)		X			
Player Rewarded (intentional)			X		
Death With or Without Damage					
Player Not Rewarded (accidental)		X			
Player Rewarded (intentional)				X	
Natural/Accidental Violence					
Damage/Death-Human Victims			X		
Damage/Death Non-Human Victims		X			
Blood/Gore (humans and non-humans)				X	
Objects (Aggressive & Accidental Violence)					
Damage and/or Destruction of Symbolic Objects	X				
Realistic Objects					
Disappear w/o Damage or Implied Social Presence	X				
Disappear w/o Damage with Implied Social Presence		X			
Damage With or Without Destruction		X			

Reprinted with permission.

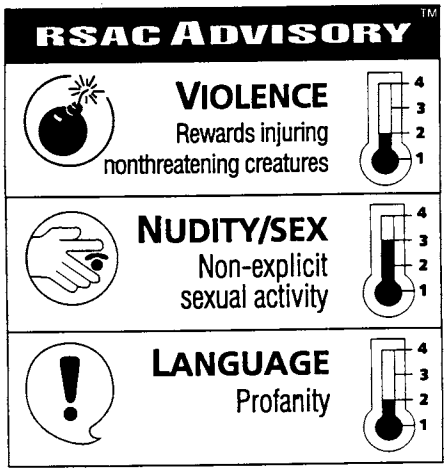

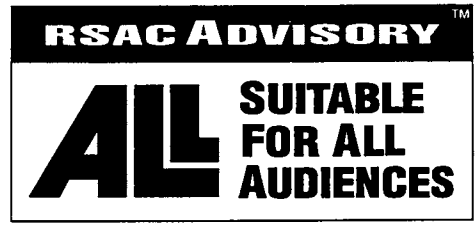

FIG. 1. Example of two RSAC Advisory labels for hypothetical computer games: one with no instances of Violence, Nudity/Sex, or Language requiring a label (i.e., Suitable for All), and one with level 2 Violence and Language and level 3 Nudity/Sex.

what information is to be used to explain that icon. In other words, the final label consists of both an icon and associated number (from 0 to 4) indicating the level of violence, and a descriptive phrase explaining why that number was assigned. As shown in the logic chart in table 3, the content description associated with any given level may vary. For example, there are six different reasons a game can receive a violence score of 2 and four different reasons it could earn a violence score of 3. In all cases, the advisory informs the consumer about the specific kind of content underlying each specific level of violence assigned. Note, for instance, the descriptive information in the violence section of the RSAC advisory label displayed in figure 1. The hypothetical game described by that label received 2 for violence level because it "rewards injuring non-threatening creatures" (e.g., the player scored points for shooting living creatures that posed no threat to other figures in the game). Another descriptive phrase paired with a level 2 violence advisory might have been simply: "Humans injured."

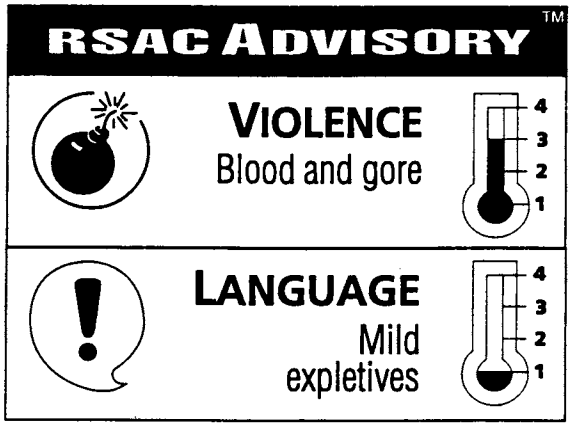

FIG. 2. Example of RSAC Advisory label attached to a particular computer game: "Doom."

Finally, figure 2 shows an advisory label assigned to one of the more violent games available in the U.S. a year or so ago (even more violent games have since reached market), a game called "Doom." As you can see, the game received the next to highest advisory for violence—in this case because it portrayed blood and gore—and a very mild rating for language (mild expletives). There were no instances of either sex or nudity in the game.

CONTENT LABELING AND TELEVISION

Now, how might something such as the RSAC content labeling system relate to a TV content labeling system to implement the V-chip? Clearly, some modifications would have to be made in the dimensions employed and in some of the questions required to assign a label. (The contextual features identified in the NCTV study would be a good place to begin developing the appropriate dimensions for television content). But for the most part, such an informational approach seems ideal for the new V-chip technology. It would be inexpensive and quick, because the producers/writers of each television show would rate their own product with the understanding that the rating procedure is public. More important, it would serve the consumer well because it has the advantage of being descriptive and informational rather than judgmental.

Joel Federman (1996) concluded his recent book on media ratings with the recommendation that whatever rating system is adopted, it should make every effort to maximize information and minimize judgment. Of course, "informational" and "judgmental" are relative terms. Since even the act of choosing to label content implies evaluation, no rating system can be purely informational. Nevertheless, because something like the RSAC content labeling system leans far more in the direction of description than evaluation, it has several valuable advantages over evaluative systems. First, and most important, it puts the decision-making power in the hands of the parents rather than some outside agency with which the parent may or may not agree. It presumes that children are different from each other and that parents know the needs and capabilities of their own children far better than anyone else can. Second, it has the advantage of consistency because the criteria for labeling any content are objective, concrete, and public. And this, in turn, means that it can be used in highly flexible ways. Parents whose primary concern might be media violence and parents whose primary concern might be language or sexuality can all use the system with confidence, adapting it to fit each of their different needs.

There is probably no such thing as a perfect solution to the problem of protecting a highly vulnerable audience such as children while simultaneously protecting people's right to say/write/film/program freely. Nevertheless, providing parents with descriptive information on which they can base informed decisions would be a big step in the right direction—a step that attempts to respond to the needs and right of all concerned parties.

Notes

* Earlier versions of this paper were delivered as The Wally Langenschmidt Memorial Lecture at the South African Broadcasting Corporation in Johannesburg, Republic of South Africa, August 28, 1996, and as an invited address to the Korean Broadcasting Commission, Seoul, Korea, June 3, 1997.

1. The National Cable Television Association, which funded the study, determined which kinds of programming would be included in the content analysis.
2. The design team consisted of Mr. Glenn Ochsnreiter of the Software Publisher's Association (Washington, DC), Mr. Jim Green of Shareware Testing Laboratories (Indianapolis, IN), and myself.

References

Brehm, J. W. (1972). *Responses to loss of freedom: A theory of psychological reactance.* Morristown, NJ: General Learning Press.

Bushman, B. J., & Stack, A. D. (1996). Forbidden fruit versus tainted fruit: Effects of warning labels on attraction to television violence. *Journal of Experimental Psychology: Applied, 2,* 202–36.

Cantor, J. & Harrison, K. (1996). Ratings and advisories for television programming. In *National television violence study, 1995–1996: Scientific papers* (pp. III.1–III.50). Studio City, CA: Mediascope.

Cantor, J., Harrison, K., & Nathanson, A. (1997). Ratings and advisories for television programming: University of Wisconsin study. In *National television violence study 2* (pp. 267–322). Thousand Oaks, CA: Sage.

Christenson, P. G. (1992). The effects of parental advisory labels on adolescent music preferences. *Journal of Communication, 42*(1), 106–13.

Christenson, P .G. (1997, May). *The effect of videogame ratings: Is there a boomerang effect?* Paper presented at the annual meetings of the International Communication Association, Montreal, Canada.

Comstock, George, with Paik, Haejung (1991). *Television and the American child.* San Diego, CA: Academic Press.

Federman, Joel (1996). *Media ratings: Design, use and consequences.* Studio City, CA: Mediascope.

Gerbner, G., & Signorelli, N. (1990). *Violence profile, 1967 through 1988–89: Enduring patterns.* Unpublished manuscript, University of Pennsylvania, Annenberg School of Communications.

Hechinger, F. M. (1992, April). *Fateful choices: Healthy youth for the 21st century—Executive summary.* New York: Carnegie Council on Adolescent Development, Carnegie Corporation of New York.

Huston, A. C., Donnerstein, E., Fairchild, H., Feshbach, N. D., Katz, P. A., Murray, J. P., Rubinstein, E.A., Wilcox, B.L., & Zuckerman, D. (1992). *Big world, small screen: The role of television in American society.* Lincoln, NB: University of Nebraska Press.

Morkes, J., Chen, H. L., & Roberts, D. F. (1997, May). *Adolescents' responses to movie, television and computer game ratings and advisories.* Paper presented at the annual meeting of the International Communication Association, Montreal, Canada.

National Institute of Mental Health (1982). *Television and behavior: Ten years of scientific progress and implications for the eighties, Vol. 1, Summary report.* Washington, DC: U.S. Government Printing Office.

National Television Violence Study (1996). Studio City, CA: Mediascope.

Executive Summary: 1994–95
Scientific Papers: 1994–1995
Content Analysis Codebooks: 1994–1995.

National Television Violence Study, Volume 2: Executive Summary. (1997). Santa Barbara, CA: University of California at Santa Barbara.

Paik, H., & Comstock, G. (1994). The effects of television violence in antisocial behavior. *Communication Research, 21,* 516–46.

Roberts, D. F. (1993). Adolescents and the mass media: From "Leave it to Beaver" to "Beverly Hills 90210." *Teachers College Record, 94,* 629–43.

Roberts, D. F. (1997, in press). From Plato's republic to Hillary's village: Children and the changing media environment. In R.P. Weissberg, C.B. Kuster, O. Reyes, & H.J. Walberg (Eds). *Trends in the well-being of children and youth.* Thousand Oaks, CA: Sage.

Sex and violence in popular culture (1995). *New York Times,* July 23–26, p. 6.

Starker, Steven (1989). *Evil influences: Crusades against the mass media.* New Brunswick, NJ: Transaction Publishers.

Wilson, Barbara J., Donnerstein, Ed, Linz, Dan, Kunkel, Dale, Potter, James, Smith, Stacy L., Blumenthal, Eva, & Gray, Tim (1996). Content analysis of entertainment television: The importance of context. Paper presented to the Duke University Conference on Media Violence and Public Policy, June 28–29, Raleigh-Durham, North Carolina.

Wilson, B. J., Kunkel, D., Linz, D., Potter, J., Donnerstein, E., Smith, S., Blumenthal, E., & Gray, T. (1996). Television violence and its context: University of California, Santa Barbara study. *National Television Violence Study, Scientific Papers* (pp. 5–268). Studio City, CA: Mediascope.

Page Intentionally Left Blank