

Statement of Dr. Ted S. Palys:

Comments on the Statement by Dr. Neil Malamuth¹

7 October 1994

Contents

Qualifications of Dr. Palys	2
Comments on Dr. Malamuth’s Statement	3
His Scientific Perspective.....	4
A Paradigmatic Effects Study.....	6
The Sample: Representativeness Issues.....	9
The Sample: Variability Issues.....	10
The Setting: Ecological Validity Issues.....	11
The Fragility of Media Influence Effects.....	14
Long Term Effects.....	15
Other Variables.....	15
Generalizing to Written Materials and to the Gay Community.....	17
Written Materials.....	17
Generalizing to the Gay Community.....	17
Rational Discussion, Emotional Issues.....	19
Final Comments.....	20
References.....	22

¹ Paper prepared for Arvay Findlay (Barristers and Solicitors), representatives for Little Sister’s Book Store and Art Emporium, plaintiffs in *Little Sister’s versus The Queen*, a case heard in the Supreme Court of British Columbia, in November 1994. Dr. Malamuth had been hired by the federal Department of Justice as an expert witness for their defense, in which they argued, in part, that the gay and lesbian sexual material imported and distributed by Little Sister’s was likely to cause significant social harm by fostering violent behaviour in those who consumed it.

STATEMENT OF DR. TED S. PALYS:

COMMENTS ON THE STATEMENT OF DR. NEIL MALAMUTH

QUALIFICATIONS OF DR. PALYS

My current position is that of Associate Professor of Criminology at Simon Fraser University, where I have been employed since 1981. Prior to that, I was Visiting Assistant Professor at the University of British Columbia department of psychology for one year. I completed my Bachelor's and Master's degrees in psychology at the University of Manitoba in 1972 and 1974, respectively, and received my doctorate (in social psychology) from Carleton University in 1979.

My primary expertise is that of a social science research methodologist. My doctoral studies were primarily in that domain, and I have taught dozens of courses in that area at every level from first year university to doctoral study. I have also published in the area — journal articles; chapters in books; presentations at symposia; invited addresses at specialist conferences — with virtually all of my work having an implicit or explicit methodological theme. A textbook I wrote — entitled *Research Decisions: Qualitative and Quantitative Perspectives* (published in 1992 by Harcourt Brace and Company, Canada) — is currently in use in university and college departments of psychology, sociology, business administration, and criminology across the country, and I am currently working on a second edition that will be published in 1996. I have also been engaged as a research consultant at various times with Solicitor General Canada, the federal Department of Justice, and the federal Department of Communications; am occasionally called upon as a reviewer for articles submitted for publication to the *Canadian Journal of Behavioural Science*, the *Canadian Journal of Criminology*, and (some time ago) the *Journal of Personality and Social Psychology*; have also reviewed grant proposals at the request of the Social Sciences and Humanities Research Council; and have been called upon by the media (particularly radio, newspapers, magazines) to comment on both methodological issues (e.g., interpretations of crime statistics and attitudinal data) and other of my substantive interests.

My substantive research areas are diverse, and include a focus on decision-making processes (such as judicial decision-making regarding sentencing and impacts of computer technology on policing), aboriginal justice issues, other human rights and social justice issues, and the pornography area. With respect to pornography (one of the two domains — along with methodology — of greatest relevance to Dr. Malamuth's statement), I have published journal articles; presented papers at symposia; taught several "special topics" seminars at the university level on the topic of philosophical, social, and

legal aspects of the pornography issue, and co-edited a book on the topic (Lowman, Jackson, Palys & Gavigan, 1986). My research in the area includes having undertaken a study of video pornography on behalf of the Special (Fraser) Committee on Pornography and Prostitution; my research report was one of those the Department of Justice chose to publish and distribute in book form as a supplement to the final Fraser Report (Palys, 1984).

COMMENTS ON DR. MALAMUTH'S STATEMENT

Dr. Malamuth is obviously a productive scholar who has achieved much in his career. He is one of the elite (as he accurately points out) among the experimental social psychological research community in the United States, and his laboratory-based research (along with that of two of his colleagues — Dr. Ed Donnerstein of the US, and Dr. James Check of Canada) is required reading for any scholar interested in the area. The general finding they have demonstrated repeatedly over a decade or more is that, under certain conditions, exposure to violent pornography can lead to increased aggressive behaviour. Perhaps because of their findings, the Donnerstein-Malamuth-Check research has been embraced by those who are concerned about the existence of sexually violent images in the media, and particularly by those who have argued for censorship. This includes varied constituencies, such as some members of the feminist community, as well as others with a more traditional “family values” bent.

Dr. Malamuth addresses three questions in his written opinion. These questions, along with his “bottom line” response to each, are summarized in the following table:

Question Addressed	Malamuth Response
(1) “What is the scientific basis for suggesting that obscene materials may be harmful?”	“A series of studies suggest that exposure to some sexually violent media may contribute in some individuals to the development and/or strengthening of attitudes or thought patterns that may be considered harmful” (p.6)
(2) “Is it your opinion that the type of harm discussed under question one may also come from the written word alone?”	“Based on the available literature, it appears that the potential harms discussed above may occur from exposure to the written word alone” (p.23)
(3) “Based on your response to [question one], is it your opinion that homosexual pornography may cause harm even if it is distinct from heterosexual pornography?”	“It is my opinion that homosexual pornography may have harmful effects even if it is distinct in certain ways from heterosexual pornography” (p.31)

In order to fully understand Dr. Malamuth’s assertions, and the particular findings he cites along the way to making them, the reader must also understand the methods by which those findings were reached. I will provide that analysis here by examining (a) the

broader perspective of how Dr. Malamuth appears to define “science”; as well as (b) the conditions in, and procedures by which the findings cited by Dr. Malamuth are generated, and the logic that underlies them.

HIS SCIENTIFIC PERSPECTIVE

The particular approach to research to which Dr. Malamuth subscribes is known as *positivism* or, more recently, as *post-positivism*. The model of research it adopts is borrowed largely from the natural sciences, and one of its central beliefs is that members of the human species can be studied in much the same manner as any other organism. The researcher is given the role of generating knowledge and understanding by theorizing about the world, and then gathering data to test those theories. The ideal method by which to do so, according to post-positivists, is the controlled experiment (e.g., see Palys, 1992).

All social scientists value theory, which involves an attempt to identify important variables and speculate on how they will interact to produce a phenomenon of interest. Post-positivist perspectives on science go on to suggest that it is in the experiment where the viability of these assertions are best tested. The challenge to the experimenter is to contrive a situation where the effects of a single variable can be isolated and observed, in much the same manner that the physicist or chemist might seek to observe a single electron particle or observe a chemical reaction in a contrived situation free of worldly contaminants. Implicit in this view is the idea that the thing being observed may be changed in *magnitude* by its removal and relocation into the laboratory, but not in *character*. The act of giving small electric shocks to a stranger, for example, may be far removed from the brutality of a sexual assault, which the experimentalist acknowledges as easily as anyone else, but the two are nonetheless assumed to be comparable in their essences, i.e., that both are manifestations of the concept “aggression” that differ only in magnitude, but not in kind.

In the case of the “effects literature” cited by Dr. Malamuth, the key “causal” variable of interest is exposure to pornography, and especially violent pornography. In order to assess its effects, the logic of the experiment requires the researcher to create two conditions (in the simplest case) that are identical in all respects except one, that one being the presence or absence of the variable whose effects one wishes to test. Most of the literature involves more complicated designs, but the fundamental principle that drives them is that a comparison needs to be made where two or more groups are equivalent in all respects, on average, on every variable but one.

Theoretical variables are abstractions, however, so that, in order to actually do an experiment, one must give life to the variables by giving them concrete operational meaning in a specific experimental context. One cannot test the abstract concept of “violent pornography”, for example, without choosing a specific example of violent pornography to use in an experiment. So, too, must “exposure” and “aggression” be given an operational existence. And although one may be interested in the relations among

these variables in “people”, one can never test “people” in general, but must decide on a particular sample of people on whom the test will be administered.

It is also important to note that an experiment is a *ceteris paribus* test of the efficacy of a given variable; that is, a test that compares what happens when a variable is present, to when it is absent, under a particular set of *ceteris paribus* conditions (that is, “all else being equal”). But even once the experiment has been performed and a “statistically significant” effect is observed, no claim is made that what has been discovered operates like that in the “real world”. That is because situations in the world are never *ceteris paribus*; any given situation brings together a web of contextual factors that may or may not be similar to those that were present in the experiment. All that one knows is that when “everything else” is held constant or otherwise equalized across groups, such-and-such a variable exerts an effect that is beyond the magnitude of what would be expected on the basis of chance variation alone.

There is a certain irony created by that state of affairs, to the extent that it becomes possible to create knowledge that is true in theory, but misleading in practice. Let me give you an example of this. I recall that as a graduate student, I was always very conscientious about making a good impression in job interviews. It was with interest, therefore, that I read in the impression formation literature in psychology that experiments had shown that people who wear glasses are typically judged to be more sincere, wise, honest and intelligent, *ceteris paribus*, than those who do not wear glasses. More specifically, if you show photographs of people either wearing glasses or not wearing glasses to the typical sample of volunteer undergraduate research participants, and ask them to rate the people in the pictures on myriad different dimensions (like intelligence, sincerity, and honesty), participants will quite reliably rate people with glasses more highly on those dimensions than people who don’t wear them. The nice thing about the experiment was that it offered tremendous control over the situation, such that you could actually have people rate the *same persons* (with and without glasses) to control for other differences, such as attractiveness. Since I wore both glasses and contact lenses at the time (though not at the same time!), I remember thinking that I may as well wear my glasses to job interviews, if that little detail could help me in the interview process.

It was something of a chuckle for me, therefore, when I read some years later that Michael Argyle of Oxford University — a traditional laboratory experimentalist who was beginning to consider the role of context in experimentation — had done an interesting extension to the basic “glasses effect” research. The extension was simply that instead of making “glasses versus no glasses” the *only* piece of information that raters knew about a person, Argyle *began* by gathering the “first impression” data (and showed that he could replicate the well-known “glasses effect”), but then simply let the scene run further, where the people in the picture would start talking about their interests, values and experiences. The end result, not surprisingly, was that now there *was no* “glasses effect”. The “glasses effect” was “true”, in the sense that it could be reliably replicated by any experimenter who wanted to perform the study, but it was only true in particular circumstances within the strict confines of the laboratory, that is, when that was the *only*

information people had on which to base their judgments. But, as soon as there was other information that could be used, then the “glasses effect” was no longer, and what previously had seemed like such an important variable (when it was the only variable under consideration), faded into nothing.

The allegorical excursion above has relevance to the laboratory research that has been constructed to address the “effects” question in the pornography area. To the extent that such experiments have been designed to test purely *theoretical* questions — for example, “Is there any evidence to indicate that exposure to violent pornography *can* influence the propensity of persons to engage in aggressive behaviour?” — then the experiment seems a reasonable method by which to make that determination. But experimenters are frequently tempted to go beyond such questions and begin speculating about what happens outside the lab; that is, when the males in this country who consume pornography are exposed to what is available, will they have an increased likelihood of engaging in aggressive behaviour?

That attention to application changes the question considerably. Instead of: “*Can* exposure to violent pornography make a difference?”, one is now asking “*Does* it?”. And, as the “glasses effect” example showed, it is not inconsistent or implausible for the two answers to those questions to be “Yes, it *can*” but “No, it *doesn't*”. Many authors (e.g., Manicas & Secord, 1988; Palys & Lowman, 1984) have suggested that while the purely theoretical “*can*” questions are eminently suitable to experimental analysis, the experiment may *not* be a particularly useful launching point for the “*does*” questions. The issue in part is one of “ecological validity”, that is, how well do the conditions of the experiment simulate the important elements of the context to which you wish to generalize (e.g., see Palys, 1978; Palys & Lowman, 1984)? To examine this issue in the pornography area, let us begin by examining some “paradigmatic” (i.e., prototypical; exemplary) study in this area in greater detail.

A Paradigmatic Effects Study

A “typical” design in the “effects” area was offered by Donnerstein & Berkowitz (1981), both of whose eminence in experimental social psychological circles equals Dr. Malamuth’s. The three are indeed valued colleagues of one another, engage in similar research emanating from a similar epistemological perspective, and have often published together. The Donnerstein and Berkowitz (1981) study may be considered “paradigmatic” insofar as its operational choices have been replicated dozens of times in subsequent research, and provide the standard against which subsequent research is often judged. Malamuth cites the study both frequently and favourably.

The research participants in Donnerstein and Berkowitz (1981) were all male-undergraduate-introductory-psychology-student-volunteers, as is true of most studies in this area. Upon showing up for his appointment at the lab, the research participant would find that another person (a woman) also had an appointment, and that the two of them would be participating in the study together. The experimenter explains (through tape-

recorded instructions) that one of them will have an opportunity to be a “learner” whose job is to try and remember certain word pairs, while the job of the other will be to assist the experimenter. An allegedly random draw is held to determine who will play each role. Unbeknownst to the male participant, however, the draw is actually “rigged”, so that it is always the woman who becomes the learner, while the male is always in the position of assistant to the experimenter. The woman, as the reader may suspect, is actually an employee of the experimenter rather than a “real” research participant, and she is trained to respond in the same, pre-programmed manner every time the experiment is run.

With their roles determined, the experimenter next states that the woman will be given some time to study the word pairings before a “test” is given. The man, in the interim, is to spend his time writing a brief essay concerning the possible legalization of marijuana. Once he has finished, the woman is brought back into the setting and is supposed to evaluate the essay. She is on the other side of a partition, however, and is not supposed to communicate directly with the man. Instead, she communicates indirectly via written note, and through the delivery of some electric shocks via finger electrodes placed on the man’s hand. She is most insulting, stating in writing that the essay is terrible. When faced with the choice of how many electrical shocks to deliver to the man, she delivers nine shocks out of a possible maximum of ten. This little interchange serves two purposes in terms of the experimental goals. First, it helps to reaffirm the “reality” of the electrical shocks to the male participant. This is important because the man will soon have an opportunity to deliver electrical shocks to the woman, and the experimenter needs the man to believe that any shocks he delivers are real. Second, this interchange is known among “effects” researchers as the “anger manipulation”, and has become a virtual requirement of effects testing, since it seems that if the female does not first anger the man, no effects of exposure to violent pornography are observed (e.g., see Donnerstein & Berkowitz, 1981).

After the anger manipulation is performed, the woman is allowed further time to study. Because this studying will take some time, and the male participant now has nothing to do, the experimenter notes something along the lines of: “By the way, there’s a friend of mine down the hall who is preparing some film clips for another experiment, and he needs people to make some ratings of them. Since we have some time to kill, would you be interested in going down the hall and helping him out for a few minutes?” Virtually all participants agree to do so.

It is at this point that the exposure manipulation occurs. Males are randomly assigned to one of four experimental conditions, where the only difference between conditions is over the type of film clip to which each is exposed. In the Donnerstein and Berkowitz (1981) study, two of these clips portrayed (1) a non-sexual and non-violent clip of a talk show; and (2) a sexually explicit but non-violent depiction of a man and a woman engaging in mutually consenting intercourse. The other two film clips both involved a scene in which a woman coed is seen studying with two male students, who then begin to make sexual advances. She resists, but is raped. The difference between the 3rd and 4th films is not in their visual content (which is identical), but in the voice soundtrack: (3) in one version,

the woman begins by protesting, but soon begins to enjoy the process (a rape myth depiction; the “positive outcome” condition); while (4) in the second version, the woman begins by resisting, and continues to resist throughout the process, experiencing all the horror of a sexual assault she is powerless to stop (the “negative outcome” condition).

After viewing one of these four film clips and completing a few rating scales (consistent with the cover story that had been offered), the male participant returns to the first experimenter, who is now ready to receive him. The woman has completed her studying, and has some electrodes attached to her fingers. The male adopts the role of assistant to the experimenter, and begins the assessment of whether the woman has remembered the list of word pairings she has been studying. Whenever she makes a mistake, it is the male participant’s role to determine how many electrical shocks she should receive, and to deliver them. Of particular interest to the experimenter is the average number of electrical shocks that the male participants deliver, and how that number varies, depending on the type of film clip viewed immediately prior to participating in this portion of the experiment.

In sum, the four groups are identical to begin with (because Participants are assigned on a purely random basis to one of the four conditions, such that we have no reason to believe that, prior to their film exposure, one group as a whole is any different from any other group as a whole). Participants are then treated identically in all respects except for one, that one difference being the type of film clip to which they are exposed. The assumption is that if the groups are all equal to begin with, any subsequent differences in average shocking frequency between the groups must be due to the one variable on which their experience varied, that is, the type of film to which they were exposed. Donnerstein and Berkowitz (1981) found that, after being exposed to the anger manipulation, the two groups who viewed the sexually violent film clip (regardless of whether it was accompanied by a “positive outcome” or “negative outcome” soundtrack) administered a higher level of average shock frequency than the two groups of participants who viewed either the sexually explicit (but non-violent) depiction, or the non-sexual/non-violent talk show clip.

The Donnerstein and Berkowitz (1981) experiment is typical of research in the “effects” literature, and, from a purely technical point of view, is also very well-designed. It meets all the criteria of “good experimentation” espoused by experimental social psychologists (e.g., see Aronson & Carlsmith, 1968; Festinger & Katz, 1953; Rosenthal & Rosnow, 1984). Not surprisingly, therefore, research of this type has been published in many of the most prestigious journals that publish experimental research of this type, and researchers such as Donnerstein and Malamuth have been lionized by professional organizations dominated by persons who share their epistemological assumptions.

Taken collectively, the literature shows that, under certain unique conditions, male undergraduate student volunteers will deliver stronger average electrical shocks to the fingertips of women after they have been exposed to violent video pornography than when they have been exposed to either sexually explicit but non-violent material, or

neutral (non-sexual, non-violent) videos. The main question to be considered here is, “Should the court be convinced, on the basis of this evidence, that the existence of harms following exposure to violent pornography has been demonstrated?” This question can be addressed by more careful consideration of how the “rules of social psychological laboratory experimentation”, when followed in as exemplary a fashion as has been the case in the Donnerstein & Berkowitz (1981) research, can create a distorted, myopic, and misleading set of findings when these are used to make predictions about the world outside the laboratory. In sum, although the “effects” researchers have indeed given us a little bit of knowledge, they may also exemplify the adage that “A little bit of knowledge can be a dangerous thing”.

The Sample: Representativeness Issues

Although their theoretical interests are in the effects of exposure to violent pornography on “people”, the specific research participants who participated in Donnerstein & Berkowitz (1981), as is the case in virtually all the studies cited by Malamuth, were male, introductory psychology students who volunteered to take part in the study. There are two matters to consider with respect to this choice of sample. The first, which is the most often cited by critics of the “effects” literature (e.g., see Byrne & Kelley, 1989; Fisher, 1986), concerns the representativeness or “typicality” of such a sample. The 18 to 21-year old male introductory psychology student volunteer is obviously anything but a “representative sample” of the general population, so one question that might be posed is whether *any* result that is obtained with such a sample can be generalized to the broader population of interest (i.e., all people, or even all males).

Personally, I do *not* see that criticism as particularly problematic here, for two reasons. First, many “effects” researchers are interested in isolating general behavioural principles that are true of everyone; and if that is so, then male introductory psychology students are as good a sample as anyone else. Just as a physician can adequately test your blood by taking only one sample from an isolated part of your body (your finger, for example) because the blood in your finger is the same as the blood everywhere else, the male introductory psychology student volunteer is not an unreasonable choice *if* the principles that govern their behaviour are the same as the principles that govern everyone else’s.

But even if we acknowledge that students are *not* typical of all people, so that the question of who is sampled is an important one, a criticism of “unrepresentativeness” is empty unless we go on to identify the sampling bias that exists, and consider what the implications of that bias might be for the results that are obtained. In that regard, we can expect several differences between male undergraduate psychology students and the general male population. Not only have such students received more extensive formal education than the average population member, but we might also speculate that they might be less likely to use “physical” means to solve conflicts or achieve goals, and more introspective about their behaviour and the motives underlying it, for example. But if that is so, then it might be argued that these people may be even *less* likely than the general population to engage in aggressive and assaultive behaviour. Indeed, the upshot of the

“unrepresentativeness” of the sample is that *if* it can be demonstrated that even such a relatively well educated and literate group can be affected by exposure to violent pornography, then it may well be that this is, if anything, a *conservative* estimate of the extent to which such effects will exist in the population as a whole (Malamuth, 1989, makes a similar argument; see p.183).

The Sample: Variability Issues

Although I am thus reserved in my criticism of the “effects” literature from the perspective of the representativeness in sampling procedures, there is another way that the choice of an all-male-introductory-psychology-student-volunteer sample can actually be quite *unfair* if we are actually seeking a “fair” test of the hypothesis that exposure to violent pornography influences the likelihood of subsequent aggressive behaviour. Indeed, the criticism here is that, rather than “testing” the proposition, the sampling choices of the investigators help to virtually “create” the effect in a self-fulfilling manner.

In order to understand this criticism, one must appreciate that, when experimenters try and test a theoretical proposition, the “success” or “failure” of the test is seen as refereed by the rules of probability theory in the “test of significance” that is performed on the data. In the effects literature, this test comes down to a comparison of whether the average level of aggressiveness evinced by persons who were exposed to violent pornography (called the “experimental” group because they receive the experimental manipulation whose effect is being tested) is “significantly” different from (i.e., more aggressive than) those in a comparison group who are equivalent in every respect to the experimental group except that they did *not* get exposed to the pornography.

It is not, however, just a matter of simply comparing the means (i.e., averages) of the two groups. Instead, probability theory recognizes that whether a given difference is worth getting excited about (that is, whether it is “statistically significant”) depends also on such matters as the amount of variation that exists naturally within the group(s) under study. Indeed, the most common statistical test used in experimental research is the analysis of variance; this procedure determines whether results are “statistically significant” through computation of an F-ratio that explicitly compares the degree of variation *between* groups (that is, the difference between experimental and comparison group means) to the average degree of variation *within* the groups.

To understand the logic of this process, let’s say that we are interested in testing whether a certain educational programme improves academic performance, and that the “reality” of the situation is that the technique can cause improvements of up to five percentage points. Is a five per cent improvement “significant”? The answer depends in part on whom we test. If the group we test has minimal variability among them — let’s say that they are all high achievers from the previous semester’s Dean’s List, so that their grade point averages are in the relatively limited range of 85 to 95 per cent — then any technique that can cause a 5 per cent increase in test scores is indeed very significant, and something on which we might want to expend funds. If, on the other hand, the group

under consideration is the whole undergraduate student body — where grade point averages might be anywhere from 40 to 95 per cent — then increases of 5 per cent are fairly trivial, and hardly worth getting excited about or expending funds for.

Analogously, we can see that the choice of Malamuth and his associates to rely virtually exclusively on the easily accessible male undergraduate introductory psychology student volunteer, may well help to *create* the very effects that the experiments are designed ostensibly to *test*. The crucial issue is not the “representativeness” of their samples *per se*, but of the natural “variability” that exists within these samples. In this regard, it is clearly the case that male undergraduate introductory psychology student volunteers evince a restricted range (relative to the general population) on almost any variable that one might consider, such that changes that might be observed following exposure to a stimulus like a violent pornographic depiction that would be considered statistically trivial if they occurred among the broader population, all of a sudden appear to be of greater magnitude when the group being tested possesses only a sliver of the variation that exists in the broader population.

The Setting: Ecological Validity Issues

As I have outlined above, the laboratory experiment involves the creation of a contrived setting designed to provide a “pure” test of a theoretical proposition among a designated sample of research participants. In the preceding section, I suggested that we should be concerned that the laboratory research in the area samples no one other than male-undergraduate-introductory-psychology-student-volunteers, both because of (1) their lack of representativeness of the broader population; and (2) use of such a restricted base can give a magnified role to factors that might otherwise be trivial. The same concerns can be brought to bear with respect to research design issues, where, in the interests of alleged experimental purity, researchers create situations that (1) are only obliquely related to the situations in the world they are ostensibly trying to help us understand; and (2) create misleading results because of the manner in which they virtually *create* the effects they are allegedly trying to *test*. These considerations have been embellished in greater detail elsewhere (see Palys, 1989; Palys & Lowman, 1984), but deserve some attention here.

Laboratory experimentalists are taught several basic principles that are to be used in the design of an experiment. These include:

- (a) **maximizing between-groups variation** (that is, the degree of difference between experimental and control groups) so as to make as clear a differentiation between conditions as possible. In the effects literature, this means choosing the most gruesome “violent pornography” one can find, choosing sexually explicit material that is extremely explicit and not at all violent, and making the “neutral” material as devoid of sexual and violent content as one can. Recall that pushing the groups as far apart as possible (by choosing film stimuli that are as different as possible) has the effect of making

the *numerator* of the F-ratio as *big* as possible, thereby increasing the likelihood of determining a statistically significant result;

- (b) **minimizing within-groups variation** (that is, the degree of “natural” or “random” variation that exists within the groups being tested) so as to make as “sensitive” a test of the hypothesis as possible. In the effects literature, this is accomplished by using relatively homogenous populations of respondents (e.g., the male-undergraduate-introductory-psychology-student-volunteer), and by attempting to exert experimenter control over the situation to standardize conditions as much as possible (e.g., by using tape-recorded instructions to minimize any variations in the reading of the instructions that would happen over time if the experimenter were reading the same instructions again and again; by having the woman confederate send her insulting statement about the male participant’s essay via written note so that the message is exactly constant across all tested groups). Recall that this has the effect of making the *denominator* of the F-ratio as *small* as possible, thereby increasing the likelihood of observing a statistically significant result; and
- (c) by **controlling all available response alternatives** so that any impetus to behaviour that is created by the conditions of the experiment are harnessed in the service of the dependent variable. Underlying this view is something of a “hydraulic” model of human behaviour where the belief is that, if a stimulus can effectively energize a behavioural response, then the astute experimenter will “dam up” all behavioural alternatives but one, so that the magnitude of all behavioural impulses will be visible in the chosen place. In the context of the effects literature, where the interest has been in determining aggressive impulses, this has meant that participants have only been given one way to express themselves, and that is by delivering electrical shocks to the woman on the other side of the partition.

All of these principles have been employed in the research cited by Malamuth in his opinion for the court. Once again, they are perceived to be principles of “good experimentation” (e.g., Aronson & Carlsmith, 1968; Festinger & Katz, 1953; Kerlinger, 1973; Rosenthal & Rosnow, 1984), and, they do make great sense *if* one is addressing a purely theoretical question of interest, where the central issue for the researcher is whether a theoretically hypothesized relationship *can* be given empirical life. Certainly it accounts for Malamuth’s (1989) chastizing comment toward Fisher and Grenier (1988), to the effect that their replication was unsuccessful because it failed to adequately “maximize the opportunity of detecting any effects that might exist” (p.185).

But my understanding is that the court is not interested solely in questions of theory, but also in their implications, if any, for practice. The interest of the court, as I understand it, is not “*Can* exposure to aggressive pornography increase the likelihood of subsequent harm?”, but, rather, “*Does* it?”. With that goal in mind, we can begin to question whether, in the interest of living up to the experimentalist principles of his discipline, Dr.

Malamuth and his colleagues have made the situation into something *other than* that which they began ostensibly to investigate.

The problems here are several. For example:

- (a) In order to control the experimental setting, “effects” researchers have a single research participant watching the experimental videos *alone*, contrary to the apparently more usual situation in the world where sexually explicit material of this type is consumed in *social* settings (e.g., stag parties; among couples) where interaction about the artificiality or absurdity of video scenarios is commonplace (see preliminary data in Palys, 1984).
- (b) While the aggressive behaviour that is of interest to us is *severe* (e.g., violence in the form of sexual assault and other sexual abuses), ethical requirements constrain laboratory investigation to behaviour that is *trivial* (small electric shocks that cause minimal pain and no damage), and hence probably result in an overestimate in the extent to which people are prepared to engage in such behaviour;
- (c) While aggressive behaviour is *discouraged* in society (e.g., one is subject to possible arrest and imprisonment for behaviour such as sexual assault), the experiment requires it to be *encouraged* in the experimental setting by an experimenter who, as an authority figure, represents the interests of legitimate science;
- (d) While we have many *alternatives* about how to respond in the real world when we are angered (e.g., by vacating the situation, by talking to the person who has angered us), the “effects” researchers offer their research participants *only one way to communicate* — via the imposition of electric shock to the fingertips of the woman on the other side of the partition.

Whether singly or in combination, the combined effect of these compromises to experimental requirements is that we change the very nature of what we are investigating in the process of achieving manipulative control, that is, making the situation into something that normally it is not. Are these shifts non-problematic? Or do we “create” effects that would not otherwise exist?

It is noteworthy that at least two studies have investigated the impact of some of the above factors. Regarding “caveat (d)”, for example, Fisher and Grenier (1994) wondered what would happen if research participants were given a broader array of response alternatives than just delivering shocks. They discovered that, when response alternatives were given, the vast majority of respondents chose to *talk* to the woman who had angered them — regardless of prior exposure condition — rather than to deliver electric shocks.

Regarding “caveat (c)” above, relevant evidence has been supplied by none other than Malamuth (1978). He describes this research in Malamuth (1984) as follows:

Following exposure to these [visual] stimuli, all subjects were insulted by a female confederate and then were placed in a situation where they could aggress against her via the ostensible delivery of electric shocks under one of two assessment conditions. Half of the subjects were assigned to read a communication that suggested it was “permissible” to behave as aggressively as they wished (disinhibitory communication); the other half were given a communication designed to make them somewhat self-conscious about aggressing (inhibitory communication). ...

The results revealed no significant differences in aggression following the inhibitory communication. (p.35).

In sum, as soon as conditions in the experiment start to better approximate those in the “real world” — that is, when people have alternative ways of responding, and where aggressive responding is discouraged — the findings of the laboratory based research appear to vanish, and no effects are found.

The Fragility of Media Influence Effects

It seems fairly clear from the literature that any “media effects” that can be attributed to the message or content of violent pornography *per se* is, in the grander scheme of things, a fairly trivial influence. Even Malamuth, who has built his reputation on the demonstration of these effects, has recognized that the media may not even be a particularly important element in the behavioural equation. His “indirect effects” model, for example, affirms that

Individual conditions and the broader social climate are postulated as the originating environmental influences on the individual. The mass media are considered one of the many social forces that may, in interaction with a variety of other cultural and individual factors, affect the development of intermediate attributes, such as thought patterns, sexual arousal patterns, motivations, emotions, and personality characteristics. (Malamuth, 1989, p.163).

Later, he concludes by stating that,

As with many behaviours, it is apparent that antisocial behaviour against women is a function of many interacting causal factors. It is very difficult to gauge the relative influence, if any, of media exposure alone. However, by itself, it is likely to exert a small influence, if any. (Malamuth, 1989, p.198).

Long-Term Effects

Although it is indeed the case that “effects” researchers such as Malamuth and his colleagues have shown that, under certain conditions, exposure to violent pornographic media *can* increase subsequent aggressiveness, the “certain conditions” under which these effects operate are in fact precariously constructed and of dubious ecological validity. One might expect, therefore, that any attempts to demonstrate “long term” effects to such an ephemeral phenomenon would inevitably be doomed to failure.

This issue was, in fact, addressed by Malamuth and Ceniti (1986), who sought “to investigate the relatively long-term effects of repeated exposure to violent and non-violent pornography on males’ laboratory aggression against women.” (p.131). Parenthetically, it might be noted that “long term” in this context meant “over a four-week period”, while “repeated exposure” was operationally defined as 10 exposures to assorted feature films, short videos, magazine pictorials, and written material. The study revealed *no* evidence of any effect of exposure — even repeatedly over a one-month period — when the male-undergraduate-university-student-volunteers returned to participate in a supposedly “unrelated” aggression study:

The results did not reveal that repeated exposure to violent or non-violent pornography had any significant effect on laboratory aggression against women. These findings appear to be inconsistent with previous data showing that exposure to violent pornography may increase males’ laboratory aggression toward women. ... The most apparent explanation for this discrepancy is that earlier investigations examined immediate effects (i.e., in the same session that exposures were presented) whereas the present experiment tested for relatively long-term effects. It may be that exposure to violent pornography might have an immediate impact on aggressive behaviour against women but this effect may dissipate quickly over time. (Malamuth & Ceniti, 1986, p.135)

Although the fragility of the media effect is re-affirmed, a positive element of the experiment was that it allowed the simultaneous testing of the media effect (a situational attribute) with an individual difference (“personality”) measure that Malamuth (1989) describes as a self-report based measure that he entitles “likelihood to rape” (or LR). Although the LR measure may itself have problems, it is interesting to note that the LR scores actually predicted laboratory aggression against the woman confederate better than did the exposure variable (i.e., whether the person had been exposed to four weeks of violent, or non-violent, pornography).

Other Variables

Although my considerations above have focussed on the evidence that Malamuth has compiled in his opinion for the court, a variety of other variables might also be considered. In particular, I would suggest that more attention be paid to the *meaning* of

these images to people, and, particularly, the role that culture and individual differences can play in the generation and interpretation of images.

Ironically, a particularly provocative article that offers a powerful counter-argument to the idea that there might be any effects to violent pornography *per se* is in a book edited by Malamuth & Donnerstein (1984). In it, Abramson and Hayashi (1984) offer a comparative analysis of Japanese and American pornography, attitudes about sex, the role of sex in the media, and so forth. They note, for example, that Japanese laws and mores completely prohibit images that most North Americans would find tame (e.g., neither pubic hair nor adult genitalia may be shown), while allowing a variety of images that many North Americans would find horrendous (e.g., the admonition against showing pubic hair has resulted in the proliferation of images of pre-pubescent girls):

Of particular note in Japanese pornography (film and novel) is the recurring theme of bondage and rape. Although movies are much less explicit than their American and European counterparts, the plot often involves the rape of a high school girl. ...In fact, one of the best ways to ensure the success of a Japanese adult film is to include the bondage and rape of a young woman. This juxtaposition of sexuality and aggression is evident in almost all forms of Japanese sexual material, including cartoons, films, and sexological museums. (Abramson & Hayashi, 1984, p.178)

If sexually violent themes/images in and of themselves somehow “caused” greater aggressiveness, then clearly sexual aggression should be rampant in Japan, and at higher levels than in the United States. But despite the pervasiveness of such sexually violent material, rates of sexual assault and/or other forms of sexual abuse are far less than in the West.

In comparison to Western nations, Japan has a substantially lower incidence of rape: in the United States there are 34.5 reported rapes per 100,000 population; in England, 10.1; in West Germany, 10.7; in France, 3.2; and, in Japan, 2.4. ... The discrepancy in the incidence between the United States and Japan cannot be attributed to variance in the laws because the laws are basically the same (although prosecution rates may vary).

If there is a direct connection between the prevalence of rape imagery and rape behaviour, Japan should have an overwhelming occurrence of rape. As indicated in the preceding paragraph's rape statistics, it does not. Consequently, it is our suggestion that mediating circumstances are involved, especially in the form of internal constraints to maladaptive behaviour. (Abramson and Hayashi, 1984, p.181)

These observations of the importance of culture in the rules of sexual practice and the interpretation of sexual and sexually violent images are particularly germane to the case

now before the court regarding Little Sister's Book Store and Art Emporium. Since it is clearly the case that Little Sister's caters primarily to the homosexual community, then a question of interest to the court concerns the applicability of any of the above literature to that community.

Generalizing to Written Materials and to the Gay Community

Written Materials

Dr. Malamuth asserts that "the potential harms [he discusses] may occur from exposure to the written word alone" (see p.23 of his opinion), and bases this prediction (since he states there are no data that reflect on the issue) on his belief that the important element is the *message communicated*, rather than the medium of communication.

The main problem with this view lies in Dr. Malamuth's apparently enduring belief that messages have some objective existence that can be communicated uniformly to all. I doubt whether there is any message that meets that criterion. Stating that the important element is the *message communicated* ignores any consideration of the person receiving the message, any attribution of meaning to its content, and any situational and/or cultural factors that might enter in to its interpretation. To say that it is above all other factors in importance is for Dr. Malamuth to ignore his own research.

For my own part, several differences between written and oral media immediately suggest themselves. First, while video pornography is often consumed in social settings (see Palys, 1984), written pornography (books, magazines) is more likely to be consumed by one individual at a time. Second, it does not require literacy to watch a video, but it does to read a book. In both cases, one could envision these variables (i.e., solo versus social consumption; literacy requirements and the propensity to read books will be related to educational attainment and cultural differences) making a difference in who seeks them out, and the range of interpretations that are made by that selection of people.

Generalizing to the Gay Community

It is actually rather ironic that Dr. Malamuth does not mention homosexual pornography until page 31 of his 33-page written opinion to the court, and that it has taken me up to page 17 of my own opinion to do the same. That in itself is a good indication of how little has been written about such issues within the gay community. When considering what the effects of exposure to homosexual pornography might be, we are clearly in the realm of speculation.

Dr. Malamuth suggests that any effects that are observed for heterosexual pornography among the heterosexual community will in all likelihood be the same for homosexual pornography among the homosexual community. Although he has no direct data on this

issue, he bases his speculation on the answers to three questions: (1) are the messages in homosexual pornography basically the same as heterosexual pornography?; (2) are the minds of homosexuals basically the same as the minds of heterosexuals?; and (3) are there problems of sexual conflict within the homosexual community?. In all three cases, he answers “yes”, and hence concludes that the same processes should prevail among the homosexual community as among the heterosexual one.

Although Malamuth is correct in saying there are no data that bear directly on this issue, the most closely related finding we know of (concerning male-male aggression following exposure to violent pornography) contradicts his opinion. Summarizing the findings from his and Donnerstein’s laboratory research, Malamuth (1984) notes that

The data show that exposure of male subjects to aggressive pornography increases aggressive behaviour against *female* but not male targets. (p.35)

Notwithstanding that result, the problem with Dr. Malamuth’s responses in my mind is his tenacious belief in the objective qualities of messages. Thus, for questions #1 (Is homo/heterosexual pornography basically the same?) and #2 (Are homo/heterosexual minds basically the same?), Dr. Malamuth pronounced the material similar, and suggests that homosexuals minds are no different, in terms of basic functioning, than heterosexual minds. But Dr. Malamuth misses the point. Although the superficial content that appears in some heterosexual and some homosexual pornography may be similar, the *meanings* associated with those images, and hence their relationship with behaviour, may be wildly different. In any case, Dr. Malamuth is also not particularly qualified to make such a judgment. The issue is not whether Dr. Malamuth is homosexual or heterosexual, but whether he has ever made any sustained effort to understand pornography from the perspective of the consumer, and he has not.

In considering possible effects, one has to keep in mind that although homosexuality is not inherently better or worse than heterosexuality as a sexuality, the historical experience of these two groups regarding their sexual activity has been significantly different. While heterosexuals have enjoyed feeling “normal” about their sexuality, the gay community has endured many years of being considered “deviant” and/or “unnatural”. Being homosexual has been an unwarranted source of stigma for many years; many gay people still feel reluctant to “come out of the closet”; and, as an oppressed lifestyle (e.g., through institutional harassment; “gay bashing”), there is probably more of a shared sense of community and interdependency among homosexuals than among the heterosexual community.

When it comes to matters of sexual violence, the situation in homosexual relations is unique to the extent that it might be more inherently egalitarian because of the gender similarity of the two people involved (e.g., see Brock & Kinsman, 1986). Taken together, all these differences between homosexual and heterosexual life experience would leave me astounded if the two communities did not attach different meanings to sexual practices and sexual images. There is clearly a requirement for more research in that area,

particularly by gay researchers, to articulate these issues. In the interim, I would suggest that the safer course would be to assume on the basis of their significantly different social history that differences in meaning do exist, until shown otherwise.

As for the third question Malamuth poses, I agree that the gay community, as is the case with all other communities, is not immune to sexual violence. But to assume that the same dynamics must therefore prevail seems inconsistent with the feminist literature that Malamuth states informs his analysis. For example, to the extent that sexual violence involves not only violence but *gender* violence, embedded in a history of patriarchal relations, then how could patterns of sexual violence among same-sex partners (who are thus equal in gender status) be a product of the same dynamic? Overall, the biggest threats to homosexuals involving violence probably involve people from outside the gay community (harassment from the intolerant; gay bashers), than from inside it.

RATIONAL DISCUSSION, EMOTIONAL ISSUES

Sexual issues can be highly charged with emotional debate, such that the decision-making processes that people go through can sometimes seem considerably less than rational. Researchers often like to think that people (1) listen to their data; (2) evaluate the data; and then (3) integrate those data into their existing opinions, whether that involves either reaffirming or changing them. One often sees people in the sexual domain allegedly following that path, but in fact doing something quite different. Indeed, it seems the order is often reversed — people do not read the data, evaluate it, and then form an opinion; instead, they have an opinion, read the data, and then accept or reject it depending on the extent to which it conforms to their opinion (e.g., see Byrne & Kelley, 1989). The research is evaluated as “good” (or “bad”) on the basis of whether they like (or don’t like) what it says.

As Byrne & Kelley (1989) stated, “judgments are routinely made on inadequate bases and the utilization of sophisticated research methodology, statistics, and professional jargon can sometimes serve only to camouflage the inadequacies of the decision-making process rather than to improve on them. ... Sexuality seems to be especially vulnerable to this problem.” (p.364). The authors note how data that would often be considered ridiculously poor can all of a sudden take on new credibility if they reveal something consistent with our biases. A good example involves data of the “case history” sort, where new calls for censorship often follow the revelation that “pornographic material” was found in the apartment of some serial murderer or convicted rapist.

Without belabouring the obvious, the general problem of selecting some small aspect of the total array of data and the need for a control group can perhaps be illustrated by a comparison. How convincing would it be to read that police officers discovered homogenized milk in the refrigerator of a convicted rapist? Without knowing whether this activity differentiated rapists from a matched sample of non-rapists, few of us would be inclined to propose laws banning dairies, forbidding milk distribution, and

directing the removal of this product from refrigerated grocery shelves.
(Byrne & Kelley, 1989, p.372).

As this suggests, the interpretation and ‘spin’ that we place on sex-related data will often be a reflection of the biases we have about certain sexual material. For example, in his written opinion for the court, Malamuth notes research by Zillmann and Bryant, in which male and female university students were given “massive doses” of explicitly sexual material, and then asked them to estimate the proportion of adults who engaged in various sexual practices, such as fellatio, cunnilingus, and anal intercourse. By way of summary conclusion, Malamuth notes that, overall, “the data indicated that subjects’ perceptions of the prevalence of sexual practices were affected by the amount of pornography they viewed.” (p.22). The statement is made as if it were some sort of self-evident statement regarding the negativity of such a finding. Zillmann and Bryant (1984), the authors of the original research to which Malamuth refers, put the same data in a somewhat different light:

Massive exposure to pornography did not necessarily distort the perception of prevalent sexual practices. Estimates of the use of oral-genital stimulation, for instance, approximate estimates based on survey data more closely the more subjects were exposed to sexually explicit materials. Massive exposure to pornography thus could be said to correct distorted views of sexuality. (p.132).

The ‘spin’ that is given to homosexual pornography and homosexual matters are particularly likely to be distorted by the biasing processes noted above, especially given the apparently widespread homophobia that exists in Canadian society today. While there may be *other* reasons for deciding that certain homosexual (or heterosexual) depictions are obscene (e.g., because individual depictions violate community standards of tolerance), my opinion is that the effects literature summarized by Malamuth in his written opinion to the court does *not* provide adequate data on which to base any conclusions.

FINAL COMMENTS

The above represents my commentary on the “effects” research outlined by Malamuth in his written opinion to the court. The following summarizes my own views concerning the three questions posed to Malamuth by the Department of Justice (compare to his “bottomline responses listed on p.3),

Question Addressed	Palys Response
(1) “What is the scientific	The laboratory data cited by Malamuth have shown that

<p>basis for suggesting that obscene materials may be harmful?”</p>	<p>violent pornography <i>can</i> enhance aggressiveness by some males against females under certain conditions, but whether it <i>does</i> beyond the confines of the laboratory, is questionable. Even in the lab, the better job the scene does of approximating reality, the more likely the effect will disappear. Overall, one gets the impression that there are few if any straightforward effects to media <i>per se</i>, and much more to be accounted for by looking at the <i>meanings</i> that those depictions have for different people, as influenced by such factors as personal history and cultural understandings.</p>
<p>(2) “Is it your opinion that the type of harm discussed under question one may also come from the written word alone?”</p>	<p>The laboratory research cited by Dr. Malamuth has been created primarily to identify any harmful effects to pornographic film, video and pictorial media. My analysis suggests that the findings, although consistent, are also fragile and generally dissipate into trivial once taken beyond the controlled confines of the laboratory. My speculation would be that the experience of consuming written pornography would in all likelihood be different than video pornography, if only because written consumption is a personal experience while video pornography is more often social, and that there would also be some selection bias, since those persons who consume books and magazines are probably more literate than those who do not.</p>
<p>(3) “Based on your response to [question one], is it your opinion that homosexual pornography may cause harm even if it is distinct from heterosexual pornography?”</p>	<p>Dr. Malamuth speculates that the same effects will arise among the gay community with homosexual pornography that exists among the heterosexual community with heterosexual pornography. This response completely ignores the feminist analysis that has allegedly guided his work (where the deleterious aspects of pornographic images are related to the reproduction of patriarchal relations in the sex act, and these gender relations are missing when same-sex partners are involved). It also ignores the different experience of homosexuals as an oppressed sexual minority, which — because the very basis of their oppression is in the definition of sexuality — will in all likelihood be associated with different images holding different meanings for homosexual and heterosexual viewers.</p>

REFERENCES

- Abramson, P.R., and Hayashi, H. (1984). Pornography in Japan: Cross-cultural and theoretical considerations. In N.M.Malamuth, & E.Donnerstein, (Eds.). *Pornography and Sexual Aggression*. New York: Academic Press, pp.173-185.
- Aronson, E., and Carlsmith, J.M. (1968). Experimentation in social psychology. In G.Lindsey and E.Aronson (Eds.) *The Handbook of Social Psychology* (2nd edition), Vol. II. Reading, Mass.: Addison-Wesley, pp.1-79.
- Brock, D.R., & Kinsman, G. (1986). Patriarchal relations ignored: An analysis and critique of the Badgley report on sexual offenses against children and youths. In J.Lowman, M.A.Jackson, T.S.Palys, & S.Gavigan (Eds.) (1986). *Regulating Sex: An Anthology of Commentaries on the Findings and Recommendations of the Badgley and Fraser Reports*. Burnaby, BC; Simon Fraser University, pp. 107-126.
- Byrne, D. & Kelley, K. (1989). Basing legislative action on research data: Prejudice, prudence, and empirical limitations. In D.Zillmann & J.Bryant (Eds.) *Pornography: Research Advances and Policy Considerations*. Hillsdale, N.J.: Lawrence Erlbaum, pp.363-385.
- Donnerstein, E. & Berkowitz, L. (1981). Victim reactions in aggressive erotic films as a factor in violence against women. *Journal of Personality and Social Psychology*, 41(4), 710-724.
- Festinger, L., and Katz, D. (Eds.) (1953). *Research Methods in the Behavioural Sciences*. NY: Holt, Rinehart & Winston.
- Fisher, W.A. (1986). The emperor has no clothes: On the Fraser and Badgley Committee's rejection of social science research on pornography. In J.Lowman, M.A.Jackson, T.S.Palys, & S.Gavigan (Eds.) (1986). *Regulating Sex: An Anthology of Commentaries on the Findings and Recommendations of the Badgley and Fraser Reports*. Burnaby, BC; Simon Fraser University, pp. 159-176.
- Fisher, W.A., and Grenier, G. (1994). Violent pornography, antiwoman thoughts, and antiwoman acts: In search of reliable effects. *The Journal of Sex Research*, 31, 23-38.
- Kerlinger, F.N. (1973). *Foundations of Behavioural Research*. (2nd Edition). New York: Holt, Rinehart & Winston.

- Lowman, J., Jackson, M.A., Palys, T.S., & Gavigan, S. (Eds.) (1986). *Regulating Sex: An Anthology of Commentaries on the Findings and Recommendations of the Badgley and Fraser Reports*. Burnaby, BC; Simon Fraser University.
- Malamuth, N.M. (1978). *Erotica, aggression and perceived appropriateness*. Paper presented at the 86th annual convention of the American Psychological Association, Toronto, Ontario. [Cited in Malamuth, 1984].
- Malamuth N.M. (1984). Aggression against women. In N.M.Malamuth, & E.Donnerstein, (Eds.). *Pornography and Sexual Aggression*. New York: Academic Press, pp.173-185.
- Malamuth, N.M. (1989). Sexually violent media, thought patterns, and antisocial behaviour. In G. Comstock (Ed.). *Public Communication and Behaviour* (Vol. 2). New York: Academic Press, pp.159-204.
- Malamuth, N.M., and Ceniti, J. (1986). Repeated exposure to violent and non-violent pornography: Likelihood of raping ratings and laboratory aggression against women. *Aggressive Behaviour, 12*, 129-137.
- Malamuth, N.M., & Donnerstein, E. (Eds.) (1984). *Pornography and Sexual Aggression*. New York: Academic Press.
- Manicas, P.T., and Secord, P.F. (1983). Implications for psychology of the new philosophy of science. *American Psychologist, 38*, 399-413.
- Palys, T. S. (1978). Simulation methods and social psychology. *Journal for the Theory of Social Behavior, 8*, 343-368.
- Palys, T. S. (1984). *A Content Analysis of Sexually Explicit Videos in British Columbia*. Ottawa: Department of Justice. Working Papers on Pornography and Prostitution. (Report No. 15).
- Palys, T. S. (1986). Testing the Common Wisdom: The Social Content of Video Pornography. *Canadian Psychology, 27*, 22-35.
- Palys, T.S. (1989). Addressing the "Third Criterion" in Experimentalist Research: Towards a Balance of Manipulative and Analytic Control. In I. Benbasat (Ed.), *The Information Systems Research Challenge: Experimental Research Methods*. (Volume 2 of the Harvard Business School Research Colloquium Series; series editors J.I. Cash, Jr., and J.F. Nunamaker, Jr.). Boston, MA: Harvard Business School.
- Palys, T.S. (1992). *Research Decisions: Qualitative and Quantitative Perspectives*. Toronto: Harcourt and Brace, Canada.

- Palys, T.S. & Lowman, J. (1984). *Methodological meta-issues in pornography research: Ecological representativeness and contextual integrity*. Paper presented at the annual meetings of the Canadian Psychological Association; Ottawa, Ontario.
- Rosenthal, R., and Rosnow, R.L. (1984). *Essentials of Behavioural Research: Methods and Data Analysis*. NY: McGraw-Hill.
- Zillmann, D., and Bryant, J. (1984). Effects of massive exposure to pornography. In N.M.Malamuth & E.Donnerstein (Eds.) *Pornography and Sexual Aggression*. New York: Academic Press, pp.115-138.