

EXTRACTING DEPENDABLE GENERALIZATIONS FROM THE PERSUASION EFFECTS LITERATURE: SOME ISSUES IN META-ANALYTIC REVIEWS

DANIEL J. O'KEEFE

This essay discusses some issues attendant to meta-analytic reviews of the research literature concerning factors influencing persuasive effectiveness. Narrowly put, my purpose is to sensitize both producers and consumers of such reviews to some important but little-discussed aspects of these reviews. But there is a broader theme to be struck, concerning the importance of tailoring meta-analytic procedures to the characteristics and concerns of the research area under investigation. The research literature on persuasive message effects obviously invites meta-analytic attention, but, as will be seen, the usefulness of the application of meta-analytic procedures to this or any substantive area depends upon how well these procedures have been fitted to the materials and questions of interest.

There is a good deal of diversity in the purposes and procedures of meta-analytic work—both generally and within persuasion research specifically. But (with the usual caveats about inevitable oversimplification) insofar as the persuasion effects literature is concerned, there are commonly two general purposes behind meta-analyses. First, one wants to know the average effect size (and the associated confidence interval) for a given variable's impact on persuasive outcomes. Second, one wants to know whether there is significant variability in the collection of effect sizes, and (if so) whether there is any discernible pattern to this variability.

In the course of pursuing these two broad meta-analytic aims in reviewing persuasion effects research, however, a number of important issues arise. The issues I want to discuss here can be grouped under two general headings: issues connected with the literature search (identification of relevant studies, etc.) and issues connected with the analysis of the collected studies (extraction of effect sizes from studies, analyses of a collection of effect sizes, etc.).

LITERATURE SEARCH ISSUES

Undertaking a meta-analysis requires identifying and collecting the relevant studies for analysis. There are a number of customary ways of locating the relevant research literature. One can scan the reference lists of studies already known to be relevant; one can examine citation indices to identify candidate studies; one can use keyword searches in appropriate publications and databases to locate candidate studies; and so on.

I want to underscore three points concerning this phase of meta-analytic work: the importance of completeness in the literature search, the value of

Daniel J. O'Keefe is associate professor of speech communication, University of Illinois at Urbana-Champaign. His work has received the Speech Communication Association's Charles Woolbert Research Award and Golden Anniversary Monograph Award. He is the author of Persuasion: Theory and research (Newbury Park, CA: Sage, 1990). A version of this paper was presented at the November 1991 convention of the Speech Communication Association, Atlanta, GA.

identifying excluded studies, and the importance of ignoring investigators' labels for experimental manipulations.

Completeness

It's natural to assume that, in retrieving relevant studies, one wants to retrieve *all* the relevant studies; that is, it's natural to assume that some complete (exhaustive) list of relevant investigations is to be sought. However, there have recently been some comments suggesting that completeness isn't necessary or even desirable in meta-analytic reviewing (e.g., Wachter & Straf, 1990, p. xx), and indeed warning against "overambition at the data collection stage" (Laird, 1990, p. 48). But it is important not to misunderstand these remarks. These comments imagine a circumstance in which there are 90 or 100 or more studies at hand. For instance, one meta-analysis of the research literature concerning aphasia treatments covered 114 separate studies (Greenhouse, Fromm, Iyengar, Dew, Holland, & Kass, 1990). In such a situation it may well be true that tracking down just one more investigation won't be all that important.

But the circumstance in the persuasion literature is quite different. The number of relevant studies may be rather small, even for comparatively well-studied questions. Consider, for instance, that in persuasion effects research, fear-appeal effects is a relatively extensively studied topic. Even so, Boster and Mongeau's (1984) meta-analysis reviewed only 25 studies concerning attitudinal effects of fear appeals (see p. 347); and Sutton's (1982) meta-analysis, with somewhat different inclusion criteria, reviewed only 21 studies (see p. 313). For other persuasion variables, there will likely be even fewer studies.

The point, thus, is that whereas some meta-analysts might (by virtue of the large number of studies available) be in a position to be unconcerned by a failure to retrieve that one elusive investigation, meta-analysts of the persuasion effects literature are unlikely to have the same luxury. The smaller the number of extant studies on a topic, the more important it is that one's literature search retrieve all the studies.

Excluded Studies

I want to plump for a particular meta-analytic reporting practice: the identification of excluded studies. Meta-analytic reviews commonly specify the criteria for including studies in the review, and often specify just which studies were included (as they should). But excluded studies are rarely mentioned, and this is unfortunate. After all, if some apparently-relevant studies are not cited in a review's references, one might well have grounds for doubt about the review's conclusions, given that (to all appearances) relevant research findings were not included in the review (were overlooked, or were not located, or whatever).

Convenient examples are provided by the several meta-analyses of the research concerning the effectiveness of the foot-in-the-door and the door-in-the-face compliance techniques (Beaman, Cole, Preston, Klentz, & Steblay, 1983; Dillard, Hunter, & Burgoon, 1984; Fern, Monroe, & Avila, 1986). One of these reviews (Fern et al., 1986) not only failed to specify the studies that were excluded, it failed to provide information about the studies that were included (such information was available separately). A second review (Beaman et al., 1983) did not specify the excluded studies (and although it cited at least some

included studies, one cannot be sure whether all of the included studies were cited). The third review (Dillard et al., 1984), in addition to citing included studies, provided examples of excluded studies: for each specified inclusion criterion, an example was given of a study failing to meet that criterion (p. 466). This last policy is better than not mentioning excluded studies at all, but is not as desirable as providing a fuller accounting.

Consider, for instance: What is one to make of the fact that the study by Arbuthnot, Tedeschi, Wayner, Turner, Kressel, and Rush (1977) was cited by Beaman et al. (1983), but not by Dillard et al. (1984)? Did this investigation fail to meet some inclusion criterion of Dillard et al., or was it not retrieved? Or, to turn the coin over, what is one to make of the fact that Hansen and Robinson's (1980) study was cited by Dillard et al. (1984), but not by Beaman et al. (1983)? Were Beaman et al. unaware of this study, or was it excluded for some principled reason?

Not specifying excluded studies has the effect of needlessly reducing the credibility of the review. When a published meta-analysis does not evince awareness of apparently-relevant studies, readers are surely entitled to reduce their confidence in the review's analysis and conclusions. Moreover, not specifying excluded studies reduces the utility of the review to future investigators. For instance, it makes subsequent literature searching more difficult than it needs to be (because researchers have to discover the excluded studies in some other way, instead of having the studies already specified).

Obviously a meta-analysis can't literally list *all* of the studies that were not reviewed. But the suggestion here is that meta-analysts should make a conscientious effort to specify those marginal studies, those seemingly-relevant studies, that were excluded. There is little or no cost to such an effort. And only by identifying excluded studies explicitly (and giving a rationale for their exclusion) can meta-analysts offer readers assurances that these studies weren't overlooked, but instead were excluded for good reason.

Investigators' Labels

In searching the literature for relevant research, it is important that meta-analysts look past the investigators' original labels for experimental manipulations. The point I want to make here is not the familiar one, that just because investigators labelled their manipulation as "X" doesn't mean that the manipulation should be taken as a manipulation of X. Instead, what I want to emphasize is that just because investigators didn't label their manipulation as "X" doesn't mean that the manipulation isn't X.

Consider, for example, that the research reported by Jarrett and Sherriffs (1953) is rarely mentioned in discussions of the relative effectiveness of one-sided and two-sided persuasive messages. But the message manipulation that Jarrett and Sherriffs used is (in the relevant respects) identical to that used in Rosnow's (1968) investigation and in Bettinghaus and Baseheart's (1969) study. In all these cases the comparison was between a message that contained only arguments favoring one side of the issue (the "one-sided" message) and a message containing roughly equal arguments for each side of the issue (the "two-sided" message). Why are Rosnow's and Bettinghaus and Baseheart's studies more familiar to students of message sidedness? One plausible account is

simply that Jarrett and Sherriffs didn't *label* their research an investigation of "one-sided versus two-sided messages" (but instead as a comparison of "propaganda" and "impartial presentation"). Nevertheless, the experimental manipulation employed by Jarrett and Sherriffs is appropriately included with those of Rosnow and Bettinghaus and Baseheart in any review of the sidedness literature.

This is not merely a problem of old studies failing to use now-current terminology; for instance, one can locate more recent investigations for which the vocabulary of message sidedness was available and seemingly appropriate (that is, the manipulations might well have been termed sidedness manipulations) but in which that vocabulary was not employed (e.g., Srull, 1983). Instead, the problem is more general, and stems from the possibility of describing an experimental manipulation in any number of different, equally legitimate ways (consider, as another example, "fear appeal," "emotional appeal," "depiction of undesirable consequences of noncompliance," and so on).

I want particularly to emphasize the implication of this point for literature searching. The fact that studies relevant to variable X may not describe their manipulation as "X" makes retrieval of relevant studies difficult. Notably, the customary ways of locating relevant studies (references in articles already known to be pertinent, keyword searches, examination of abstracts, citation checks, etc.) are insensitive to the possibility of redescribing some manipulation. So far as this problem is concerned, there is no substitute for direct exposure to, and knowledge of, the primary research literature. (This, I think, should give some pause to those who appear eager to embrace the role of roving meta-analyst, ready to meta-analyze any research area whatever—including areas in which they have but meager previous contact with the primary literature.)

ANALYSIS ISSUES

Once a collection of relevant studies is in hand, the meta-analyst faces the tasks of obtaining an effect size from each study and analyzing the collection of effect sizes (e.g., combining these into an average). Here I want to raise some issues concerning how effect sizes are extracted from studies, how the unit of analysis is to be identified, and how the collection of effect sizes is to be analyzed.

Effect Sizes in Factorial Designs

Much persuasion effect research takes the form of an experimental design with multiple independent variables, and the most common form of statistical analysis is analysis of variance. In such cases, the *F*-value for the variable of interest typically provides the basis for the computation of the effect size.

However, the size of the *F* for the factor of interest varies depending upon what other factors are included in the analysis. (Including additional factors in the analysis—if such factors have any variance at all associated with them—increases the *F*-ratio for the factor of interest.) Where sufficient information is provided, the meta-analyst will have the option of computing the *F* in alternative ways (corresponding to alternative statistical analyses of the original data). The issue is how the meta-analyst ought to proceed in such a circumstance.

One might, of course, simply use whatever *F*-value was reported in the original investigation. But to unthinkingly use reported *F*-values is to put one's

meta-analysis at the mercy of whatever statistical analyses the original investigators happened to employ. It thus invites unnoticed inconsistency in the computation of effect sizes (and so muddies the comparison of effect sizes). If study A's ANOVA includes sex as a factor, but study B's does not, then simply using the reported F -values can produce misleading results (e.g., the two studies' effect sizes may appear to be identical when they should properly be seen as different—or vice versa).

The question thus becomes whether there is some principled way to proceed in circumstances such as these. One possibility is simply to recast every analysis as a one-way ANOVA. That is, one would put all the other factors (that had been included in the original analysis) into the error term—regardless of their nature—and recompute the F for the factor of interest. (Such a policy is recommended by, e.g., Cooper, 1989, pp. 106–107.)

A different possible principle has been articulated by Johnson (1989, p. 16). His (tentative) recommendation is that any individual-difference variables (such as sex) should be included in the error term, but factors representing experimental manipulations (other than the factor of interest) should not. Crudely put, the idea is that any variance the experimenter put in (that is, created through some manipulation) ought to be removed before computing the effect size of interest. Since different studies might manipulate different additional variables, and since the effects of these manipulations may vary in impact, the best way to ensure a reasonably common basis of comparison across studies is to use an error term that does not include the sums of squares attributable to such variables. By contrast, where individual-difference variables are concerned, the best way to produce comparable effect sizes will be to include such variables in the error term (because in designs that don't have such variables as separate factors, those variables' effects will typically still be present, contained in the error term).

In practice, of course, one will not always be able to perform such recomputations. And there need not be any single procedure adopted uniformly either across or within meta-analyses; for instance, one might report alternative results based on different means of extracting the effect sizes. But the general points are surely plain: In no case should reported F -values be used unthinkingly; investigators should be sensitive to the possibilities for alternative ways of computing effect sizes in these research designs; and investigators should be explicit about the procedures followed.

Identifying Units of Analysis

Meta-analysis involves the analysis of a collection of effect sizes, and in that sense the smallest unit of analysis in most meta-analyses (the "case") is the effect size (just as in primary research the smallest unit of analysis—the case—is typically taken to be the individual respondent). Now in most applications of meta-analysis, each individual study contributes one and only one effect size, and hence the unit of analysis becomes (in effect) the individual study. (I bypass here the issues connected to circumstances in which each study contributes more than one effect size.) That is, one has a collection of studies whose characteristics (effect size, sample size, etc.) are recorded for analysis (parallel to primary research, where one records information for each respondent), and

these different studies are treated as independent cases (just as respondents are in primary research).

This natural way of proceeding—treating the “study” as the unit of analysis in meta-analysis—encounters various complications. Here I want to underscore a particular difficulty encountered in the persuasion effects literature, namely, the repeated use of the same messages across different studies (for some previous discussion of this matter in another context, see Jackson, O’Keefe, Jacobs, & Brashers, 1989, pp. 370–371).

As an example, one might consider that in research on the persuasive effects of rhetorical questions, several of the studies have used (virtually) the same messages (Burnkrant & Howard, 1984; Petty, Cacioppo, & Heesacker, 1981; Swasy & Munch, 1985). An appropriate meta-analytic treatment of this research area will surely want to be sensitive to this fact, and not simply assume that each different report studies a different message. Where one’s interest is in generalizing across messages, it will be important to notice the difference between three studies of three different messages and three studies of the same message—and one’s choice of a unit of analysis can reflect that. (My general recommendation would be to combine all the effect sizes that are based on a given message.)

The general point to be emphasized is that the particular characteristics of the persuasion effects literature suggest that certain customary ways of proceeding in meta-analysis—treating each study as providing an independent effect size estimate—may not be entirely appropriate, and certainly should not be unthinkingly adopted. (For a more general discussion of units of analysis in meta-analysis and related issues of independence of studies, see Cooper, 1989, pp. 76–79.)

Analyzing Effect Sizes

I want to raise some worries about common ways of analyzing collections of effect sizes. As a means of priming the relevant intuitions, consider the following situation. A researcher has conducted an investigation of the persuasive effects of a particular message variable using a single-message design (i.e., a design in which the experimental manipulation involves versions of just one stimulus message); the study had 100 respondents. Another 100 respondents become available to the investigator. The question is, how should the investigator spend those additional respondents, assuming an interest in generalizing effects across messages? Specifically, should those new respondents be exposed to the same message as the previous participants, or should a different message (or messages) be used? I think that (nowadays, anyway) the widespread intuition among investigators would be that the researcher should use new message materials; the evidence for any claimed effect will be stronger if the effect can be found to obtain across message replications.

Examples of this sort have been used previously (in Jackson et al., 1989, p. 377) for the purpose of underwriting the policy of preferring multiple-message research evidence over single-message research evidence. But I want to use this sort of example for another purpose, by considering just *why* it is that one will have the intuition that multiple messages are to be preferred.

Briefly put, I think the basis for the preference for multiple messages is simply the recognition of the possibility that the effect of a manipulation may vary from

message to message—that the effect of a manipulation may not be constant from one message to the next. But once one recognizes this possibility, then arguably one should begin to think differently about how to treat observed effect sizes.

Consider, initially, the (apparently simple) question of how to average effect sizes. The common meta-analytic practice is to weight by sample size (that is, human sample size, the number of respondents on which each effect size is based) or some analogue thereto (e.g., *df*). And this practice makes a good deal of sense, since it gives greater weight to those effect-size estimates that are presumably more accurate (more accurate, because based on larger samples).

However, if one recognizes the possibility that the effect of the variable in question may vary from message to message, then this procedure is not unproblematic. An increase in the number of respondents for a given message presumably gives a better estimate of the size of the treatment effect for *that* particular message, but this doesn't automatically mean a better estimate of the size of the treatment effect across messages.

Think about a simple example, involving two studies, each with a single-message design. Study A has 50 respondents, study B has 250. The message used in study B is not intrinsically more important or interesting, or any better an example of messages of interest, than is the message in study A; insofar as generalization across messages is concerned, the two studies are equally valuable (in the sense that there's no reason to be more interested in the results for one study's message than for the other's). But the common meta-analytic procedures would give study B's results 5 times more weight than study A's as estimates of the average effect—which seems in some ways to be a misleading way of characterizing the findings of interest.

One crude way of addressing this awkwardness is to provide two mean effect sizes, one based on weighting effect sizes by their sample sizes, the other based on weighting effect sizes by the number of messages on which they are based (see, e.g., O'Keefe, 1987). This policy is not without its difficulties, but at least it offers readers a clearer view of the phenomenon under discussion.

Thus far I've been discussing some complexities attendant to computing the mean effect size across a collection of effect sizes, but similar issues arise elsewhere within meta-analysis. Consider, for instance, the task of constructing a confidence interval around a mean effect size. The usual meta-analytic procedures for this task will yield confidence intervals that vary substantially as a function of the total number of human respondents—but the number of effect sizes (studies, messages, implementations) makes little difference. That is, the usual procedures for constructing these confidence intervals are largely indifferent to the number of estimates involved (see, e.g., Cooper, 1989, p. 110; or Hedges & Olkin, 1985, p. 231). So, for instance, these procedures will yield confidence intervals that are (*ceteris paribus*) very nearly the same if one has 4 estimates with a sample size of 500 in each or 20 estimates with a sample size of 100 in each.¹

These customary procedures (for computing a mean effect size or placing a confidence interval around a mean effect size) will seem to make good sense *if* one has a particular vision of the underlying phenomenon. Specifically, the common sample-size-weighting procedure appears to envision a circumstance in which each observed effect size reflects simply the combined effects of the

single underlying true effect of the treatment variable and sampling error (plus, of course, the operation of any moderator variables).

But this is not the only possible image of the phenomenon. Imagine, instead, that one thinks in the following way: "There is no single true or population effect of the 'treatment' across [messages]. Rather, there is a *distribution* of true effects; each [message] has its own unique true effect." Thinking this way, one may compute (across messages) "an average true effect of the treatment as an index of overall efficacy"—but one should recognize that "it is quite possible for the average true effect to be greater than zero, whereas the true effect of the treatment is negative in nearly half the [messages]" (Hedges & Olkin, 1985, p. 190). Notice that the common preference for multiple-message data over single-message data seems very much in harmony with this line of thinking, with this way of conceptualizing the phenomenon under investigation: Because we recognize that there is a distribution of unique true effects, we realize that examining just one message may give a misleading picture of the treatment's general effect (average true effect).

Readers familiar with meta-analytic techniques will have discerned Hedges and Olkin's (1985) random effects model in the preceding paragraph. That is, there are meta-analytic procedures available for analyzing effect sizes conceptualized in that way. And one might well suppose that such procedures are better-suited than more familiar meta-analytic methods are to the analysis of the persuasion effects literature.

But my point here is not to argue for some particular analytic procedure as intrinsically preferable to all others (not even within the persuasion effects literature, much less communication research generally). Instead it is to underscore the diversity of available meta-analytic procedures. One need not follow customary or familiar meta-analytic paths—and one certainly ought not follow such paths blindly. There is not one unique correct way to perform a meta-analysis. Indeed, Hedges and Olkin (1986, p. 21) have pointed out that one danger of Hunter, Schmidt, and Jackson's (1982) manner of presentation is precisely that unwary readers can get the impression that there is one and only one way that meta-analysis is done. Producers and consumers of meta-analyses should recognize the diversity of analytic procedures available, should be sensitive to the consequences and potential weaknesses of the procedures followed, and should be attentive to the fit between the phenomena under study and the procedures chosen.

CONCLUSION

When meta-analytic procedures were first becoming prominent, one sometimes saw the fear expressed that naive consumers of meta-analysis would be deceived by the apparent "precision" and "objectivity" of this quantitative procedure—that they would be led to be uncritical of its application, led to have unwarranted confidence in the derived conclusions, and so on (e.g., Cook & Leviton, 1980, p. 455; Strube & Hartmann, 1982, p. 132). Of course, this is not, and was never intended as, a good argument against the use of meta-analysis; instead, it's an argument for creating informed consumers. And the primary aim of this essay has been the encouragement of better-informed consumers—

and producers—of meta-analytic work, especially in the area of persuasion effects research.

But there is also a broader theme here concerning the interpenetration of method and belief. We have seen, in a variety of ways, how the particular characteristics and concerns of persuasion effects research ask for correspondingly particular meta-analytic treatment. Certain familiar, customary ways of proceeding in meta-analysis (treating studies as the units of analysis, determining mean effect size through sample-size-weighting, and so on) ought not be mindlessly applied to this research literature. But of course this is not something distinctive about persuasion research. The application of meta-analytic procedures to any particular area of study will require close attention to the peculiarities that arise there. Facile assumptions about “the” correct way to do meta-analysis, unreflective adoption of customary meta-analytic procedures, lack of familiarity with and sensitivity to the materials and questions of the substantive research area—these are surely the makings of meta-analyses that fall rather short of what is wanted.

ENDNOTE

¹In fact, for a given weighted-average correlation value (termed z by Cooper, 1989, p. 108, and termed z_r by Hedges & Olkin, 1985, p. 231), the customary confidence interval around that figure (as constructed using, e.g., the procedures described on p. 231 of Hedges & Olkin, 1985, or equation 15 in Cooper, 1989, p. 110) will actually be just a touch smaller in the case of 4 estimates with $n = 500$ in each than it will in the case of 20 estimates with $n = 100$. In general, as the number of estimates increases (for a given total N and a given weighted-average correlation value), the width of the confidence interval also increases ever so slightly.

REFERENCES

- Arbuthnot, J., Tedeschi, R., Wayner, M., Turner, J., Kressel, S., & Rush, R. (1977). The induction of sustained recycling behavior through the foot-in-the-door technique. *Journal of Environmental Systems*, 6, 355–368.
- Beaman, A.L., Cole, C.M., Preston, M., Klentz, B., & Steblay, N.M. (1983). Fifteen years of foot-in-the-door research: A meta-analysis. *Personality and Social Psychology Bulletin*, 9, 181–196.
- Bettinghaus, E.P., & Baseheart, J.R. (1969). Some specific factors affecting attitude change. *Journal of Communication*, 19, 227–238.
- Boster, F.J., & Mongeau, P. (1984). Fear-arousing persuasive messages. In R.N. Bostrom (Ed.), *Communication yearbook 8* (pp. 330–375). Beverly Hills, CA: Sage.
- Burnkrant, R.E., & Howard, D.J. (1984). Effects of the use of introductory rhetorical questions versus statements on information processing. *Journal of Personality and Social Psychology*, 47, 1218–1230.
- Cook, T.D., & Leviton, L.C. (1980). Reviewing the literature: A comparison of traditional methods with meta-analysis. *Journal of Personality*, 48, 449–472.
- Cooper, H.M. (1989). *Integrating research: A guide for literature reviews* (2nd ed.). Newbury Park, CA: Sage.
- Dillard, J.P., Hunter, J.E., & Burgoon, M. (1984). Sequential-request persuasive strategies: Meta-analysis of foot-in-the-door and door-in-the-face. *Human Communication Research*, 10, 461–488.
- Fern, E.F., Monroe, K.B., & Avila, R.A. (1986). Effectiveness of multiple request strategies: A synthesis of research results. *Journal of Marketing Research*, 23, 144–152.
- Greenhouse, J.B., Fromm, D., Iyengar, S., Dew, M.A., Holland, A.L., & Kass, R.E. (1990). The making of a meta-analysis: A quantitative review of the aphasia treatment literature. In K.W. Wachter & M.L. Straff (Eds.), *The future of meta-analysis* (pp. 29–46). New York: Russell Sage Foundation.
- Hansen, R.A., & Robinson, L.M. (1980). Testing the effectiveness of alternative foot-in-the-door manipulations. *Journal of Marketing Research*, 17, 359–364.
- Hedges, L.V., & Olkin, I. (1985). *Statistical methods for meta-analysis*. New York: Academic Press.
- Hedges, L.V., & Olkin, I. (1986). Meta-analysis: A review and a new view. *Educational Researcher*, 15, 14–21.
- Hunter, J.E., Schmidt, F.L., & Jackson, G.B. (1982). *Meta-analysis: Cumulating research findings across studies*. Beverly Hills, CA: Sage.

- Jackson, S., O'Keefe, D.J., Jacobs, S., & Brashers, D.E. (1989). Messages as replications: Toward a message-centered design strategy. *Communication Monographs*, *56*, 364-384.
- Jarrett, R.F., & Sherriffs, A.C. (1953). Propaganda, debate, and impartial presentation as determiners of attitude change. *Journal of Abnormal and Social Psychology*, *48*, 33-41.
- Johnson, B.T. (1989). *DSTAT: Software for the meta-analytic review of research literatures*. Hillsdale, NJ: Erlbaum.
- Laird, N.M. (1990). A discussion of the aphasia study. In K.W. Wachter & M.L. Straf (Eds.), *The future of meta-analysis* (pp. 47-52). New York: Russell Sage Foundation.
- O'Keefe, D.J. (1987). The persuasive effects of delaying identification of high- and low-credibility communicators: A meta-analytic review. *Central States Speech Journal*, *38*, 63-72.
- Petty, R.E., Cacioppo, J.T., & Heesacker, M. (1981). Effects of rhetorical questions on persuasion: A cognitive response analysis. *Journal of Personality and Social Psychology*, *40*, 432-440.
- Rosnow, R.L. (1968). One-sided versus two-sided communication under indirect awareness of persuasive intent. *Public Opinion Quarterly*, *32*, 95-101.
- Srull, T.K. (1983). Affect and memory: The impact of affective reactions in advertising on the representation of product information in memory. In R.P. Bagozzi & A.M. Tybout (Eds.), *Advances in consumer research* (Vol. 10, pp. 520-525). Ann Arbor, MI: Association for Consumer Research.
- Strube, M.J., & Hartmann, D.P. (1982). A critical appraisal of meta-analysis. *British Journal of Clinical Psychology*, *21*, 129-139.
- Sutton, S.R. (1982). Fear-arousing communications: A critical examination of theory and research. In J.R. Eiser (Ed.), *Social psychology and behavioral medicine* (pp. 303-337). New York: Wiley.
- Swasy, J.L., & Munch, J.T. (1985). Examining the target of receiver elaborations: Rhetorical question effects on source processing and persuasion. *Journal of Consumer Research*, *11*, 877-886.
- Wachter, K.W., & Straf, M.L. (1990). Introduction. In K.W. Wachter & M.L. Straf (Eds.), *The future of meta-analysis* (pp. xiii-xxviii). New York: Russell Sage Foundation.

