JOHN U. FARLEY, DONALD R. LEHMANN, and LANE H. MANN*

Generalized knowledge comes from cumulating results across studies, a process known as meta-analysis. Efficiently increasing generalized knowledge in a defined area—estimates of price or advertising, for example—is one important goal for research. Because (1) most meta-analyses are based on highly inefficient and unbalanced natural experiments or designs and (2) additional studies are costly, carefully selecting the next study is important. The authors demonstrate that, rather than simply selecting a study that uses currently underrepresented design variables, a procedure that reduces collinearity among design variables will produce far superior improvements in knowledge.

Designing the Next Study for Maximum Impact

There are two fundamentally different ways to enhance knowledge. One approach relies heavily on creativity: generating new paradigms, models, methods, variables, and applications. This approach has the potential to produce major breakthroughs, albeit along with the nontrivial possibility of total failure. The other approach is less grand in scope and focuses more on refining knowledge within an existing paradigm. This approach involves replication and systematic changing of variables in an attempt to increase certainty in the estimates of the impact of (predetermined) variables. Although work of the first type is desirable, this article focuses on the more modest goal of providing guidance for designing a study that increases the precision of knowledge in an established field of inquiry. Specifically, this article suggests an approach for designing the next study in a defined field that is likely, a priori, to be the most useful. An implicit assumption behind our approach is that data are dear and/or expensive. The factorial design mentality, which is relevant when dealing with individual subjects in a limited within-study design (e.g., 2 × 2 × 2), simply does not apply when each data point comes from a complex design and represents a study that costs $100,000 or more and requires years to complete. Therefore, we concentrate on designing the next study in a field to optimize improvement in knowledge.

*John U. Farley is C.U. Starr Senior Research Fellow, Amos Tuck School, Dartmouth University, and China Europe International Business School (e-mail: jfarley@Dartmouth.edu). Donald R. Lehmann is George E. Warren Professor, Graduate School of Business, Columbia University (e-mail: DLehmann@Research.GSB.Columbia.edu). Lane H. Mann is Director of Market Research, Fidelity Management Trust Company (e-mail: Lane.Mann@FMR.com). The authors thank Hooman Estelami for his assistance with the analysis.

THE "NATURAL" EXPERIMENTAL DESIGN IN META-ANALYSIS

Meta-analysis in marketing usually involves summarizing parameters drawn from a fairly large number of studies, which means focusing on relatively well-studied areas that usually develop with limited attention to how later studies contribute to existing knowledge. The basic approach involves three stages:

Step 1. A model is selected as the basis for study, such as $Y = B_0 + XB + e_1$. For example, X might be advertising and Y sales, so B represents the impact of advertising on sales. A meta-analyst gathers n estimates of B from different models that may contain other explanatory variables and involve different products or estimation technologies.

Step 2. Substantive (type of product or service, study setting/situation) and methodological (measurement method, estimation method) characteristics that might be related to systematic variation of B are identified and recorded in a design matrix, Z. This matrix may include both main effects and appropriate interactions. Each row of the matrix represents the m characteristics of the parameter value in question.

Step 3. The impact of situational characteristics on the parameter are estimated, so that $B = C_0 + ZC + e_2$. The estimates of C represent information about the impact of the design variables on B (i.e., $C_i = 0$ implies the parameter generalizes—is equal—across condition i; the variance of C_j indicates how certain we are of a systematic, generalizable difference across conditions).

In general, the design matrix Z is configured as a series of dummy (1 – 0) variables that represents the presence or absence of the various situational characteristics. This matrix, dimensioned by n results (estimated B values) and m factors, represents the "natural experimental design" that characterizes the body of knowledge. Because the contributing studies generally fall into no orderly pattern, most meta-analysis designs are nonorthogonal, and it may be impossible to estimate some main effects, much less a complete set of complex interactions. Rather than regarding such deviations from orthogonality as nuisances in estimating C, we use the content of Z to decide how best to add rows to improve estimates of C.

Empty Cells and Sparse Data

Two problems in the configuration of Z are of obvious importance in estimating C. By their nature, most meta-analyses have a large number of empty cells. The first instinct is to fill these or simplify/scale back the meta-design to reduce the number of cells. However, scaling back the design limits the eventual scope of the generalizations, and as we will see, simply filling empty cells may be nonoptimal. We also could try to pick studies to fill cells with large within-cell variance. However, because of the sparse nature of the natural experimental design (and the large number of empty cells), reliable within-cell variances are not available.

Small column sums of Z indicate sparse data, levels of factors that will be either difficult or impossible to assess. (Because coding of dummy variables is arbitrary, a large sum also indicates sparse data for the 0-coded level of a variable.) In the extreme, a column sum of 1 indicates that the particular level of a factor occurs in a specific study, in which it is perfectly confounded with other idiosyncratic characteristics of that study.

Collinearity

Another problem with a typical meta-analysis design matrix Z is collinearity. Some collinearity is due to simple variable redundancy (i.e., U.S. studies on retail promotions all use scanner data). However, more complex and subtle patterns often are encountered in practice that affect the invertibility of (Z′Z). For example, two near-identical columns might be, for all practical purposes, confounded or collinear.

Collinearity among characteristics generally has been treated as a nuisance, and steps have been developed to eliminate it (Farley and Lehmann 1986). However, the reliability of the estimates of C is, in large part, determined by Z, so to improve knowledge we must add observations to Z that reduce uncertainty about C. The key to this is reducing collinearity in Z (Belsley 1980).

SELECTING THE NEXT STUDY

Because meta-analyses reflect the generally unplanned nature of research in a particular field, one option is to encourage researchers to conduct the next study in a "best" way by providing some guidance a priori. Encouraging such behavior is not necessarily easy or consistent with what some perceive as academic freedom. However, organizations such as the Marketing Science Institute have found that publishing priorities for research in a particular area has provided useful guidance without stifling initiative. Editorials in journals and calls for papers also have some impact.

If our goal is to increase knowledge in a defined field, selecting the next study should follow a logical progression, as follows:

Ensure there is variance on design variables. To ensure that the variance in the estimate of the impact of a variable is finite, the design variables must be represented in one or, preferably, more studies.

Deal with perfect collinearity. Select the n + 1 study so that perfect empirical collinearity/confounding of design variables is eliminated. A standard regression program (such as SAS) identifies the confounded variables in equation form. The next observation should be one that fails to satisfy as many of these empirical redundancies as possible.

Minimize variance in estimates of the design dummy variable effects. When perfect redundancy is eliminated, one obvious approach is to fill underrepresented cells (e.g., by picking a design profile that maximizes the number of 1s in columns of the matrix for which the totals of the first n studies are relatively small). This method makes intuitive sense and is easy to implement, and we have recommended it in the past. Unfortunately, it is not optimal. We recommend instead picking the next observation such that the sum of the variances of the parameters of the meta-analysis model (Var C) is reduced. We operationalize this by minimizing collinearity.[1]

A trade-off exists between reducing the variance of the coefficient (C_j) with the largest variance (which a criterion such as minimizing the product of the variances would suggest) and improving several variances at the "expense" of others (e.g., improve knowledge of the impact of product category and leave the uncertainty about the effect of measurement method unresolved). In the absence of theoretical or practical considerations, we suggest minimizing the sum of the variances of the effects of the design variables, which is equivalent to minimizing the trace of the $(Z'Z)^{-1}$ matrix. Such minimization is not necessarily accomplished by simply increasing the observations in underrepresented conditions. In essence, our method identifies which observation would most reduce collinearity in the meta-analysis design matrix Z.[2]

TWO SIMPLE EXAMPLES

In both examples in Table 1, four binary (1,0) design variables are used. In Example A, for the 16 observations, each variable is at each level 8 times, which makes the low fre-

[1]Variance minimization and collinearity reduction may not be identical if the design variables are negatively correlated (Mela and Kopalle 1998).

[2]An alternative criterion is to minimize the condition index, the ratio of the largest to the smallest eigenvalue. Because the smallest eigenvector is orthogonal to the largest, picking an observation to match it simultaneously reduces the maximum and increases the minimum eigenvalue. A different approach/objective is to focus on (minimize) the variance of the predicted value of the dependent variable in the meta-analysis. This variance is equal to $S_Y^2 \, |R|/|R_{ZZ}|$, where R is the correlation matrix of B plus all the Z variables. Note that, because R_{ZZ} is a subset of R, it is not clear what rule follows from this objective. However, if the within-cell variances are equal, then there is no advantage in picking any particular design for the next study. Conversely, if the true cell variances are unequal (heteroscedastic) because of, for example, unequal measurement error, then minimizing the calculated standard error of prediction $S_{B.Z}$ is accomplished by adding more observations to the smallest variance cell. This will increase collinearity among the Zs and, thus, obscure the impact of the design variables. Thus, though (1) we want to reduce observed variance within cells if possible, and (2) for some purposes (e.g., policy decisions), minimizing error of prediction from the meta-analyses is an appropriate goal, we concentrate on minimizing variances in the effects of the design variables to increase confidence in the contingent generalizations that emerge from meta-analysis. Alternatively, one could try to minimize the variance in the prediction for a particular combination (c) of design variables. In this case, Silvey (1969) shows that the optimal next obeservation is $(\mathrm{It}[x'x]^{-1})c$. We thank Wilfried Vanhonacker for this reference.

Table 1

TWO EXAMPLES OF SIMPLE DESIGN MATRICES

EXAMPLE A				
	VARIABLE			
Observation	1	2	3	4
1	1	1	1	1
2	1	1	0	0
3	1	1	1	1
4	1	1	1	1
5	1	1	1	1
6	1	1	0	1
7	1	0	1	1
8	1	0	1	0
9	0	1	1	0
10	0	0	0	0
11	0	0	0	0
12	0	0	0	0
13	0	1	1	1
14	0	0	0	0
15	0	0	0	1
16	0	0	0	0
Column Sum	8	8	8	8
EXAMPLE B				
1	1	1	1	1
2	1	1	1	1
3	1	1	1	1
4	1	1	1	0
5	1	1	1	0
6	1	1	1	0
7	1	1	1	0
8	1	1	1	0
9	1	1	1	0
10	1	1	0	0
11	1	1	0	0
12	1	0	1	0
13	1	0	1	0
14	1	1	0	0
15	0	0	0	1
16	0	0	0	0
Column Sum	14	12	11	4

Table 2

CORRELATION MATRICES OF ORIGINAL 16 OBSERVATIONS

EXAMPLE A			
	2	3	4
1	.5	.5	.5
2		.5	.5
3			.5
EXAMPLE B			
	2	3	4
1	.65	.56	–.22
2		.23	.00
3			.08

Table 3

IMPACT OF ADDITIONAL OBSERVATIONS: EXAMPLE B

	Determinant	Sum of Variances	Variance by Variable 1	2	3	4
Original 16	162	244	98	65	48	33
Possible Configuration for an Additional Observation						
1111	204	240	98	65	48	29
1110 (1)	186	238	96	63	46	32
1101 (3)	270	213	86	64	36	27
1100	216	228	90	65	40	33
1011	297	197	78	44	47	27
1010	243	211	81	48	48	33
1001 (4)	411	153	51	42	34	26
1000	321	162	50	44	36	33
0111	375	167	53	45	37	33
0110	393	147	46	39	32	31
0101	265	209	82	48	48	31
0100	267	191	73	39	47	33
0011 (2)	258	224	90	65	40	30
0010	240	211	83	63	33	33
0001	216	233	97	64	47	25
0000	162	244	98	65	48	33

Note: (i) corresponds to the i[th] eigenvector.

quencies for some levels not an issue. Furthermore, all variables overlap and are correlated equally, as we show in Table 2, in that each variable has 1s in 6 of the 8 cases in which each other variable is 1.

Example B is more interesting and typical. First, the four design variables occur with different frequency. Second, the pattern of confounding is different: Variables 1 and 2 and 1 and 3 are quite correlated, and variable 4 is fairly independent of the others (see Table 2). Notice also that, in an information sense, variable 1 is actually the least informative; whereas there are 14 cases in which it is a 1, there are only 2 in which it is a 0. By that same logic, variables 2 and 4 are equally well accounted for, with 12 cases of one value and 4 of the other.

We compute the value of alternative additional observations for Examples A and B. We assess improvement in the variance of C, which is proportional to the sum of the diagonal elements in the inverse of $Z'Z$ and, thus, related to the increase in the determinant of $Z'Z$.

In Example A, as expected, none of the possible 16 configurations of observations results in a dramatic improvement. By contrast, in Example B, the choice of the next observation makes a significant difference (see Table 3). Several interesting results appear. First, whereas an observation that contains the variable with the smallest frequency (0001) has a small benefit (the sum of variance decreases from 244 to 233), that which contains only the most common variable (1000) makes a large improvement (162 versus 244). Second, an exact replicate of a prior study (1010) is more helpful than a previously unexplored combination (0011). Here, the optimal next observation is either (1001) or its complement (0110), which reduces the variance in C by almost half. They do so not by filling sparse cells but by reducing the collinearity between variable 1 and variables 2 and 3.

A METHOD FOR SELECTING THE BEST NEXT STUDY

A combinatorial search to find the optimal next observation becomes cumbersome for the large number of design variables that usually characterize meta-analyses. Using a foldover design may reduce the set to be considered to the same size as the number of observations in the original meta-analysis.[3] The set, however, is not simply the complement of the existing studies, because some combinations of factors are mutually exclusive. For example, a single observation cannot be for both a consumer durable and industrial product simultaneously.

Our recommendation is simple: Pick the next observation to match the smallest eigenvector of (Z′Z). Because the first eigenvector represents the most common observation (i.e., frequently purchased U.S. study), in Example B, (1110) is essentially another study/replication that is similar to previous ones. Because we already have information about it, we need something to break the confounding among the variables it represents. The solution is to choose an observation orthogonal to it.

The math behind picking an observation to match the smallest eigenvalue is straightforward. The variance of the estimates of the effects of the design variables is

$$\text{Var C} = S_{BZ}\,(Z'Z)^{-1}, \quad (1)$$

where S_{BZ} is the standard error of estimate from the meta-analysis regression. One way to compute $(Z'Z)^{-1}$ is as a cofactor matrix multiplied by the inverse of the determinant (Z′Z), so that a large determinant leads to a small variance of C. Because the determinant equals the product of the eigenvalues and because eigenvalues that occur later in principal components analysis typically are small fractions, the product is influenced heavily by the smallest eigenvalue. By making the smallest eigenvalue proportionally bigger (by choosing an observation to match it), the product is likely to be increased the most (for a proof, see Silvey 1969).

Converting eigenvectors (which are, in practice, usually vectors of fractions of all variables) into discrete observations also requires some effort. We adopt the simple rule that any variable that has a positive weight is coded 1, and any with a negative value is coded 0. If the result is an infeasible solution—for example, it suggests that the next study be of both a durable and a nondurable product—we recommend choosing the variable from the group with the largest weight on that component. (In the unlikely case that all weights for a factor group are negative, the one with the least negative weight is chosen.)

EVALUATING AN EXISTING META-ANALYSIS DESIGN: ECONOMETRIC ADVERTISING MODELS

As an illustrative example, we use a meta-analysis of parameters of the effect of advertising on sales from 90 models in 50 books and articles, drawn from Assmus, Farley, and Lehmann's (1984) work, as our starting point. These models focus on sales, not share, and therefore differ from the logit models that recently have become more popular. The design (described in detail by Assmus, Farley, and Lehmann 1984) involves 43 design variables that describe the research environment, model specification, measurement, and estimation.

There is a heavy representation of mature or maturing markets (80%, on which advertising can be expected to have little impact) in the United States (72%, in which advertising to sales ratios tend to be higher; Leff and Farley 1980). By contrast, almost no studies exist in the nonindustrial world for products in the early growth (3%) or decline phases of the life cycle or for weekly (1%) or retail (3%) advertising data. Elimination of definitional singularities in the design variables—for example, all models are fit at either the brand or product level, and all models are either additive or multiplicative—involves dropping some levels.

However, even with this step completed, the reduced design matrix is empirically singular. SAS diagnostics indicate that the classic studies of Lydia Pinkham (research environment) and pooled data (measurement) contribute to singularity, so they also must be eliminated, which produces a design matrix containing 35 study features that are not absolutely singular. Only 23 of the 35 eigenvalues in the reduced, nonsingular design matrix pass the conventional test of a ratio of 50 or less between smallest and largest, so collinearity is still an empirical problem in estimation.

The Best Next Study

To assess our method, we constructed a 91st observation to match each of the 35 eigenvalues of Z′Z. We evaluated each in terms of their impact on the determinant of Z′Z and the sum of the variances of the coefficients of the 35 design variables (measured as the trace of $Z'Z^{-1}$). The results are both encouraging and mildly surprising. The encouraging result is that adding the observation on the basis of the last (35th) eigenvector produced the biggest improvement, with a trace of 135.

The surprising nonmonotonic relation between the eigenvector and the improvement in efficiency (variance) associated with an observation related to it (possibly due to instability in the inversion of a nearly perfectly collinear matrix) raised the question of whether the smallest eigenvalue approach was effective. We therefore generated 500 random new observations and assessed the variances, assuming each was added separately to Z. (Complete enumeration was computationally burdensome, given 35 variables.) The range of variances was between 134 and 358, which means the smallest eigenvalue-based observation did as well as the best of 500 random next observations. Put differently, one well-chosen observation reduces the variance in the estimates of the effects of the design variables by almost two-thirds, whereas a poorly chosen one has little impact. Furthermore, the smallest eigenvalue produced variances of 135/228, or 59% the size of a random observation. Thus, not only is the smallest eigenvalue a good clue as to what the next observation should be, but the choice also makes a dramatic difference in knowledge improvement.

To highlight this, we constructed characteristics of the best next study on the basis of both the best of the 500 random observations and the smallest eigenvalue and then compared it with the typical study of the original 90 (see Table 4). Both the best of 500 random and the smallest eigenvalue method recommend more extensive models (price and/or product included), more sophisticated methods (multiequation, multistep estimation), and work on durables

[3]Rajeev Kohli and James Wiley both independently made this observation, and we thank them for it.

Table 4
BEST NEXT ADVERTISING STUDY

Design of Typical Study in Assmus, Farley, and Lehmann (1984)	*Design of Best of 500 Random Next Studies*	*Design of Optimal Next Study (Study 91)*
Market share	Market share	Total sales
Mature	Growth	Introduction
Food	Cars, other durables	Cars, other durables
United States	United States	United States
Single equation	Multiequation	Multiequation
Carryover included	Carryover not included	Carryover included
Price not included	Price included	Price included
Product not included	Product included	Product included
Other exogenous variables not included	Other exogenous variables	Other exogenous variables
Brand level	Product level	Product level
Ordinary least squares	Multistep estimation	Multistep estimation
Time series	Time series	Time series
Bimonth/quarter	Bimonth/quarter	Bimonth/quarter
Multiplicative	Additive	Additive
Aggregate advertising	Aggregate advertising	Journal advertising

early in the life cycle while maintaining the use of bimonthly time series data and, somewhat surprisingly, U.S. data. As Table 4 shows, best does not mean opposite; approximately half the design variables are the same as the typical study and half differ, exactly as the concept of orthogonality suggests. Therefore, contrarian studies that alter all aspects of the design, close replications, and simply choosing studies that include less-represented characteristics are not optimal approaches.

Where Has the Field Taken Itself?

A computer search of econometric work on advertising since 1984 (the publication date of the original meta-analysis) produced 27 listings, indicating a decline in the rate of publication in the field during the post-1985 period. Nine of the articles included either explicit estimates of advertising elasticities or the elements necessary to compute them (see Table 5).

The average short-term elasticity reported was .24, nearly identical to those reported by Assmus, Farley, and Lehmann (1984). Models without a carryover coefficient produced elasticities .32 larger than those models with carryover, which again are nearly equal to those reported in the 1984 meta-analysis.

The published studies tend to move the field in the direction we have identified as the next best study. The addition of the new categories of services and fresh foods expanded the research environment in new directions (e.g., perishable goods). Model specifications incorporated other marketing variables, though additive and other variable elasticity models were not used. Progress was not made in assessing the impact of disaggregate advertising. There is limited overlap in the journals represented by the original set of 90 articles and those containing the 9 new articles, however, and there is no indication that these developments were part of an orderly strategy. To apply our method effectively, we would need to update the calculations as each observation is added to the database, so the optimal new study would not be the same as the optimal first one. Because such a stepwise approach might not produce the best three next observations, we could extend the procedure to consider groups of observations.

Table 5
SUBSEQUENT STUDIES

Bjorndal, Trond, Fjell G. Salvanes, and Jorun H. Andreassen (1992), "Demand for Salmon in France: The Effects of Marketing and Structural Change," *Applied Economics*, 24 (September), 1027–34.

Helmuth, John A. (1987), "Dealership Automobile Demand: Advertising Elasticity and Optimality," *Adron Business and Economic Review*, 18 (Spring), 37–44.

Johnson, L.W. (1985), "Alternative Econometric Estimates of the Effect of Advertising on the Demand for Alcoholic Beverages in the United Kingdom," *International Journal of Advertising*, 4, 19–25.

Jones, Eugene and Ronald W. Ward (1989), "Effectiveness of Generic and Brand Advertising on Fresh and Processed Potato Products," *Agribusiness*, 5 (5), 523–36.

Popodopolis, Socrates (1987), "Strategic Marketing Techniques in International Tourism," *International Marketing Review*, (Summer), 71–84.

Radfar, Mehran (1985), "Effect of Advertising on Total Consumption of Cigarettes in the UK," *European Economic Review*, 29 (November), 225–31.

Tellis, Gerard J. and Doyle L. Weiss (1995), "Does Advertising Really Affect Sales? The Role of Measures, Models, and Data Aggregation," *Journal of Advertising*, 24 (Fall), 1–12.

Wagner, Joachim and J. Matthius Graf von der Schulenburg (1992), "Unobservable Industry Characteristics and the Innovator–Advertising–Maze Evidence from an Econometric Study Using Panel Data for Manufacturing Industries in the FRG 1979–1986," *Small Business Economics*, 10 (December), 315–26.

Zidack, Walter, Henry Kennucan, and Upton Hatch (1992), "Wholesale and Farm-Level Impacts of Generic Advertising: The Case of Catfish," *Applied Economics*, 24 (September), 959–68.

Note: These are the nine identified studies published after the 1984 meta-analysis that included advertising elasticities in their elements.

DOMAIN EXPANSION

The method presented here focuses on refining knowledge within a domain defined by factors for which variance has been observed. Because many insights can result from extending the domain to new variables, we offer two suggestions for their choice:

1. When designing the meta-analysis, researchers should try to conceptualize and report other variables that might matter and that they would like to have. For example, if all available studies involve packaged goods, researchers should include categories for theoretically or practically interesting alternative types of markets (durables, services, business-to-business).
2. When selecting a variable to add to the analysis, researchers should consider its likely explanatory power. Variables correlated with included variables offer little room for improvement, whereas variables nearer to orthogonal have greater potential.

CONCLUSION

We have presented a procedure to assess the potential value of a next study for improving knowledge of the impact of various technical and environmental conditions on a key model parameter (such as advertising elasticity). The approach can help provide both authors and editors some objective criteria of "contribution to knowledge" for new work in a relatively mature field. Our results suggest three key conclusions:

1. The choice of the characteristics of the next observation can make a dramatic difference in the reliability of the estimate of the parameters of the design variables in meta-analysis and, thus, on knowledge development.
2. Simply filling empty/sparse cells or including design variables that have been measured infrequently is not necessarily optimal.
3. Reducing collinearity among design variables is the key to improving knowledge.

Notice that, though the field's general resistance to exact or close replications is supported a priori, a posteriori they still add knowledge, albeit less than the optimal study does.

These results assume that all studies are equally costly and all variables are equally important or weighted. There is probably a cost hierarchy from least to most expensive, such as

1. Reanalyzing technical issues using data from existing studies (a much underutilized option),
2. Expanding the scope within studies in terms of (a) design and (b) analysis. Studies that examine multiple products and multiple settings help, and
3. Creating a new study. Although a new study is typically the most useful because it has more design freedom, it is also expensive and, therefore, not necessarily the most cost effective.

Note that one could incorporate formally the weights that reflect either cost or variable interest/importance by multiplying them by the Z matrix and then applying our method to the revised Z matrix.

This article has limitations. First, researchers must believe that meta-analysis has value. Problems such as unrepresentative samples (which can be caused by publication bias; Rust, Farley, and Lehmann 1990) render meta-analyses less useful. Still, there is presently no better way to cumulate knowledge formally, as is evidenced by its use in life-or-death situations involving medical treatments.

Second, the method focuses on expanding a body of knowledge rather than on creating new theories or fields of study. The boundaries of the field studied are determined by the phenomenon studied (which leads to the dependent variable in the meta-analysis design) and the space spanned by the design (independent) variables in the design matrix. Still, constructing the design matrix may suggest levels on which there are no observations (e.g., advertising studies in undeveloped or centrally planned economies, Internet advertisements).

Third, the method will not explicitly suggest exploring a field in really new ways. It may be that, for example, studying customer relations or value chains is more useful for understanding markets than examining the impact of advertising. Although philosophically the thrust of the method would encourage ventures into uncharted waters, it does not provide direct guidance as to which ventures to investigate.

Limitations aside, the approach suggests that authors, reviewers, and editors should assess how much a study adds to a body of knowledge about model parameters. An author benefits from having an objective way to position the contribution of an article in a relatively well-developed field. Articles then could be judged primarily on a priori design factors, such as contribution to knowledge, rather than on a posteriori factors, such as the strength, or worse, the statistical significance, of their results.

The method involves extending and making more precise a particular (quantitative) relation. The procedure neither robs the researcher of freedom of inquiry nor limits experimentation with really new models or methods. Rather, it helps guide attention to one means toward efficient knowledge production.

REFERENCES

Assmus, Gert, John U. Farley, and Donald R. Lehmann (1984), "How Advertising Affects Sales: Meta-Analysis of Econometric Results," *Journal of Marketing Research*, 21 (February), 65–74.

Belsley, David A. (1980), *Conditioning Diagnostics: Collinearity and Weak Date in Regression*. New York: John Wiley & Sons.

Farley, John U. and Donald R. Lehmann (1986), *Meta-Analysis in Marketing: Generalization of Response Models*. Lexington, MA: Lexington Books.

———, ———, and Alan Sawyer (1995), "Empirical Marketing Generalization Using Meta-Analysis," *Marketing Science*, 14 (September), G36–46.

Leff, Nathaniel and John U. Farley (1980), "Advertising Expenditure in the Developing World," *Journal of International Business Studies*, 11 (Fall), 65–79.

Mela, Carl and Praveen Kopalle (1998), "The Impact of Collinearity on Regression: The Asymmetric Effect of Negative and Positive Correlations," working paper, Business School, Notre Dame University.

Rust, Roland T., Donald R. Lehmann, and John U. Farley (1990), "Estimating Publication Bias in Meta-Analysis," *Journal of Marketing Research*, 27 (May), 220–26.

Silvey, S. D. (1969), "Multicollinearity and Imprecise Estimation," *Journal of the Royal Statistical Society*, Series B, 31, 539–52.