

USING CONJOINT ANALYSIS TO DETECT DISCRIMINATION: REVEALING COVERT PREFERENCES FROM OVERT CHOICES

Eugene M. Caruso
University of Chicago

Dobromir A. Rahnev
Columbia University

Mahzarin R. Banaji
Harvard University

In an effort to continue the development of methods to understand social cognition, we adopt a technique called *conjoint analysis* that mathematically deduces preferences from the implied tradeoffs people make when choosing between sets of attributes at varying levels. We asked 101 students to make a series of choices between prospective teammates in a trivia contest who varied on three dimensions relevant to the decisions (education, IQ, experience) and one dimension that was irrelevant (body weight). Although participants stated explicitly that weight had little impact on their decisions, weight actually accounted for more than 25% of the variance in their revealed preferences. Additional analyses demonstrated that participants gave up about 11 IQ points to have a thin rather than overweight teammate. We suggest that conjoint analysis can be a valuable tool for detecting and quantifying the social costs of covert attitudes that are not in sync with overt values.

Eugene M. Caruso, Booth School of Business; Dobromir A. Rahnev and Mahzarin R. Banaji, Department of Psychology. We acknowledge the support of a Graduate Research Fellowship from the National Science Foundation to Eugene M. Caruso, grants from the Harvard College Research Program and a Thomas Temple Hoopes Prize to Dobromir A. Rahnev, and grants from the NIMH, Third Millennium Foundation, and the Mind Science Institute to Mahzarin R. Banaji.

We would like to thank Max Bazerman, Dolly Chugh, and Kristina Olson for insightful comments on previous drafts of the manuscript, Katie Lancaster for assistance with the execution of the experiments, Wendy Berry Mendes for valuable statistical advice, and members of The Understand Seminar at Harvard for helpful discussions throughout this project.

Correspondence concerning this article should be addressed to Eugene M. Caruso, University of Chicago Booth School of Business, 5807 South Woodlawn Avenue, Chicago, IL 60637. E-mail: ecarus@chicagogsb.edu.

The well-documented limitations of introspection (e.g., Nisbett & Wilson, 1977) have led behavioral scientists to pursue alternative techniques for measuring mental operations to which individuals may not have conscious access. The past few years have seen a surge of studies that use indirect or implicit measurement techniques to capture various psychological constructs (for an overview, see Fazio & Olson, 2003). Although these techniques have produced an impressive array of findings and insights, there are still issues that need to be addressed to maximize the applicability and usefulness of indirect measures. For instance, interpreting the effects obtained from indirect measures is often ambiguous, as it is difficult to quantify the magnitude of a particular bias in meaningful and relevant ways (Nosek, Greenwald, & Banaji, 2007).

In an effort to develop an indirect measure that can capture and quantify the consequences of one's attitudes, we adopt a technique, *conjoint analysis*, that may hold the potential to expand the set of indirect measures and overcome some of the limitations of existing methods. The technique, which has been widely used in marketing research for more than 35 years, boasts high reliability, good predictive validity, and the ability to quantify relative preferences among the attributes tested. By presenting participants with realistic descriptions of problems similar to everyday choices they may actually face, the technique is well-equipped to provide access to social cognition that may not operate in conscious awareness.

Conjoint analysis was designed in part to bypass the limits of introspection, and we feel that its application in the domain of social judgment holds the potential to shed light on the serious consequences that result when one's decisions violate one's stated preferences. Although previous research on social judgments has shown that people are often inaccurate in their assessments of how much various personal attributes actually affect their attitudes about others (e.g., Nisbett & Bellows, 1977), these studies have not directly quantified the tradeoffs between those attributes. The ability to estimate precisely the simultaneous contributions of each attribute as a function of the importance of the other attributes affords conjoint analysis unique advantages over previous methods that have been used to understand social judgment.

Moreover, research on social attitudes has almost exclusively focused on the cost to the target of prejudice. Here we wish to test the reverse possibility that reliance on implicit attitudes and stereotypes may also have consequences for the bearer of such mental states themselves (Chugh, 2004). Capturing the magnitude of social category biases to those who hold the stereotypes could prove to be an effective means of discouraging reliance on such information. To the extent that people are unaware of the tangible personal costs that their attitudes carry, they are unlikely to be motivated to assess those attitudes and strive to change them. By allowing for the measurement of biases that are hard to reveal and hard to scale using other techniques, conjoint analysis can provide empirical insight into otherwise hidden preferences.

CONJOINT ANALYSIS: A BRIEF OVERVIEW

Conjoint methods of analysis were first developed in the 1960s by mathematical psychologists and statisticians Luce and Tukey (1964). Because people tend to be better at giving well-ordered preferences when evaluating options together ("con-

jointly") rather than in isolation, the method relieves a respondent from the difficult task of accurately introspecting about the relative importance of individual attributes for a particular decision (Green & Rao, 1971). That is, an individual may be hard pressed to state exactly how important the dimension of "price" is in a decision involving her honeymoon, and be unwilling or unable to place an accurate dollar amount on a month-long sojourn in Paris. However, she may more easily determine whether she prefers a month in Paris for \$5,000 to a month in Brussels for \$2,000. Stating a preference for one vacation or the other reveals a value for each. A decision favoring Paris implies that, all else being equal, she values a month in Paris at least \$3,000 more than a month in Brussels.

The ability to capture and quantify the relative importance of such real-world tradeoffs has led conjoint analysis to be the most broadly used technique in marketing research for measuring and predicting consumer preferences (Green, Krieger, & Wind, 2001; Orme, 2006). We tested whether the technique can be expanded beyond preferences for consumer products to study preferences for social categories. Before describing the current study, we first present a detailed overview of the general method.

Attributes and Level

Conjoint analysis asks participants to give an overall evaluation of product bundles that vary systematically on a number of attributes. For instance, to understand consumer preferences for televisions, a researcher could select price, brand, and screen size as the most important attributes. Next, the researcher chooses realistic levels of each attribute. The price points for TVs could be \$750, \$1000, \$1350, and \$1600; the available brands could be Sony, Panasonic, and Toshiba; the screen sizes could be 32", 37", and 42". Based on the selected attributes and levels, the researcher then creates profiles of possible product combinations for participants to evaluate in terms of preference, acceptability, or purchase likelihood.

Design

The specific set of these profiles must meet some critical requirements to produce optimal results. The minimal advisable number of profiles that participants need to evaluate is typically a multiple of the difference between the total number of levels and the number of attributes used. The set of profiles should be relatively balanced (each level of an attribute should appear an equal number of times) and orthogonal (each level of an attribute should appear together with each level of every other attribute approximately the same number of times).¹ The researcher can also choose to exclude the pairing of some levels and attributes to increase the realism of the test (a 42" Sony TV for only \$750 may be unrealistic). By having participants evaluate all chosen profiles sequentially, the technique can compute the importance of each single attribute that was presented.

1. A purely orthogonal design is not always possible, given the number of attribute levels and desired number of profiles. Measures are available to assess the quality of a particular design relative to a hypothetical orthogonal design (e.g., D-efficiency; Kuhfeld, Tobias, & Garratt, 1994).

Analysis

Linear regression (typically Ordinary Least Squares regression) is used to analyze the responses to the conjoint questions and to produce beta scores for each individual level of each attribute. These beta scores, also known as part-worth utilities, can be interpreted on two levels. Within every attribute, levels with higher part-worth values are preferred to levels with lower part-worth values. Among the attributes themselves, subtracting the level with the lowest beta score from the level with the highest beta score within each attribute provides an indicator of the relative importance of the attribute as a whole. The importance of the attributes can then be compared to one another to determine, for instance, the percentage of the total variance in preference that each attribute explains.

Variations

Several variations of conjoint analysis exist, each with specific strengths and weaknesses (Orme, 2003). For the present purposes, we will focus on two variants of Conjoint Value Analysis, both of which present all of the attributes in every conjoint question. In the *single concepts* variation, participants evaluate a series of single profiles and indicate the strength of their preference for each one. In the *pairwise comparisons* variation, participants evaluate a series of two profiles at a time and indicate the strength of their preference for one relative to the other. We use both of these variations, in addition to explicit ratings of preference, to test the stated and revealed importance of another person's weight in a social context.

METHOD

PARTICIPANTS

One hundred and one Bulgarian university students took part in a study on how people make choices. They imagined playing a trivia game of general knowledge in which teams of two competed to win a cash prize, and indicated their preference for a number of potential teammates that varied on several dimensions.

MATERIALS

Attributes and Levels. The potential teammates varied on four attributes. Three of these attributes—educational level, IQ, and previous experience with this trivia game—were meant to be diagnostic of ability to succeed in the game. The final attribute—weight—was meant to be nondiagnostic of actual ability. However, because of the pervasive stigmas and biases toward overweight people (see Puhl & Brownell, 2001), we expected that participants' choices would reveal a preference for thin teammates.

Each of these attributes contained a number of different levels, as shown in Table 1. For example, each teammate had an IQ of either 82, 93, or 104. The weight attribute was manipulated by including a picture of each teammate below the three

TABLE 1. Attributes and Levels for Trivia Teammates

Attribute	Level 1	Level 2	Level 3	Level 4
Educational Level	High School	B.A.	M.A.	
Intelligence (IQ)	82	93	104	
Previous Experience	Has never played	Has played 3 times	Plays every week	
Weight (Pictures)	Overweight Female	Overweight Male	Thin Female	Thin Male

other attributes. We used pairs of pictures from an existing database of photos that are used in weight Implicit Association Test (IAT) studies (Nosek, Banaji, & Greenwald, 2006), in which the same face was manipulated to appear to be either thin or overweight with as little change as possible to any structural features of the face itself.

We chose four of these pairs of faces to use in this study.² Each participant saw one thin male, one thin female, one overweight male, and one overweight female. So that no participant would ever see both the thin and overweight version of the same face, we ran two separate conditions. In one, participants saw the thin versions of Faces 1 and 2 and the overweight versions of Faces 3 and 4. In the other, participants saw the overweight versions of Faces 1 and 2 and the thin versions of Faces 3 and 4. Because the faces themselves varied on dimensions other than weight, it was crucial to have both conditions to isolate the effect of weight itself. Importantly, no significant differences between conditions were found on any of the key measures, so we collapsed across condition in the data reported below.

PROCEDURE

After reading the instructions, all participants first completed a series of 24 pairwise comparison questions. Two complete profiles of teammates were presented side-by-side and rated on a 9-point scale.³ Participants then rated 24 individual profiles of potential teammates in isolation, and indicated how likely they would be to choose this particular person as a teammate on a scale ranging from 0 (*not at all likely*) to 100 (*very likely*). Finally, participants were asked explicitly to rate how important each of the four attributes was for their decisions on scales ranging from 1 (*not at all important*) to 9 (*extremely important*).

RESULTS

A specialized software package designed for conjoint analysis studies (Sawtooth Software, 2006) generated the part-worth utilities and the importance of each attri-

2. Twenty-one university students from Bulgaria rated these pictures on perceived weight. The thin faces were rated as significantly lighter ($M = 140$ pounds) than the overweight faces ($M = 237$ pounds), $t(20) = 9.97, p < .0001$.

3. The profiles used for both the pairwise comparison and single concept procedures were generated from the Sawtooth Software Market Research Tools platform (Sawtooth Software, 2006; see www.sawtoothsoftware.com for detailed information).

TABLE 2. Importance of Attributes: Percentage of Variance Explained

	Education	Intelligence	Experience	Weight
Pairwise Comparisons	19%	25%	28%	28%
Single Concepts	19%	26%	30%	25%
Importance of Attributes: Explicit Ratings				
Explicit Ratings	4.95	6.40	5.24	2.48

bute for each participant. Four participants who did not complete all of the single concept questions, and four participants who based all ratings on just a single attribute, were excluded from those analyses.

Revealed Importance of the Attributes

Table 2 summarizes the importance of each attribute as computed by the two conjoint procedures. As can be seen, participants' responses revealed that the weight of the teammate accounted for 28% of the variance in the pairwise comparisons and 25% of the variance in the single concepts. No other single attribute accounted for significantly more variance ($ts < 1$), as weight was on par with IQ and previous experience.

Revealed Preference for Weight

To test for the direction of the weight preference, the part-worth utilities of the two thin and the two overweight faces were combined for each participant. We then coded each participant as either preferring a thin (positive value) or overweight (negative value) teammate. A single sample t -test with a hypothesized population mean of zero (i.e., no weight preference) revealed a clear preference for thin teammates compared to overweight ones in both the pairwise comparisons, $t(98) = 4.19$, $p < .001$, $d = 0.42$ and the single concepts, $t(92) = 3.86$, $p < .001$, $d = 0.40$.⁴ Seventy-three percent of people in the comparisons, and 61% of people in the single concepts, preferred a thin to an overweight teammate ($\chi^2s > 4.70$, $ps < .03$).

Explicit Preferences

Participants' stated preferences told a different story from their revealed ones. As shown in Table 2, participants reported that weight was the single least important factor in their choice (all pairwise $ts > 8.00$, $ps < .001$).

4. Technically, one underlying assumption for these t -tests is that all participants placed equal importance on all attributes other than weight. Although this assumption is unlikely to have been strictly satisfied because of the natural variation in the importance of the other attributes, such added noise would only make the weight bias harder to detect through these tests.

Reliability of Measures

To assess the reliability of the two conjoint procedures, we correlated the importance of each attribute across the two measures at the individual level. As can be seen in Table 3, inter-measure correlations between the pairwise comparisons and single concepts were significant for each attribute.

Once again, explicit ratings showed a somewhat different pattern. Table 3 shows that explicit ratings correlated significantly with the single concepts for all attributes *except* weight. Explicit ratings did correlate significantly with the pairwise comparisons on the weight attribute, but this correlation itself was significantly lower than the correlation between the pairwise comparisons and single concepts on this attribute ($z = 4.35, p < .001$). Overall, this pattern suggests that participants' stated preferences were a much better match for their revealed ones on the attributes of intelligence and experience than on the attribute of weight.⁵

Quantifying the Bias

Because conjoint analysis is a relative measure that yields beta weights for all of the levels within each attribute, it was possible to compare the importance of the weight preference to the importance of the other quantifiable attribute (IQ). To do so, we took the average beta score for all the participants at each level of IQ and computed the magnitude of the difference between the highest and the lowest score. We followed the same procedure for the weight attribute to determine the magnitude of the preference for thin to overweight faces. We then computed the importance of weight in terms of the importance of IQ (for comparisons and single concepts separately) by dividing the importance of weight by the importance of IQ and multiplying the result by the span within the IQ attribute ($104 - 82 = 22$ IQ points). Our pretest of the perceived weight revealed an average difference of 97 pounds between the thin and overweight faces, and this difference in weight equated to a difference in 12.31 IQ points in the pairwise comparison and 10.53 IQ points in the single concept procedure. Averaging across procedures, participants' behavior showed a relinquishing of 11.42 IQ points in order to have a thin, as opposed to overweight, trivia teammate.⁶

GENERAL DISCUSSION

Conjoint analysis is a flexible technique that can be adapted to reflect the way people make choices in the real world. Imagine an employer who is looking to hire

5. We expected the correlations between the explicit and conjoint measures to be high for education as well, but they were not (on average) higher than the correlations with weight. Debriefing revealed that many participants were not familiar with the educational levels presented, as Bulgaria had only recently incorporated these levels into its educational system.

6. It is important to keep in mind that the magnitude of this implied tradeoff depends on the magnitude of the difference between the attribute levels presented, as the importance of an attribute tends to increase when the span of its levels increases. We attempted to provide a relatively narrow and realistic range of IQ scores (less than one standard deviation above and below the measured mean IQ in Bulgaria of 93). However, another way to quantify the extent of the weight bias without this caveat would be to claim that the 11.42 point difference represents 52% of the range of available IQs to which participants were exposed.

TABLE 3. Inter-Correlations between all Measures

	Education	Intelligence	Experience	Weight
Pairwise Comparisons—Single Concepts	0.41**	0.73**	0.62**	0.71**
Pairwise Comparisons—Explicit	0.15	0.48**	0.57**	0.25*
Single Concepts—Explicit	0.23*	0.35**	0.58**	0.17

Note. **Correlation is significant at the 0.01 level (2-tailed). *Correlation is significant at the 0.05 level (2-tailed).

a new manager. The process of culling through a stack of résumés or interviewing individual candidates is similar to the process that our participants went through when evaluating potential teammates. The applicants likely differ on some standard criteria (education, experience, skills) that the employer needs to integrate to make the appropriate hiring decision. Research that manipulates the social category of an applicant has found that résumés with stereotypic White names receive more interviews than résumés with stereotypic Black names (Bertrand & Mullainathan, 2004). A similar study using conjoint analysis would allow for the quantification of that bias in terms of the other characteristics of the applicants, and would further allow for an estimate of bias at the level of the individual respondent.

By presenting participants with realistic choices in a format similar to one that they may actually encounter in their lives, conjoint analysis boasts high ecological validity and the ability to successfully predict individual behavior (e.g., Louviere & Timmermans, 1992). We believe that the technique's numerous strengths make it well-suited to cross over into the realm of the psychological study of attitudes and preferences more generally.

For one, the technique's ability to be adapted for use in different contexts affords tests of attitudes and behavior in specific, rather than generalized, circumstances. And although conjoint analysis is an indirect measure of attitudes, it does involve conscious and deliberative thought processes from its respondents. Because conjoint methods do not require participants to understand the structure of their preferences but only to be able to express them, they may provide a critical bridge between explicit and implicit measures of attitudes and biases (Nosek et al., 2007).

Of course, such a claim requires that the attitudes captured through conjoint analysis are in fact implicit, and more work needs to be done to justify such a claim. As an initial test of this possibility, we have found that the difference between revealed preferences from conjoint analysis and explicitly stated preferences does in fact correlate significantly with responses a standard weight-IAT ($r = .20$ for single concepts and $r = .28$ for pairwise comparisons, $ps < .05$; Rahnev, Caruso, & Banaji, 2007), which causes us to speculate that the disparity between the stated and revealed preferences in the current study stems in part from people's inability to have full knowledge of their own minds.

One reason that such attitudes might be particularly resistant to conscious awareness is that the importance participants place on the non-socially sensitive attributes changes dynamically in response to their automatic reaction to the socially sensitive one. That is, when comparing a thin teammate with a low IQ but lots of experience to an overweight teammate with a high IQ but no experience, participants may justify their preference for the thin teammate by deliberately reasoning that experience with the game is more important. When they subsequently compare a thin teammate with a Master's degree but no experience to an

overweight teammate with a high school degree but lots of experience, they may come to generate reasons why an advanced degree is likely to lead to success in the game. Indeed, people have been shown to recruit seemingly acceptable criteria (a job candidate's education) to justify decisions that would otherwise be seen as biased (hiring a White candidate over a more qualified Black candidate; Norton, Vandello, & Darley, 2004). Because people take selective advantage of inputs that best serve their interests (Dunning, 2001), participants may capitalize on this elasticity in their judgments to select thin teammates without being consciously aware of the influence that the teammates' weight had in their judgments (see also Dovidio & Gaertner, 2000; Schweitzer & Hsee, 2002).

The negative consequences of such selective ignorance might be felt by more than just the targets of discrimination. The decision makers themselves in our study paid a cost—11.42 IQ points, to be specific—for their bias, just as a hiring manager may pay a cost for hiring the less capable White candidate. In another study using conjoint analysis, we found that undergraduates behaved as if they were willing to give up \$3,249 (22% of the available range of salaries) to work for a male, as opposed to female, boss (Rahnev et al., 2007). By revealing the magnitude of the "stereotype tax" people have to pay for attitudes of which they may not be aware (Chugh, 2004), the process of conjoint analysis may prove to be a useful tool for increasing awareness of biases and thereby reducing the costs associated with them.

As with any method, conjoint analysis has its limitations. The method assumes simple additivity of preferences; the typical analyses imply that participants are following a compensatory decision rule; the specific attributes and levels included critically affect the quality and interpretation of the results; the recommended number of attributes that can be used within a given profile is relatively small. Although a detailed discussion of these drawbacks is beyond the scope of this initial investigation, applications of the method in social and cognitive psychology can capitalize on the rich conjoint literature in the field of Marketing to choose the specific method, procedure, and analysis plan to suit the needs of the study at hand (see Bradlow, 2005; Green et al., 2001).⁷

One exciting possibility is to specify interaction terms in a conjoint model. By incorporating two (or more) social categories into a single study, the relative importance of each individual attribute, as well as the interaction of them, could be revealed. For people who are relatively favored members of one category but relatively unfavored members of another (White females and Black males, for instance), conjoint analysis could help to determine the specific contexts in which it is advantageous to be in each category. In this way, a conjoint analysis among potential voters could shed light on the implied tradeoffs people are making when faced with a choice between candidates of different races, genders, and ages. We believe that the continual development, refinement, and application of indirect measurement techniques in general—and of conjoint analysis in social attitudes in particular—will prove valuable in understanding and quantifying the impact of numerous biases. In doing so, we hope such progress will raise awareness of both the social costs we impose on others and the personal costs that we incur ourselves when we act on attitudes of which we are unaware.

7. Further insight may also be gained from the *policy-capturing* methodology, a technique similar to conjoint analysis that has been used extensively to study the evaluative judgments of decision makers in organizational contexts (see Aiman-Smith, Scullen, & Barr, 2002; Karren & Barringer, 2002).

REFERENCES

- Aiman-Smith, L., Scullen, S. E., & Barr, S. H. (2002). Conducting studies of decision making in organizational contexts: A tutorial for policy-capturing and other regression-based techniques. *Organizational Research Methods, 5*, 388-414.
- Bertrand, M., & Mullainathan, S. (2004). Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American Economic Review, 94*, 991-1013.
- Bradlow, E. T. (2005). Current issues and a "wish list" for conjoint analysis. *Applied Stochastic Models in Business and Industry, 21*, 319-323.
- Chugh, D. (2004). Societal and managerial implications of implicit social cognition: Why milliseconds matter. *Social Justice Research, 17*, 203-222.
- Dovidio, J. F., & Gaertner, S. L. (2000). Aversive racism and selection decisions: 1989 and 1999. *Psychological Science, 11*, 315-319.
- Dunning, D. (2001). On the motives underlying social cognition. In A. Tesser & N. Schwarz (Eds.), *Intraindividual processes* (pp. 348-374). Oxford: Blackwell.
- Fazio, R. H., & Olson, M. A. (2003). Implicit measures in social cognition research: Their meaning and use. *Annual Review of Psychology, 54*, 297-327.
- Green, P. E., Krieger, A. M., & Wind, Y. (2001). Thirty years of conjoint analysis: Reflections and prospects. *Interfaces, 31*, 56-73.
- Green, P. E., & Rao, V. R. (1971). Conjoint measurement for quantifying judgmental data. *Journal of Marketing Research, 8*, 355-363.
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology, 74*, 1464-1480.
- Karren, R. J., & Barringer, M. W. (2002). A review and analysis of the policy-capturing methodology in organizational research: Guidelines for research and practice. *Organizational Research Methods, 5*, 337-361.
- Kuhfeld, W. F., Tobias, R. D., & Garratt, M. (1994). Efficient experimental design with marketing research applications. *Journal of Marketing Research, 31*, 545-557.
- Louviere, J. J., & Timmermans, H. J. P. (1992). Testing the external validity of hierarchical conjoint analysis models of recreational destination choice. *Leisure Sciences, 14*, 179-194.
- Luce, R. D., & Tukey, J. W. (1964). Simultaneous conjoint measurement. *Journal of Mathematical Psychology, 1*, 1-27.
- Nisbett, R. E., & Bellows, N. (1977). Verbal reports about causal influences on social judgments: Private access versus public theories. *Journal of Personality and Social Psychology, 35*, 613-624.
- Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review, 84*, 231-259.
- Norton, M. I., Vandello, J. A., & Darley, J. M. (2004). Casuistry and social category bias. *Journal of Personality and Social Psychology, 87*, 817-831.
- Nosek, B. A., Banaji, M. R., & Greenwald, A. G. (2006). <http://implicit.harvard.edu/>.
- Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2007). The Implicit Association Test at age 7: A methodological and conceptual review. In J. A. Bargh (Ed.), *Automatic processes in social thinking and behavior* (pp. 265-292). New York: Psychology Press.
- Orme, B. (2003). *Which conjoint method should I use?* Technical working paper. Sequim, WA: Sawtooth Software.
- Orme, B. (2006). *Getting started with conjoint analysis: Strategies for product design and pricing research*. Madison, WI: Research Publishers LLC.
- Puhl, R., & Brownell, K. D. (2001). Obesity, bias, and discrimination. *Obesity Research, 9*, 788-805.
- Rahnev, D. A., Caruso, E. M., & Banaji, M. R. (2007). *Conjoint analysis: A new method of investigating stereotypes*. Unpublished manuscript, Harvard University.
- Sawtooth Software. (2006). *Sawtooth Software Market Research Tools (Version 4.10)*. [Computer software]. Sequim, WA.
- Schweitzer, M. E., & Hsee, C. K. (2002). Stretching the truth: Elastic justification and motivated communication of uncertain information. *Journal of Risk and Uncertainty, 25*, 185-201.