

Stereotypes: from an illusory stability to an illusory belief account

Sofia Correia dos Santos (sosantos@fpce.ul.pt)

Faculdade de Psicologia e Ciências da Educação da
Universidade de Lisboa, Alameda da Universidade, 1649-
013, Lisbon, Portugal

Sérgio Moreira (d8281@students.iscte.pt)

Higher Institute of Social Sciences and Organizational
Studies (ISCTE), Avenida das Forças Armadas, room 224,
1649-026 Lisbon, Portugal

Leonel Garcia-Marques (garcia-marques@sapo.pt)

Faculdade de Psicologia e Ciências da Educação da Universidade de Lisboa, Alameda da
Universidade, 1649-013, Lisbon, Portugal

Abstract

The categorization literature has shown that concepts are unstable and context-dependent knowledge structures (e.g., Barsalou, 1987; 1989). Three studies show similar findings for stereotypes. Study 1 revealed only moderate levels of within-participants stability in the content of social categories. Study 2 showed that irrelevant attributes of social groups are incorporated into the groups' stereotype as long as they are contextually salient. Study 3 provided preliminary evidence suggesting that stereotype context sensitivity may be accounted for a belief illusion account.

Keywords: Stereotypes; Conceptual instability; social cognition

According to early abstractionist positions, knowledge structures play an important role in achieving cognitive economy (e.g., Rosch, 1978). Fulfilling such a role demands cognitive structures that are both constant and persistent. It is assumed that the need for cognitive stability coupled with the scarcity of cognitive resources leads to the neglect of much of the detailed information about exemplars of non-social categories. So, perceivers are highly insensible to variant characteristics of exemplars and are likely to use the invariant characteristics already represented in the category mental representation to go beyond the information given. In short, perceivers become chronic *abstractionists* (Frank & Bransford, 1971; Posner & Keele, 1968). Such abstractionist tendencies thus envisaged mental representations of objects as enduring mental entities that exhibit an impressive degree of constancy in the face of environmental turmoil. One corollary of such approach is that knowledge structures, at least under ideal measurement conditions, should exhibit high reliability across relatively extended periods of time within the same individual. Nonetheless, empirical evidence assessing the stability of common concepts and categories indicates much more instability than a classic abstractionist position would suggest. The same individual on two different occasions (24 hours or two months apart) exhibits only modest reliability in defining and characterizing common concepts (Barsalou, Spindler, Sewell, Ballato & Gendel, 1987), retrieving exemplars from common categories (Bellezza, 1984), classifying instances into categories (McCloskey & Gluksberg, 1978), and rating the typicality of instances relative to their parent categories (Barsalou, Sewell &

Ballato, 1986). Other research shows that common categories are largely context-sensitive, in that the immediate linguistic context biases both how typical an instance is judged to be as well as how fast it can be accessed (Roth & Shoben, 1983).

Should we expect mental representations of social categories to show the same degree of fluidity as non-social categories? At first sight, the answer to this question may seem quite obvious. Just consider the famous "Princeton trilogy" (Gilbert, 1951; Karlines, Coffman, & Walters, 1969; Katz & Braly, 1933). Although the level of consensus decreased somewhat across studies, successive generations of Princeton University students conveyed only slightly more benevolent versions of basically the same stereotypes. Moreover, abstractionist tendencies should make stereotypes self-perpetuating and highly resistant to change (e.g., Hamilton & Troler, 1986).

However, stereotype stability has been typically assessed by the correspondence between the attributes chosen to describe the social category across different studies (e.g., Devine & Elliot, 1995). Consequently those studies do not provide an answer to one key issue: Are stereotypes stable over time within the same individual? Although, some of us might be tempted to suggest a positive answer, the fact is that we simply lack the relevant empirical evidence to answer this question (for one exception, see Rothbart & John's, 1987 unpublished longitudinal study, described in Rothbart & John, 1993). So, we may have been neglecting important attributes of social information processing, such as the malleability and context sensitivity of stereotypes.

Studies in this paper were designed with three main goals in mind. Study 1 assesses whether stereotypes show the same type of malleability that has now been demonstrated for mental representations of non-social categories using longitudinal methodology (for a review, see Barsalou, 1989; Barsalou & Medin, 1986). Study 2 takes further the idea of malleability assessing to what extent are stereotypes context sensible knowledge structures. Study 3 provides preliminary evidence about what processes may be responsible for stereotypes malleability.

Study 1

The likelihood of selecting the same attributes to characterize a social category at two different points in time

(Barsalou et al., 1987) seems to be a particularly appropriate way to assess stereotype stability. In fact, attribute or property selection was the first procedure used to study stereotypes empirically (Katz & Braly, 1933), it was used in the series of studies that assessed stereotypes in different generations of the same student population to infer the temporal persistence of stereotypes (Gilbert, 1951; Karlins, Coffman & Walters, 1969; Katz & Braly, 1933), and it has remained quite popular ever since (for a review see Dovidio, Brigham, Johnson & Gaertner, 1996). However, none of these studies asked the same participants to characterize a social group at two different points in time. In Study 1, we asked participants about the cultural stereotype content – what “people in general” think about the social groups at stake (e.g., Devine & Elliot, 1995). Our primary goal was to assess stereotype stability as reflected in the communality of stereotypic attributes selection within individuals over time.).

Method

Participants Participants were 46 University of Lisbon students who volunteered for the study at the researcher’s request.

Design The design of the study was a 3 immigrants social categories (Gypsies, Gays, and African) x 2 verification task sessions (session 1 and session 2) within-participants factorial.

Pre-testing the adjective check-list A different group of 31 students from the same population that otherwise did not participate in the study were asked to give descriptions of three social groups (Gypsies, Gay people, and African immigrants). Participants were instructed to generate a list of attributes for each group, based on their cultural stereotypes. The nine most frequently mentioned attributes were selected for each group (excluding those overlapping in meaning). Whenever possible, attribute antonyms were added to the list. This task produced a final list of 43 personality traits¹.

Procedure All participants were tested twice with the second session following the first session by two weeks. To identify participants’ answers across sessions, we asked each participant to indicate his or her birthday date and that of his/her mother, assuring anonymity. Participants were given a booklet containing the instructions and experimental materials. Participants then had to choose and write down, from the full list of 43 traits, the five that best described each of the target groups (see Katz & Braly, 1933). At the second session approximately 2 weeks later, participants were again given the same instruction they received in the first session and completed the checklist again.

Results and Discussion

Aggregate Sample (Within-Item) Stability When the checklist methodology is used, stereotype stability is typically assessed by the correspondence between the attributes chosen to describe the social category across different sessions (e.g., Devine & Elliot, 1995). We followed this procedure to compare attributes chosen across the two sessions. Across sessions agreement was very high (the within-item correlations varied from .94 to .97). As in previous studies (e.g., Devine & Elliot, 1995) within item analysis (aggregating across participants) apparently support stereotype stability.

Within-participants stability To determine the degree of overlap in the attributes used to describe social groups by each participant across sessions, a common-element correlation was used (Bellezza, 1984). This measure of correlation represents the proportion of common to total items and varies between the values of 0 and 1. Mean overlap scores ranged in value from .48 to .60, indicating that only approximately half a participant’s trait selections for a category in one session were also chosen in the second session. These moderate levels of within-participant stereotypes are generally similar to those found with non-social categories, using a similar experimental paradigm (Barsalou et al., 1987).

To sum up, the overlap values indicated that there was only moderate correspondence between the category’s attributes selected in the two sessions. These results reverse the picture we obtained when stability was assessed within-item aggregating across participants. The evidences reported stand in stark contrast to abstractionist expectation. On the other hand, our evidences are well predicted by partial retrieval exemplar models (e.g., Smith & Zárate, 1992) and by global matching memory models (e.g., Ratcliff & McKoon, 1988). According to these models, salient contextual information may be incorporated in the knowledge assembling. Consequently, more than merely instable, these memory models see knowledge structures as highly context sensitive. In line with this prediction, study 2 tests if contextually salient stereotype irrelevant attributes may be incorporated in the stereotype assembling.

Study 2

Should we expect stereotypes to be also affected by contextually salient attributes, even when they are irrelevant to them? Study 2 tests this hypothesis by priming a non-stereotyped concept immediately before stereotype assembling by means of an unrelated linguistic task.

¹ Since participants almost always generated personality traits in the pre-test the final list included only traits.

Method

Participants Participants were 104 University of Lisbon students, who volunteered for the study at the researcher’s request.

Design The design was a 2 priming condition (intelligence or friendliness) x 2 stereotypes (computer programmer or childcare professional) x 2 personality traits (stereotyped and non-stereotyped) mixed factorial design, with the last factor within-participants.

Procedure First, participants completed the linguistic task by thinking in abstract and define in their own words one of two concepts, intelligence or friendliness. After, following Katz & Braly (1933), they selected, from a list of 45 traits², the five traits that best describe one of the two occupational groups included in the study, computer programmers or childcare professionals. Finally, they completed 14 relevant 9-point trait rating scales for the same group described. The 14 trait dimensions selected were the most frequently used to describe the stereotypes of the two groups, in the pre-test, because they could easily represent expectations subjects have about each group in respect to desirable and undesirable traits. Dimensions concerning primed concepts were also included.

Priming Task Primes were chosen from a pretest in which 20 other students were asked twice to perform traits judgments on 9-point trait rating scales (the fourteen ones mentioned above in procedure) based on each group stereotype. Based on these ratings, we then selected the least frequently mentioned trait for a group (irrelevant) that was simultaneously one of the most consensually agreed trait for the other group (relevant). So, by using computer programmers and childcare professionals as target groups, the primed concepts (intelligence and friendliness) became stereotype congruent or irrelevant, depending on the group. Priming task was presented as an unrelated experiment from Language Department of Lisbon University. Instructions and questions were formatted differently from the ones used in the supposed next experiment to convince participants that the experiments were indeed separate and unrelated.

Table 1: Percentage of stereotyped and non-stereotyped traits chosen for each priming condition.

Prime	Stereotyped trait		Non-stereotyped trait	
	Intelligent (Comp. progr.)	Friendly (Child pro.)	Friendly (Comp. progr.)	Intelligent (Child pro.)
Intelligent	92%	85%	11%	35%
Friendly	96%	96%	50%	11%
Fisher exact p test	p=.840		p=.006	

² Similar procedures to the ones described in study 1 were also followed to pre-test this new adjective check-list.

During the first part, participants were asked to give familiarity judgments for several words. This task was used as a way of giving more credibility to the priming task. In the second part, they were asked to define two words: a neutral trait (e.g., conservator) and a word that corresponded to the trait we wanted to prime (intelligence or friendliness).

Results and Discussion

Since we did not find any significant effects of prime on non-matching traits choices and ratings, these results will not be referred.

Stereotype Attribute Selection Task To determine the effect of primes on trait choices we used Fisher exact probabilities tests. Results (see Table 1) showed no effect of stereotyped primes on choices of stereotyped traits matching primes (p=.84). In contrast, non-stereotyped primes did increase the number of choices of non-stereotyped traits matching primes (p=.0061).

Trait Rating Task We computed a 2 primes (intelligence or friendliness) x 2 stereotypes (computer programmer or childcare professional) x 2 traits (stereotyped and non-stereotyped) mixed model ANOVA on traits ratings. A significant three-way interaction emerged [F (1,100)=13,01, p<.0005, Mse=1,159], showing that there is no effect of stereotyped primes on ratings of stereotyped traits matching primes (see Table 2). But non-stereotyped primes greatly augmented ratings of non-stereotyped traits matching primes. So, as predicted by non-“enduring abstractionist” accounts, the outcome of stereotype assembling is, at least in part, constituted by context-dependent information, even though they are irrelevant and not commonly associated with it.

Study 1 showed that stereotypes are instable knowledge structures and study 2 suggests that this instability may, at least in part, be due to context sensitivity. One way to explain stereotypes context sensitivity is by assuming that contextually salient information is incorporated in the assembled stereotype independently of its association with the category in memory. This is in accordance with the Source of Activation Confusion model (SAC; e.g., Ayers & Reder, 1998). The SAC predicts that, once we only have consciousness of using a concept and not of its source of activation, highly accessible distracters can be accepted as true memory without retrieval of specific memory traces. Our explanation is also in accordance with dual processes models such as Jacoby and Hollingshead (1990).

Table 2: Ratings of stereotyped and non-stereotyped traits, for each priming condition.

Prime	Stereotyped trait		Non-stereotyped trait	
	Intelligent (Comp. progr.)	Friendly (Child pro.)	Friendly (Comp. progr.)	Intelligent (Child pro.)
Intelligent	7.92	7.61	5.54	6.31
Friendly	7.77	7.96	6.54	5.65
ANOVA	F(1,100)=13,01, p=0,0005, Mse=1,159			

Specifically, once recognition may be based in processing fluency, most accessible concepts can be accepted as true memories. Once we used as material people names and objects instead of groups and attributes, third study can only be considered as a preliminary test to that belief account as a possible explanation for stereotypes context sensitivity. In fact, in order to rule up alternative explanations, we controlled the information participants had memorized, using material easier to manipulate and to control experimentally.

Study 3

Should we expect person mental representations to be affected by contextually salient objects, even when they are not associated to them in memory, as predicted by processes like source of activation confusion and processing fluency? If so, those same processes can be a viable explanation for stereotype context sensitivity too. Study 3 does a preliminary test to this hypothesis using an adaptation of Jacoby and Whitehouse (1989) and McElree, Dolan and Jacoby (1999) paradigms.

We predicted for contextually salient objects lower response accuracy and faster response times in the conditions where participants were unable to use memory monitoring outcome as a base for a recognition task.

Method

Participants 115 students from the University of Lisbon participating in this study for a course credit.

Design The design was a 5 lists x 2 response deadlines (1 second or 3 seconds) x 4 targets (Francisco and Alexandre and Bernardo and Guilherme) x 3 type of items (old target matched and old person mismatched and new) x 2 priming (prime and no prime). The priming is nested with the last two levels of type of item and all factors except the first two are within-subjects.

Material Four lists of 35 objects were constructed for this study. Each list belonged to a specific target and was composed by three types of objects: 1) 15 old objects (i.e., objects exclusive to a specific target presented in the memorizing and recognition phase of the study); 2) 10 new objects (i.e., objects that were only presented in the recognition phase); and 3) 10 old mismatch objects (i.e., objects that were associated in the memorizing phase to a target different from the one in the recognition phase). To assure that there were no material effects, we use 5 different quartets of lists as a between subjects condition.

Procedure Participants were instructed to memorize four lists of objects, each one belonging exclusively to one target person. After, participants were requested to do a recognition task for a total of 140 objects (35 for each target). To make the objects salient (i.e., to enhance objects processing), half of the new and old mismatch items were

primed subliminally with themselves immediately before the recognition trial. To control the use of memory monitoring outcome, for each of the 140 trial half of the participants had 1 second response deadline and the other half 3 seconds. For each trial both response accuracy and response time were recorded. To assure that participants clearly understood and were familiarized with the recognition task, a training phase preceded the test phase.

Results and Discussion

Response Accuracy We have run two mixed measures ANOVAs with 2 (prime) x 2 (response deadline) x 5 (stimulus replications), one for the new items and another for old mismatch items conditions. The ANOVA showed a significant prime main effect for the new items condition, $t(1, 105) = 13.568$, $p = 0.038$, $Mse = 0.006$ (one-tailed) ($M_{prime} = 0.86$ vs $M_{no-prime} = 0.88$) and a significant prime response deadline interaction for the old mismatch items condition, $F(1, 105) = 2.669$, $p = 0.036$, $Mse = 0.011$. Planned comparisons for the old mismatch items results showed a significant effect of prime in the 1 second response deadline condition, $t(1, 105) = 3.695$, $p = 0.023$ (one-tailed) ($M_{prime} = 0.433$ vs $M_{no-prime} = 0.472$) that disappears in the 3 seconds response deadline condition, $t(1, 105) < 1$ ($M_{prime} = 0.561$ vs. $M_{no-prime} = 0.543$).

In general, the data is accommodated by our predictions. Specifically, the results show that in the 1 second response deadline condition the prime condition led to more false recognition than the no-prime condition, for both new and old mismatch items. This supports the notion that, under more severe time constraints, participants base recognition on processing fluency. In the 3 seconds response deadline condition the prime effect remained reliable for the new items condition but disappeared for the old mismatch items. Once participants had time to trigger memory monitoring in the 3 seconds response deadline condition, one could expect that the prime should no longer affect both types of items. However, what our results suggest is that monitoring will only succeed when there is sufficiently diagnostic information available in memory – information that contradicts the fluency derived from priming (i.e., the case of old mismatched items). Naturally, this interpretation claims for further test of the specificities of memory monitoring processes.

Response time We have run four mixed measures ANOVAs 2 (prime) x 5 (lists), one for each combination of item type (i.e., new vs. old mismatch) with response deadline (i.e., 1 second vs. 3 seconds)³, on response time (RTs) for yes responses (i.e., false recognitions). The results (see Table 3) are similar for new and old mismatch objects.

³ Since data distributions for the two response deadlines are quite different we ran separated analysis for 1 second and 3 seconds conditions.

Table 3: Mean RTs for yes responses for 1 a 3 seconds response deadline conditions.

	1s		3s	
	new	mismatch	new	mismatch
Prime	636 (22)	656 (13)	1312 (113)	1194 (42)
No-prime	682 (15)	738 (10)	1273 (88)	1219 (34)

Note: Standard deviation in parenthesis

In 1 second deadline condition participants took less time to respond in the prime condition than in the no prime, $F(1,33)=4,225$, $p=0,048$, $Mse=7473$, and $F(1,51)=49,989$, $p=0,000$, $Mse=3762$, for new and old mismatch items. However, prime effect on RTs is no longer significant in the 3 seconds response deadline condition, both for new and old mismatch items (both $F_s < 1$).

These results are consistent with our predictions. “Yes” responses are facilitated by the prime only in the condition where participants base recognition on processing fluency (i.e., 1 second response deadline condition). Whenever participants memory monitoring is efficient, the effects of the prime in the “yes” responses cease.

To sum up, the accuracy and RTs data supports the assumption that representations context sensitivity may be due to processes like source of activation confusion and processing fluency. More specifically, this preliminary study suggests that contextually available cues may be incorporated in category assembling independently of their association with the category in memory.

General Discussion

Study 1 and 2 suggest that considerable instability in stereotype knowledge is possible. Moreover, study 2 also suggests that stereotype instability may mainly be due to stereotypes context sensitivity. At last, study 3 provides some preliminary evidence showing that context sensitivity may be due to processes like source of activation confusion and processing fluency.

Although our evidences are in line with the research literature in the cognitive psychology field, convergent data on social cognition field for stereotypes is considerably inexistent. Some exceptions can be found. For instance, Garcia-Marques & Mackie (1999) have shown that the incongruent exemplars presented previously to stereotype assembling affect variability judgments. Coats & Smith (1999) provided data showing that subtypes descriptions are affected by the most accessible exemplars. Garcia-Marques, Santos and Mackie (in press) showed with a different paradigm that manipulation of context stability affects the stability of stereotype assembling across sessions. These studies strengthen the claim that stereotypes, as non-social knowledge structures, are context sensitive. Consequently, also strengthen the urge to study possible explanations for this phenomenon.

Processes responsible for stereotypes context sensitivity

Abstractionist positions clearly clash with the claim that stereotypes are context sensitive knowledge structures. However, this claim fits quite well with non-“enduring abstractionist” views. We propose that, according to these alternative theoretical accounts, there are at least two possible processes that may be responsible for stereotypes context sensibility: parallel matching retrieval and belief illusion.

Some memory models (e.g., Ratcliff and McKoon, 1988; Smith and Zárate, 1992) assume that malleability is to be expected and, moreover, that is greatly due to context sensitivity. Specifically, these models state that cognitive economy does not impose input simplification and, consequently, does not impose cognitive stability, at least in any austere degree. Instead, cognitive economy may be achieved through highly efficient retrieval processes like parallel matching. In other words, concept assembling may result from parallel matching process between salient contextual information and information in memory associated to the contextual cues. Consequently, the assembled concept is expected to be highly contextualized. Retrieval assumptions of this nature are present, for instance, in Smith and Zárate (1992) exemplars model (partial retrieval process) and Ratcliff and McKoon (1988) global matching memory model (compound cues process), and are also totally in line with the recently developed situated cognition (Barsalou, 2002) and social situated cognition (Smith & Semin, 2004) views.

Stereotypes context sensitivity can also be explain by a belief illusion account. For instance, the SAC model (e.g., Ayers & Reder, 1998) and Jacoby and Hollingshead (1990) predict that, because of the source of activation confusion and processing fluency respectively, highly salient information may be accepted as true memory even without retrieval of specific memory traces. Applied to stereotypes, contextually salient cues can become part of the assembled stereotype independently of their association with the stereotype in memory. Similarly to the previous account, this would result in instable and context sensible representations. These two processes are probably not mutually incompatible explanations for stereotypes context sensitivity. For instance, whenever a person for some reason (e.g., lack of resources) is not allowed to integrate memory monitoring outcome in the assembled stereotype, we predict that stereotype context sensibility may be mainly due to a belief illusion. However, whenever memory monitoring outcome can be used in the assembling process, we predict that both context sensible parallel matching retrieval and belief illusion are responsible for stereotypes context sensitivity.

Summing up, first two studies described a phenomenon almost unexplored – the context sensitivity of stereotypes. Although there is some recent work that strengthens this claim, further tests still remain to establish it broadness. Third study provides preliminarily evidence that allows regarding belief illusion account as a possible explanation of

stereotypes context sensitivity. Also, further tests are needed to extend these evidences to the stereotype literature and to clearly establish the conditions constraining this process occurrence.

Sometimes we live in illusion. Some illusions seem to be inherent to living adaptively. This work describes one of those cases – we have been fooled by stereotypes stability as a fact and, ironically, now that we began to question this fact, not just stereotypes seem flexible as this flexibility seems to be a result of an illusory belief!

References

- Ashmore, R.D., & DelBoca, F.K. (1981). Conceptual approaches to stereotypes and stereotyping. In D.L. Hamilton (Ed.), *Cognitive processes in stereotyping and intergroup behavior*. Hillsdale, NJ: Erlbaum.
- Ayers, M.S., & Reder, L.M. (1998). A theoretical review of the misinformation effect: Predictions from an activation-based memory model. *Psychonomic Bulletin & Review*, 5, 1-21.
- Barsalou, L.W. (1989). Intra-concept similarity and its implications for inter-concept similarity. In S. Vosniadou and A. Ortony (Eds.), *Similarity and analogical reasoning*. New York: Cambridge University Press.
- Barsalou, L.W. (2002). Being there conceptually: Simulating categories in preparation for situated action. In N.L. Stein, P.J. Bauer, & M. Rabinowitz (Eds.), *Representation, memory, and development: Essays in honor of Jean Mandler*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Barsalou, L.W. & Medin, D.L. (1986). Concepts: Static definitions or context-dependent representations? *Cahiers de Psychologie Cognitive*, 6 (2), 187-202.
- Barsalou, L.W., Sewell, D.R., & Ballato, S.M. (1986). *Assessing the stability of category representations with graded structure*. Unpublished manuscript.
- Barsalou, L.W., Spindler, J.L., Sewell, D.R., Ballato, S.M., & Gendel, E.M. (1987). *Assessing the stability of category representations with property generation*. Unpublished manuscript.
- Bellezza, F.S. (1984). Reliability of retrieval from semantic memory: Information about people. *Bulletin of the Psychonomic Society*, 22, 511-513.
- Coats, S., & Smith, E.R. (1999). Perceptions of gender subtypes: sensitivity to recent exemplar activation and in-group/out-group differences. *Personality and Social Psychology Bulletin*, 25 (4), 515-526.
- Devine, P.G. & Elliot, A.J. (1995). Are racial stereotypes really fading? The Princeton trilogy revisited. *Personality and Social Psychological Bulletin*, 21, 1139-1150.
- Franks, J.J., & Bransford, J.D. (1971). Abstraction of visual patterns. *Journal of Experimental Psychology*, 1, 65-74.
- Garcia-Marques, L., & Mackie, D. M. (1999). The impact of stereotype-incongruent information on perceived group variability and stereotype change. *Journal of Personality and Social Psychology*, 77, 979-990.
- Garcia-Marques, L., Santos, A.S., & Mackie, D.M. (in press). Stereotypes: Static definitions or dynamic knowledge structures? *Journal of Personality and Social Psychology*.
- Gilbert, G.M. (1951). Stereotype persistence and change among college students. *Journal of Abnormal and Social Psychology*, 46, 245-254.
- Hamilton, D.L., & Trolie, T.K. (1986). Stereotypes and stereotyping: An overview of the cognitive approach. In J. Dovidio, & S.L. Gaertner (Eds.), *Prejudice, discrimination, and racism*. New York : Academic Press.
- Jacoby, L. L., & Hollingshead, A. (1990). Toward a generate/recognize model of performance on direct and indirect tests of memory. *Journal of Memory & Language*, 29, 433-454.
- Jacoby, L.L., & Whitehouse, K. (1989). An illusion of memory: False recognition influenced by unconscious perception. *Journal of Experimental Psychology: General*, 118, 126-135.
- Karlins, M., Coffman, T.L., & Walter, G. (1969). On the fading of social stereotypes: Studies in three generations of college students. *Journal of Personality and Social Psychology*, 13, 1-16.
- Katz, D. & Braly, K.W. (1933). Racial stereotypes of 100 college students. *Journal of Abnormal Social Psychology*, 28, 280-290.
- Katz, D., & Braly, K.W. (1935). Racial prejudice and racial stereotypes. *Journal of Abnormal Social Psychology*, 30, 175-193.
- McCloskey, M.E. & Glucksberg, S. (1978). Natural categories: Well defined or fuzzy sets? *Memory and Cognition*, 6 (4), 462-472.
- McElree, B., Dolan, P.O., & Jacoby, L.L. (1999). Isolating the contributions of familiarity and source information to item recognition: A time course analysis. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25, 563-582.
- Posner, M.I., & Keele, S.W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, 77, 353-363.
- Ratcliff, R. & McKoon, G. (1988). A retrieval theory of priming. *Psychological Review*, 95, 385-408.
- Rosch, E. (1978). Principles of categorization. In E. Rosch, & B.B. Lloyd (Eds.), *Cognition and categorization* (pp. 27-48). Hillsdale, NJ: Erlbaum.
- Roth, E.M. & Shoben, E.J. (1983). The effects of context on the structure of categories. *Cognitive Psychology*, 15, 346-378.
- Rothbart, M., & John, O.P. (1993). Intergroup relations and stereotype change: a social-cognitive analysis and some longitudinal findings. In P. M. Sniderman, P. E. Tetlock & e. G. Carmines (Eds.), *Prejudice, politics and the American dilemma*. Stanford, CA: Stanford University Press.
- Smith, E.R., & Semin, G.R. (2004). Socially situated cognition: Cognition in its social context. *Advances in Experimental Social Psychology*, 36, 53-117.
- Smith, E.R., & Zárate, M.A. (1992). Exemplar-based model of social judgments. *Psychological Review*, 99, 3-21.