

A 100 years of certitude? Social psychology, the experimental method and the management of scientific uncertainty

S. Alexander Haslam* and Craig McGarty

Division of Psychology, Australian National University, Australia

[This paper, which is jointly authored by one of the current editorial team, was entirely dealt with by the previous editorial team.]

For at least 100 years the experimental method has been used to add scientific rigour to the process of conducting social psychological research. More specifically, experiments have been used to reduce methodological uncertainty surrounding the causal relationships between variables. In this way the method has proved particularly useful in demonstrating the impact of social contextual variables over-and-above individual differences. However, problems with the method have arisen because over time experimentalists have tended (1) to define uncertainty too narrowly, (2) to emphasize uncertainty reduction, but (3) to neglect the equally important process of uncertainty creation. This has contributed to the normalization of social psychology as a science but also made the discipline more conservative and circumscribed. It is argued that experimentalists need to address broader metatheoretical and political uncertainties in order to rediscover the experiment's potency as a tool of revolutionary and progressive science.

Social psychology is commonly regarded as having been founded in 1898—the publication date of Triplett's experimental research into the effects of competition on children's performance in a reel-winding task.¹ Triplett was inspired by the informal observation that, despite some individual differences, racing cyclists generally achieved better times on laps of a circuit when they had other cyclists pacing them. Consistent with this observation, he found that children wound fishing reels faster in the presence of other children winding reels in the same room.

Having reached the end of the first full century of empirical work in the field and the start of a new millennium, it would therefore appear timely to consider the contribution of the experimental method to the character of social psychology. The review is also worth conducting in light of the sustained attacks that critics both inside and outside social psychology have made upon the experimental method. In some instances experiments are seen more as a symptom than as a cause of broader

* Requests for reprints should be addressed to Alex Haslam or Craig McGarty, Division of Psychology, The Australian National University, Canberra, ACT 0200, Australia (e-mail: Alex.Haslam@anu.edu.au or Craig.McGarty@anu.edu.au).

¹ Haines and Vaughan (1979) have pointed out that Triplett's study was not in fact the first experiment in social psychology; they give the credit for this to Binet and Henri, 1894.

Table 1. Types of uncertainty associated with social psychological research and common forms of uncertainty management

Type of uncertainty	Typical cause of uncertainty	Typical goal of uncertainty management	Typical means of achieving uncertainty management goal
Ontological			
Statistical			
Descriptive	Individual differences; random variation; measurement error	Uncertainty measurement	Use of descriptive statistics to describe dispersion around mean
Inferential	Random variation; error	Uncertainty measurement	Use of inferential statistics to make probability statements
Methodological			
Internal	Equivocal status of independent variable	Uncertainty reduction	Experimental control
External	Sampling error; equivocal status of dependent variable	Uncertainty reduction	Appropriate sampling; theory development
Epistemological			
Social/political	Assumptions, values and theories of researchers and research community	Uncertainty creation	Debate; challenging research

problems in social psychology, but criticism still typically points to the central role of the experiment in establishing the nature and priorities of the discipline (e.g. see Billig, 1997; Gergen, 1997; Harré, 1997).

Our goal in this study is not, however, to respond to such criticism directly. As a more constructive (and original) offering, we present instead a model of the *research process* that attempts to account for some of these problems and that points to ways in which their aetiology can be better understood. At the heart of our argument is an assertion that significant problems with social psychology are not products of defective methodology, but rather arise from a limited and faulty conceptualization of the research process.

In essence, the argument we seek to develop is that psychological (and all other) research can be construed broadly as a process of *uncertainty management*, and that problems with that research arise when uncertainty is badly managed. The two particular forms of bad management we focus on are first, a tendency to define uncertainty too narrowly, and second, a tendency for researchers to emphasize uncertainty reduction but to neglect the equally important process of uncertainty creation.

As an extension of these arguments, we contend that doing good, useful and impactful research is not primarily a matter of choice of methodology. Moreover, we

suggest that debates about the appropriateness and utility of the experimental method may at times distract researchers from social psychology's more pressing problems. In this way, such debates may actually become *part* of the problem—at least to the extent that they focus attention on only one of the forms of uncertainty that researchers confront and not necessarily the one that is most interesting or far-reaching.

In developing this argument we start by outlining the uncertainty management model of the research process (following Haslam & McGarty, 1998; see Table 1). We then move on to look at some examples which examine this process in action.

The uncertainty management model

The nature of statistical and methodological uncertainty

In general terms, *uncertainty* exists when people suspect that there is a high probability that their beliefs about physical and social reality are inaccurate. The uncertainty management model of research practice in psychology suggests that there are two main types of uncertainty that are regularly dealt with by psychologists: *statistical uncertainty* and *methodological uncertainty*. Broadly speaking, statistical uncertainty relates to doubts over *what* a particular finding is and methodological uncertainty relates to what that finding *means*. The customary partition of results and discussion sections in reports of empirical research reflects this division. That is, results sections generally seek to define the appropriate level of statistical uncertainty and discussion sections deal with other forms of uncertainty, with a particular focus on methodological uncertainty.

As Tukey (1991) pointed out, it is for the purpose of quantifying statistical uncertainty that almost all psychological research involves the use of statistics. Specifying forms of statistical uncertainty further, it is apparent that statisticians routinely distinguish between two types of statistics: those that describe the properties of particular sets of data and those that allow researchers to make inferences about the probable relationship between a set of data and the underlying reality to which it relates. Each of these statistics is associated with a different form of statistical uncertainty: *descriptive uncertainty* and *inferential uncertainty*, respectively.

Descriptive uncertainty arises from the fact that when psychologists study a phenomenon they need to make multiple observations and these observations invariably produce different data. As a result of this, any descriptive statement must always incorporate some doubt. In all psychological research there is always a certain amount of descriptive uncertainty arising from the fact that it is rare for all research participants to respond in an identical fashion. Indeed, even if the *same* person completes the same task on two separate occasions it is common for their performance to be different on the two occasions.

To develop these ideas in terms of a working example, we can think of researchers who are interested in the productivity of people in the workplace and who want to examine the stress levels of employees who work under artificial light (e.g. Mayo, 1933; Roethlisberger & Dickson, 1939). Part of such an investigation might involve

administering some psychological tests to workers whose stress is quantified by assigning them a score between 0 and 100. To summarize, the level of stress in this sample the researchers might then compute the group's mean score on this scale. Descriptive uncertainty in this case arises from variation in observations around this measure of central tendency and may be attributed to random error or other factors (e.g. measurement error or individual differences). Whether descriptive uncertainty arises from random error or not, it can be quantified by some measure of dispersion—for example, standard deviation around the mean. The greater the dispersion in the scores obtained from the research participants, the greater the uncertainty and the less confident any researchers would be in the ability of the mean to describe those participants appropriately. Here uncertainty would only be removed completely if all participants obtained an identical score.

Descriptive uncertainty is random variation in the sample over which the experimenter has no control and for which no particular explanation is proposed. Inferential uncertainty, on the other hand, relates to a judgment about whether it is plausible that results like those obtained could have been produced by a random process (such as drawing random samples from the same population). Inferential uncertainty prevails when any meaningful variation present in a piece of research cannot be distinguished from the random variation expected in a relevant population. When inferential uncertainty is low, researchers suspect that it is not plausible that the variations obtained are because of random variation of a specified type. Inferential uncertainty manifests itself in many different forms and the particular form depends largely on the research question that is being addressed. Broadly speaking, it depends upon three things: (1) the amount of information that researchers have to indicate that something interesting is going on; (2) random error; and (3) the sample size (or, more specifically, the square root of the sample size).

In the sort of research to which our artificial lighting example relates, researchers would almost definitely want to make a number of inferences and each of these would be associated with some uncertainty. For example, they might want to know whether people who work under artificial light differ in their levels of stress from those who work under natural light. Inferential uncertainty could be eliminated by sampling everyone in these two populations and establishing whether they differ. In practice, of course, because this strategy is impractical (or impossible) researchers would only take a sample of the people from different types of working condition. They would then set about quantifying the amount of uncertainty associated with a statement suggesting that the two samples have not been drawn from the same population. As with all inferential statistics, this quantification of uncertainty is an estimation. It is stated as a probability, where a lower probability signifies greater certainty about the inference being made.

In dealing with descriptive and inferential uncertainty, a researcher is often able to quantify the uncertainty accurately, but unable to remove it. Statistics allow researchers to make precise descriptive and inferential statements, but the elimination of statistical uncertainty is best achieved by methodological (and occasionally statistical) procedures. Among other things, statistical uncertainty can be minimized by increasing the power of tests—for example, by increasing the sample size, or by reducing the amount of random variation. In some instances, random variation can

be reduced by using a within-participants, rather than a between-participants, design or by developing a more reliable measure of the phenomenon under investigation.

However, it is important to note that if statistical uncertainty were *completely* removed from research, it is likely that the psychological issues being addressed by the researcher would be extremely uninteresting. This is because the research would have to involve *complete sampling* of a *completely uniform* population. It is actually hard to think of any research using psychological statistics that would satisfy these requirements and would not be utterly trivial. This observation reflects an extremely important tension in psychological research—the much-vaunted fact that statistical significance is not the same as psychological or theoretical significance. Obtaining a statistically significant result tells researchers and their audience that descriptive and inferential uncertainty are tolerably low. But in the most extreme case statistical significance can amount to a statement of the blindingly obvious (e.g. that all first-year psychology students have one head). This is clearly problematic, at least to the extent that the role of psychological research is to contribute to and extend knowledge, not just to reproduce it.

Accordingly, an obsession with the reduction of statistical uncertainty can be counterproductive. This is one reason why in recent years the wisdom of a hypothesis-testing approach to statistical inference has come into question (Cohen, 1994, 1995; G. Hammond, 1996; Judd, McClelland & Culhane, 1995; Syvanteck & Ekeberg, 1995). Responses to the problem of how to handle statistical uncertainty include accompanying treatments of inferential uncertainty (significance) with details of confidence intervals (Smithson, 1999) and treatments of effect size (Cohen, 1977). The aim of including information about effect size is to be sure not only that chance is not a plausible cause of research findings, but also that they are of a sufficient size to be interesting or useful (criteria which will vary from area to area and problem to problem).

Yet statistical uncertainty is only one form of uncertainty faced by researchers. A second form, in many ways more serious, relates to confidence that research procedures are directly addressing the question in which researchers are interested and are allowing them to answer it. This second form of uncertainty can be referred to as *methodological uncertainty* and, as with statistical uncertainty, it has two key forms (Smithson (1999) adds a third which he calls ‘design uncertainty’). *Internal uncertainty* relates to a researcher’s confidence that any piece of research shows what it is believed to show. This term has most in common with the notion of internal validity: the idea that in experimental research the observed effect of a particular manipulation has been correctly interpreted. As outlined by Campbell and Stanley (1963), a large number of features of experimental design can contribute to uncertainty of this form. A researcher’s ability to understand why the manipulation of an independent variable brings about change in a dependent variable can be compromised by the effects of maturation, history, reactivity, selection and various other experimental confounds. In a study of occupational stress, a researcher’s ability to confidently interpret the effects of making one randomly selected group of workers work under non-artificial light would be reduced if participants’ behaviour were affected by their knowledge that they were taking part in psychological research (as it was, for example, in the Hawthorne studies; Roethlisberger & Dickson, 1939).

However, unlike internal validity, the concept of internal uncertainty is not only relevant to experimental and quasi-experimental research. This form of uncertainty is also prevalent in survey research where independent variables are not directly manipulated. In survey research, whenever a relationship between two variables is observed it is inappropriate to draw any conclusions about the effect that one variable has on the other. This means that internal uncertainty—pertaining to the interpretation of that relationship—must remain high. Nonetheless, such uncertainty can be reduced in cases where this type of research fails to reveal a relationship between two variables. This is because a lack of correlation allows the researcher to *conclusively rule out* certain interpretations (provided that suppressor variables are not present). For example, if research fails to find a relationship between lighting conditions and stress then internal uncertainty about the psychological interpretation of such a relationship is reduced. In this case if someone suggests that workers become stressed because they get headaches from poor lighting, we can be reasonably certain they are wrong.

Nonetheless, certainty about the fact that manipulation of an independent variable is responsible for an effect does not eliminate debate about internal uncertainty. This is because the psychological status and impact of experimental manipulations is always open to question. In Triplet's studies, when additional reelers were included to accompany a solo reeler, was this a manipulation of co-presence, camaraderie or freedom from isolation? And did it achieve its effects by reducing feelings of loneliness, by changing the context of social comparison, or by introducing an element of competition? Debates of this form surround just about every robust effect in social psychology.

The second type of methodological uncertainty is *external uncertainty* and this arises where researchers are not sure that the results they have obtained can be generalized to the population of interest. In many respects this is the most generic form of uncertainty encountered in psychological research, partly because it is related to all of the other forms of research uncertainty that we have discussed so far. Obviously, researchers do not want to generalize their results unless inferential and internal uncertainty are low: they want to be confident that they are not owing to chance and that they reflect the impact of an independent variable rather than some confound. It is also easier to generalize results when descriptive uncertainty is low (mainly because the researcher will be more confident about what those results are).

In the case of our artificial lighting example, researchers will generally have less external uncertainty if they show that the result applies under a wide range of conditions. For instance, they will be more confident that they can generalize the results if these are obtained using representative random sampling techniques and a wide range of testing conditions and if they have been predicted on the basis of a well-developed theory (e.g. one which asserts that unnatural environments are psychologically damaging). On the other hand, external uncertainty will be high if researchers use naïve empiricism to generalize research findings without recourse to theory (Turner, 1981), or if they draw conclusions on the basis of samples that are not representative of the population of interest on theoretically relevant variables.

In social psychology, debate about external uncertainty is just as heated as debate surrounding issues of internal uncertainty. Here, concerns about generalizability

typically focus on issues of sampling and theory development. Ideally, then, researchers seek to explain the causal links they observe in experimental samples in terms of well-specified theoretical principles and then make generalizations by applying those principles to a broader population.

Again, though, each of these elements is open to scrutiny, as Triplett's research has shown. To what broader population are the theoretical principles tested in studies of fishing-line reelers relevant? How do predictions of social facilitation square with contradictory evidence of social loafing (Harkins, 1987; Haslam, 2001; Karau & Williams, 1993)? Is the analysis of social facilitation too circumscribed, or not circumscribed enough? Once more, questions of this form are raised in relation to just about every robust effect in social psychology.

It is important to note that the concept of methodological uncertainty that we have introduced here differs from the concept of validity in focusing attention on the role (and power) of researchers and the research community in the research process. Among other things, it reminds us that conclusions about seemingly objective features of research can change depending on the state of scientific knowledge in a field. This, for example, is what happened in the Hawthorne studies as researchers started to question some of the early explanations of the effects of manipulating factory working conditions (e.g. lighting levels). Certain inferences about these findings *always were* invalid: it was only the researchers' level of uncertainty that changed. However, it was the level of uncertainty about the research, not its validity, that best characterized the state of scientific knowledge at any particular point in time. Our terminology therefore reinforces the fact that scientific knowledge is contingent on human understanding (and ignorance), not facts in the abstract.

Managing statistical and methodological uncertainty

Notwithstanding the new terminology used to frame them, the issues discussed above will be familiar to many readers. Indeed, their familiarity is very much our point. The issues are familiar because most of what researchers do, both when they conduct research and when they subject their own and other people's research to scrutiny, comes down to issues of uncertainty management. Even more precisely, we can state that provided that statistical uncertainty is sufficiently low—so that researchers are reasonably confident they are not discussing a chance finding—then most remaining debate surrounds issues of methodological uncertainty.

We have argued that these four forms of uncertainty represent the backbone of most researchers' training in methodology and statistics. Accordingly, it is upon these that most research in psychology focuses. Importantly too, psychologists have developed reasonably well-defined and widely shared methods for managing them. Indeed, generally speaking, training in issues of statistics and methodology orients researchers to the twin goals of (1) measuring statistical uncertainty, and (2) minimizing methodological uncertainty (Haslam & McGarty, 1998).

Both experimental and non-experimental methods typically allow researchers to measure statistical uncertainty very well. However, it is in relation to the second of these goals—minimizing methodological uncertainty—that the experimental method has proved particularly important in psychology. This arises from the fact that in

experiments internal uncertainty can be reduced because, where there is low inferential uncertainty, effects observed on dependent measures can always be attributed to the manipulation of an independent variable. On the other hand, correlational studies (i.e. surveys and quasi-experiments) must always admit some internal uncertainty because they do not allow researchers to draw similar conclusions about causation.

We noted above that researchers' ability to quantify statistical uncertainty is high, but that it is often difficult to control. The very opposite is true for methodological uncertainty. No procedures have ever been agreed upon for measuring or estimating methodological uncertainty, but a lot is known about how to reduce it. Internal uncertainty can be minimized or eliminated by using appropriately controlled experimental designs, while external methodological uncertainty can be reduced through appropriate sampling and the development of psychological theory.

Attempts to highlight these different forms of uncertainty raise obvious questions about the manner in which they should be managed *collectively*. In practice, strategic decisions about the management of uncertainty generally involve *trade-offs*: being sure about some things increases uncertainty about others. For example, if researchers reduce descriptive uncertainty by using a very homogenous sample of participants in their research, they may increase external uncertainty about the applicability of their findings to a broader population. However, as Tajfel and Fraser (1978) noted, if researchers attempt to reduce external uncertainty by studying a large number of people in a naturalistic setting, internal uncertainty may be increased because it is impossible to exert control over the independent variable (e.g. through random assignment of participants to conditions).

Beyond ontological uncertainty

In philosophical terms, the uncertainties we have dealt with thus far can be thought of as *ontological* (i.e. they relate to the more tangible features of the research process, such as the methodology, the statistical techniques, the sampling method). However, there are also other *epistemological* uncertainties that are far less tangible. The existence of these is reflected in questions such as 'Does the research produce genuine or worthwhile knowledge?' and 'Towards which goals is research oriented?' Among other things, such uncertainties take both a *social* and a *political* form (K. R. Hammond, Harvey, & Hastie, 1992; Parker & Spears, 1997). Social uncertainty might be reflected in the question 'Is this form of knowledge acceptable in my research community?', while political uncertainty is addressed in the questions 'Which groups in society benefit from this research?' and 'What values does this research transmit?'.

Reflecting again on studies of workers' productivity, we can see that researchers might have uncertainty about whether the research agenda is appropriate and whether the approach taken on an issue ultimately serves particular political interests (Haslam, 2001; Pfeffer, 1997; Tannenbaum, 1966; Thompson & Warhurst, 1998). In this way the research might be criticized for addressing topics from a position which ultimately serves the narrow economic interests of management rather than

the interests of the broader community (e.g. by defining productivity and performance in financial rather than social terms; Pritchard, 1992).

It is clear, however, that questions of this form are asked less commonly during the research process than those that pertain to issues of methodological and statistical uncertainty. Moreover, answers are demanded even less frequently. Yet, given that the approach researchers take on *any* psychological issue will tend to reflect a particular ideological slant (however implicit), it seems appropriate to consider how the answers that their research provides are constrained by questions of this form that are *not* asked (Jevons, 1973). Certainly, psychologists have been reluctant to pursue these points, possibly mindful of the answers and of their inability to handle this particular form of uncertainty very effectively. Indeed, here uncertainty is far easier to create than to reduce.

Yet, lest this strategy be seen as wholly prudent, it is worth pointing out that a number of researchers interested in the social psychology of science have argued that progress in research is often achieved by *creating uncertainty* rather than by reducing it (e.g. Campbell, 1986; McGuire, 1997; Moscovici, 1976; Smithson, 1993). Uncertainty can be created when researchers are forced to confront new uncertainties using new methods (Kuhn, 1962). It can also be created when a researcher goes against the grain by questioning received wisdom and challenging knowledge that has been taken for granted within a particular worldview. This is what happened when Galileo observed that the earth went round the sun and when Freud demonstrated that the unconscious plays a central role in human behaviour. In many ways uncertainty reduction is therefore the business of 'normal science', while uncertainty creation is the stuff of scientific revolution. In effect, then, a focus on ontological uncertainty necessarily renders psychology a conservative science.

Uncertainty management in social psychology

The above arguments set out a novel and integrated framework for thinking about research in psychology. We also argue that they provide a set of interrelated terms that allow researchers to go beyond standard prescriptions for research practice and to think more broadly about their role as scientists. In relation to this point, the argument we want to explore in the remainder of this study is that while there are clear advantages associated with conducting normal social psychological science in the manner we have outlined, there are also significant dangers.

Reflecting first on some of the advantages of received ways of managing uncertainty in social psychology, we would suggest that the ability of experiments to reduce internal uncertainty has directly contributed to major advances in the discipline. One particularly clear illustration of this point is provided by reflecting on the manner in which the experimental method has allowed researchers to eliminate individual difference and personality-based confounds which represent potential threats to the interpretation of almost all the phenomena in which social psychologists are interested. For example, without the experimental method it would be impossible to offer compelling support for the view that prejudice and stereotyping are more than just the products of deviant personalities (e.g. as argued by Adorno, Frenkel-Brunswik, Levinson, & Sanford, 1950). This is because any purely correlational

relationship between a social contextual variable and a given behaviour can always be seen as personality-based—as was the case with the study of prejudice prior to pioneering experiments by Sherif and Tajfel. Tajfel (1977, 1981) argued that what is striking and interesting about intergroup behaviour is the consistency of behaviour observed across individuals who belong to the same group. However, while individual difference accounts struggle to explain this uniformity of behaviour, these could never be wholly discounted without recourse to experimental methodology which showed that people responded in more or less discriminatory ways when randomly assigned to different experimental conditions.

Importantly, though, as we have already intimated, a key feature of the way in which social psychology (and psychology as a whole) has developed over time is that it has come to place a particular premium on the *reduction* of ontological (i.e. methodological and statistical) uncertainty. It is this premium that has contributed in no small part to the dominance that the experimental method has assumed in our discipline. In itself this is no bad thing. It is not obvious that ambiguous or imprecise analysis is in any way preferable to research in which the causal link between variables is clearly specified and fully explored. It also seems desirable that researchers seek and apply relatively consensual criteria in order to determine what constitutes a tolerable amount of statistical uncertainty.

However, as we suggested above, there are clear dangers associated with a scientific strategy which elevates uncertainty reduction to the status of first principle. Three which are particularly significant are: (1) a tendency to see the reduction of methodological (and sometimes statistical) uncertainty as an end in itself; (2) the associated tendency to see the research process as purely a matter of dealing with ontological uncertainties; and (3) a tendency to neglect the contribution that uncertainty creation plays in the research process. We argue that these tendencies have become more pronounced as social psychology has progressed, and that while this has played a major role in normalizing social psychology as a science, it has also made that science increasingly more conservative, more safe and more dull.

There are at least two ways in which this can be illustrated. The first is by reflecting upon the likely treatment that so-called classic papers in social psychology would receive if they were submitted for peer review today. The second is to look at the sort of methodological prescriptions for psychological research that have gained in prominence as the discipline has developed.

In relation to the first of these points, it is interesting to speculate on the likely reaction of contemporary journal reviewers and editors to Triplett's (1898) social facilitation studies, Asch's (1956) conformity studies, Festinger and Carlsmith's (1959) cognitive dissonance studies, Milgram's (1963) obedience studies on Tajfel, Flament, Billig, and Bundy's (1971) minimal group studies. It is telling to reflect at the outset that most of this research could (and probably would) be rejected on grounds that it failed to provide acceptable estimates of statistical uncertainty. Most of these studies used only crude comparisons of means to make their point or simply examined the percentage of participants who made a particular response relative to some normative baseline. More damaging still for their prospects in the current publication climate, none of these studies provided any explanatory closure of the form associated with contemporary prescriptions for carrying out psychological

research. Indeed, the impact of these papers derives from the fact that they increased epistemological uncertainty rather than reduced ontological uncertainty. This was because they challenged received wisdom and raised new questions for researchers to answer. Of course, in every case a considerable amount of subsequent research was conducted with a view both (1) to providing answers that *did* reduce uncertainty, and (2) to ensuring that ontological uncertainty had been managed appropriately in the original studies (e.g. that the findings were reliable and valid). Nonetheless, the fact remains that in every case the studies themselves only made tentative steps in this direction—steps that in some instances proved to be completely wrong.

The fact that these studies are rightly perceived as classics is therefore at odds with received prescriptions for uncertainty reduction. Again, this is because their primary contribution was to create a particular form of uncertainty.² Indeed, we can see that by doing this, they contributed not to normal science but to scientific revolution.³ One of the key problems arising from pressures to reduce uncertainty is therefore that they may preclude progress of this nature. More provocatively, one might suggest that this is actually their latent, if not their manifest, function.

The same points can be reinforced if we turn to reflect upon the methodological fashions and strictures that have gained favour in experimental social psychology over time. In many instances, we see that these make it unlikely that researchers would ever even be tempted to conduct research of the form carried out by Triplett, Asch, Milgram and others (Fiske & Leyens, 1997). One very simple reason for this is that demands to reduce ethical uncertainties have also increased over this period. This is a topic that attracts heated debate, but nonetheless it is a clear example of the role that bigger-picture political uncertainties play in the research process.

One development that we want to consider in slightly more detail concerns experimental social psychologists' willingness to examine social interaction. Perhaps the first point to note here is that a willingness to broach issues of interaction is central to most definitions of social psychology (e.g. after Allport, 1935; see McGarty & Haslam, 1997). Indeed, we can see that it was a concern to look at the effects of social interaction that underpinned the research by Triplett that provided a starting point for this discussion.

If we think about a topic like stereotyping, for example, it is clear that stereotypes achieve their force precisely because they become widely shared through processes

² While arguing that uncertainty creation has a role to play in the research process, it is clearly not the case that all forms of uncertainty creation are equally productive. Generally speaking, frivolous or ill-informed questioning of empirical practice is less likely to contribute to theoretical progress than that which is based on a detailed examination of, and a clearly articulated intellectual and theoretical response to, the issues and research in question. Nonetheless, it is difficult (perhaps impossible) to specify *a priori* exactly which forms of uncertainty creation are most productive. Among other things, this is because there is a form of meta-uncertainty that arises from the fact that there is a political or social dimension to the question 'Which assumptions or values should be challenged?'. This is quite neatly demonstrated by Hopkins, Reicher, and Levine's (1997) argument that contemporary social cognitive theorizing about prejudice is itself prejudiced. Clearly, some researchers would argue that the political uncertainties created by this analysis constitute a worthwhile contribution to the field, but others would disagree. However, there is nothing in the *form* of Hopkins *et al.*'s contribution that allows one to deliver a verdict on its worth as that verdict must itself be a product of the *political process* through which science evolves.

³ None of these studies led to paradigm shifts of the sort discussed by Kuhn (1962) in his analysis of scientific revolution (e.g. as associated with the work of Galileo, Newton or Einstein). Nonetheless, the research was revolutionary in the more general sense in that it led to significant changes in research focus and thinking from which there was no turning back. It is in this sense that the term is used here.

of social interaction (Haslam, Turner, Oakes, McGarty, & Reynolds, 1998; Oakes, Haslam, & Turner, 1994; Stangor, 2000). As Tajfel (1981) noted, stereotyping would be worthy of only fleeting consideration if everybody held completely different stereotypes of other groups and their own. This is because under such circumstances stereotypes would bear no relationship either to issues of conflict and prejudice on the one hand or cooperation and empowerment on the other.

However, if we look at the history of empirical research into stereotyping from the beginning of the twentieth century (e.g. Katz & Braly, 1933) to the present day, it is abundantly clear that experimental social psychologists have moved away from studying issues of stereotype sharedness (see Condor, 1988; Haslam *et al.*, 1998; Reicher, Hopkins, & Condor, 1997). In part this is a reflection of important theoretical developments in the field, regarding the emergence of a cognitive approach to this and other related topics. However, hand-in-hand with these theoretical developments, methodological injunctions have served to preclude investigation of issues related to social interaction. In particular, we can see that methods like the Katz–Braly checklist, which still provide the best insights into shared stereotype content, are often shunned because participants' responses are less amenable to analysis using inferential statistics. Even more widespread is a reluctance to examine social interaction because behaviour that is shaped by that of another person violates statistical assumptions of independence. However, a concern to examine individuals in isolation means that measures of consensus—which should perhaps be of central interest to researchers (Haslam, 1997; Tajfel, 1981)—are often neglected or completely ruled out.

Here, then, we see that a concern to reduce methodological and statistical uncertainty has played a major role in restricting the scope of the work researchers undertake, and the theoretical possibilities they test between and explore when they do. This is a scenario that has been played out across experimental social psychology (and perhaps experimental psychology as a whole), and it is one reason why the ground on which researchers tread currently appears to be more tightly defined and constricted than it was, say 20 or 30 years ago. What has happened here, we suggest, is that the reduction of ontological uncertainty has been raised to the status of first principle, so that the means through which the science of social psychology has been normalized have become ends in themselves.

Many of the above points are quite controversial and, although they are a topic for informal discussion among researchers, they are rarely addressed in journals. However, we can examine the above propositions more formally by attempting to discern trends in the content of research publications over time. For this purpose we sought to examine the prominence given to issues of statistical and methodological uncertainty in empirical social psychology papers published over the last 40 years. To do this we analysed the content of the first 15 articles published in each of the calendar years 1968, 1978, 1988 and 1998 in the *Journal of Personality and Social Psychology* (the journal which is most clearly associated with the publication of experimental social psychology findings over that period of time) which contained the presentation of original experimental results. The mean number of experiments in each paper was taken as a measure of concern for methodological uncertainty. Within this sample, the amount of space in each single-experiment article given to the

reporting of results was taken as a measure of concern with statistical uncertainty.⁴ To examine whether researchers' willingness to study social interaction has changed over time, we also looked at the proportion of experiments in each of these four periods that involved participants interacting directly with one or more other people (either confederates or participants).

Clearly, our measures of methodological and statistical uncertainty do not represent 'pure' measure of these constructs. However, both serve this purpose reasonably well. If researchers are concerned to measure statistical uncertainty they should devote more energy to reporting the results of (ever-more complex) statistical procedures, and if they are concerned to minimize methodological uncertainty they should conduct more studies to eliminate confounds and establish generalizability.

The results from these analyses are presented in Figs 1–3. In each case, it is apparent that the emerging patterns are consistent with the arguments we have put forward. Specifically, the number of studies per publication has risen dramatically over time ($r(58) = \cdot63, p < \cdot001$; heteroscedasticity is a consideration here), as has the number of column centimetres given to the presentation of statistical analysis ($r(25) = \cdot43; p < \cdot03$).⁵

On the other hand, the proportion of experiments that examine social interaction has fallen over time ($\chi^2(3) = 13\cdot43, p < \cdot01$)—a decline that owes much to the total demise of confederate studies. In view of the fact that some of the most challenging and dramatic early experimental research in social psychology involved the use of confederates (e.g. Asch, 1956; Milgram, 1963), this demise could itself be seen as consistent with our general claim that experimental research has become less provocative and more hide-bound as it has become more concerned with issues pertaining to ontological uncertainty.

Again, though, we are certainly not arguing that attention to ontological uncertainty is bad in itself. However, what would appear to be problematic is the fact that advances in this sphere have been made at the expense of advances in the management of other forms of uncertainty. The steady decline of interactive research thus suggests that many researchers are now less concerned, or less able to deal with issues that ought to inform their science—especially when it is defined as the science of human interaction.

At least three processes may have been at work here. First, and most benignly, researchers' original interest in interaction and confederate studies may simply have been a passing fad or fashion (as suggested by Fiske & Leyens, 1997). Second, ethical or financial constraints may have made dramatic and time-consuming research less appealing. Third, an awareness of the power of interaction may have led

⁴ Because the number of single-experiment papers in the present sample declined dramatically over time—itsself a reflection of an increased concern with methodological uncertainty—this analysis took the form of a comparison between single-experiment studies reported in 1968 ($N = 13$) and those reported in the other three years (i.e. 1978, 1988, 1998) combined ($N = 14$).

⁵ The number of column centimetres was calculated for each study by taking the space devoted to (1) results sections (including heading); (2) complete paragraphs of sections of results and discussion sections that included any statistical copy, references to tables of results or presentations of effects; and (3) tables of results. Notes were not included in the total. Papers that involved some experimental manipulations but which did not have clearly labelled results or results and discussion sections were not included. Changes in page size and layout between 1978 and 1988 mean that the amount of statistical analysis may actually have increased rather more than these figures indicate.




Figure 1. Mean number of experiments per experimental study in articles sampled from the *Journal of Personality and Social Psychology*.




Figure 2. Mean number of columns cms. devoted to data analysis in single-experiment studies sampled from the *Journal of Personality and Social Psychology*.

researchers to control for its effects by designing it out of their research. Somewhat paradoxically, then, this final argument suggests that it is the very knowledge that interaction is important that has led to research outcomes which obstruct its investigation.

However, whatever the reasons for the patterns revealed in Figs 1–3, it would appear that increased attention to ontological uncertainty has deflected researchers' attention away from certain forms of epistemological uncertainty. Moreover, as a result of this, it may also be the case that the prospect of revolution (or simply excitement) at a theoretical and metatheoretical level has been diminished. One could also argue that this has made researchers less willing to raise or confront issues of political or social uncertainty (see Hopkins *et al.*, 1997). As noted above, these issues relate to the assumptions built into every stage of the research process, from the




Figure 3. Proportion of experiments involving interaction in articles sampled from the *Journal of Personality and Social Psychology*.

decision about which question to address, to the choice of dependent measure. In the case of Triplet's research, for example, why is it that we regard the performance of the solo reeler as the baseline rather than that of the reeler in a group? Why is it that we define performance in terms of the individual reelers' speed rather than the comfort and cohesiveness of the group as a whole? Why are we interested in this form of competitive performance at all?

The uncertainty management model suggests that these big-picture uncertainties are every bit as important as the methodological and statistical uncertainties that are becoming the ever-tightening focus of standard treatments of the research process. Moreover, all researchers have to deal with uncertainties of this form. However, the above patterns suggest that there may be a growing tendency to do this through strategies of avoidance or displacement.

There are several areas in social psychology where this strategy palpably fails. The most obvious of these perhaps relates to the study of uncertainty itself and the concomitants of uncertainty such as truth, ignorance and accuracy. This is because these constructs are themselves part of the subject-matter of social psychology. Such problems can again be illustrated with reference to research into stereotyping—this time considering debate about stereotype accuracy. Here a range of views has been put forward in answer to the question of whether social stereotypes are accurate (see Oakes *et al.*, 1994; Oakes & Reynolds, 1997, for summaries).

An early view that has remained popular in social psychology was that stereotypes are necessarily distorted representations of reality (e.g. Klineberg, 1950; Zawadzki, 1942). An alternative view was that stereotypes could be wrong but that they contained a kernel of truth (Schuman, 1966). In trying to make sense of these competing views, Judd and Park (1993) have suggested that advances in psychological measurement mean that the prospect exists to establish the accuracy of stereotypes by comparing specific stereotypes with objective measurements (for earlier statements to this effect, see Fishman, 1956; Klineberg, 1951; Schuman, 1966).

Thus, just as it is possible to ask a group of Britons to estimate the height of Germans and determine the accuracy of this stereotypic estimate by comparing it to the actual height of Germans, so it would be possible to establish the accuracy of British estimates of the efficiency of Germans by comparing estimates with the objective level of the psychological construct of efficiency in the German population. Efficiency, like any other psychological construct, would be measured by standardized psychological tests that enable comparisons between populations.

We will not deal here with our own solutions to the particular scientific problems raised by this sort of example (see McGarty, 1999; Oakes *et al.*, 1994) but instead focus on the uncertainties that must be addressed by researchers working in this field. First, there are descriptive uncertainties surrounding problems of measurement that apply to any research domain. In this regard Judd and Park (1993) express optimism but not complete confidence that psychological science will be up to the task of establishing stereotype accuracy. On the other hand, Oakes and Reynolds (1997) state expressly that the methods of psychological science can never achieve this goal. One reason for this is that 'efficiency' is a dynamic behavioural construct so that a person's efficiency itself varies with the context in which they behave. If, for example, efficiency is a norm for Germans, then Germans may become more efficient at those particular times when a social self-categorization as German becomes salient (Turner, 1985). To the extent that it exists, it seems inappropriate to reduce *meaningful* variation of this form into a simple statement of descriptive uncertainty.

In this example, though, perhaps the weightiest uncertainties relate to the political implications of lending scientific credibility to the beliefs that some groups hold about other groups (although, of course, behavioural sciences such as anthropology, which contribute enormously to the received ethnographic wisdom about groups, do this almost as a matter of course; see Campbell, 1967). To illustrate this point, imagine that researchers conclude that Group A's negative stereotype about Group B is accurate but that Group B's positive stereotype of Group A is not (e.g. Schuman, 1966; Triandis & Vassiliou, 1967). Will Group A's factually correct view be used to justify negative treatment of members of Group B (in the way that the scientific knowledge of the 1930s was used to justify the extermination of Jews in Nazi Germany)? Should we conclude that the members of Group A are more objective than those in Group B? Should we seek to disabuse Group B of their rosy view? Should we consider other forms of social intervention, to ensure that the truth we have identified is internalised by the subjects of our enquiry?

Social psychologists routinely offer prescriptions of this form (which is one reason why there is a large body of research on the topic of stereotype change; e.g. Hewstone, 1989; Hewstone, Macrae, Griffiths, & Milne, 1994). However, it is less clear that they routinely reflect on the political implications of such prescriptions (see Hopkins *et al.*, 1997; Oakes, Haslam, & Reynolds, 1999). These large scale uncertainties are not associated with the experimental method *per se*, but it is clear that there is considerable danger if researchers think that they have been dealt with merely because they have conducted research which has addressed and reduced *other* uncertainties. Our message is simple: no amount of attention to ontological uncertainties (i.e. issues of methodology and statistics) will ever militate against the need to address epistemological uncertainties. As intimated by Billig (1996, p. 37),

the gods of knowledge will never be appeased simply by offerings of impeccably harvested experimental results.

Conclusion

Having made these various points, two questions remain to be addressed by way of conclusion. First, we need to ask whether the experimental method is such a threat to the research process that we need to consider abandoning it all together. Second, we need to reflect more broadly on the strategy that is likely to redress some of the problems that have been identified here.

In relation to the first point, it is clearly tempting to suggest that social psychologists simply need to forsake experimental methodology because this imposes an unnecessary and restrictive straightjacket on research progress. It is clear, too, that this is a view that has been endorsed by a number of researchers in the field, most recently by discourse analysts (e.g. Billig, 1997). This view is based on an argument that the experimental method is directly responsible for the atrophy of scientific creativity and the narrowing of political perspective.

We think this strategy and argument are wrong. In the first instance this is because the problems discussed above arise from pressures to normalize social psychology that exist in every science. In social psychology this has taken a particular form—and is realized in a concern to minimize methodological uncertainty and to accurately measure (and sometimes to minimize) statistical uncertainty. The experimental method has been important in our discipline because it is capable of achieving these objectives particularly well. But other methodologies can also be marshalled to achieve similar ends, and they too can become slaves to uncertainty reduction in much the same way. Experimentalists do not have a monopoly on the discovery of what Walker (1997) refers to as ‘impeccable trivia’.

To blame the experimental method for the woes of social psychology is therefore like blaming horses for the Charge of the Light Brigade. If researchers’ primary concern is only to reduce methodological uncertainty, then the price they will pay is to have a social psychology of trivia. And this will be true whatever method they use, though it may be true that they will be more attracted to experimental methodology—for the same reasons that members of the Light Brigade preferred horses to donkeys.

So, if the method of social psychology is not the problem, what is, and how can it be dealt with? The problem, we suggest, is one associated with all sciences where pressures towards normalization start to exceed the capacity for the data and phenomena undergoing investigation to be normalized. In social psychology, as in society at large, there is strong pressure for us all to agree about the interpretation of the world around us, and people often argue that such agreement must emerge if we eliminate the methodological noise or uncertainty that contributes to our disagreement. The review process, with its pressures towards consensualization and universal acceptance, is one example of this belief at work.

However, sometimes we have to step back from this normalization process and recognize that no amount of methodological tinkering will explain away our disagreement. Sometimes, the uncertainties we face are not methodological, they are

metatheoretical and political (see Fiske & Leyens, 1997; Oakes *et al.*, 1994, Chap. 8). While a variety of ideological and other pressures lead us to skirt around this fact, a concern for scientific progress dictates that we recognize and embrace it.

Social psychology, like all science, is about argument as well as acceptance, conflict as well as consensus, revolution as well as stability. For this reason, the way forward lies not in abandoning experimental research or dismissing it as inherently worthless. But just as importantly, it is not about trying to conceal real differences within a morass of methodological minutiae. It is about taking our differences at face value and being prepared to confront some of the bigger metatheoretical and political uncertainties that underlie them.

Importantly, as many of the pioneers of experimental social psychology proved, in this way the experimental method need not be a barrier to scientific progress. For its potential to be realized, though, we have to recognize that we can, and must, use experiments for more than just reducing methodological uncertainty. Moreover, by taking this step, we need not resign ourselves to the stagnant, circumscribed and pointless science that the critics of experimental social psychology lament (e.g. Billig, 1997; Gergen, 1997; Harré, 1997). Neither do we have to resign ourselves to the fact that the most interesting and important issues in social psychology will be researched by people who are not social psychologists.

As Tajfel (1972) argued, when the so-called crisis in social psychology was reaching its peak, it is not the case that experiments *per se* are bad or good, trivial or non-trivial. As with any other method, the quality and utility of experiments is always constrained by the quality and utility of the ideas and theories they test. A trivial idea makes for a trivial experiment. But experiments can test powerful ideas and they can also be tools of revolutionary science. The more willing we are to put them to this use the more likely it is that we will rediscover and share in a science that is dynamic, challenging and relevant. It would be interesting too.

Acknowledgements

The authors would like to thank Michael Cook, Mike Smithson, Kate Reynolds, John Turner and three anonymous reviewers for their detailed comments on previous versions of this article. Parts of this study are based on a paper presented at the 4th meeting of the Society of Australasian Social Psychologists, Christchurch, 16–19 April 1998 in a symposium convened by J. Michael Innes entitled ‘100 years after Triplett: The development of social psychology in its various forms’.

References

- Adorno, T. W., Frenkel-Brunswick, E., Levinson, D. J., & Sanford, R. N. (1950). *The authoritarian personality*. New York: Harper.
- Allport, G. W. (1935). Attitudes. In C. M. Murchison (Ed.), *Handbook of social psychology* (pp. 789–844). Worcester, MA: Houghton-Mifflin.
- Asch, S. E. (1956). Studies of independence and conformity: A minority of one against a unanimous majority. *Psychological Monographs: General and Applied*, 70, 1–70.
- Billig, M. (1996). *Arguing and thinking* (2nd ed.). Cambridge: Cambridge University Press.
- Billig, M. (1997). Discursive, rhetorical and ideological messages. In C. McGarty & S. A. Haslam (Eds.), *The message of social psychology: Perspectives on mind in society* (pp. 36–53). Oxford: Blackwell.
- Binet, A., & Henri, V. (1894). De la suggestibilité naturelle chez les enfants. *Revue Philosophique*, 38, 337–347.

- Campbell, D. T. (1967). Stereotypes and the perception of group differences. *American Psychologist*, 22, 817–829.
- Campbell, D. T. (1986). Science's social system of validity-enhancing collective belief change and problems of the social sciences. In D. W. Fiske & R. A. Schweder (Eds.), *Metatheory in social science: Pluralisms and subjectivities* (pp. 108–135). Chicago: University of Chicago Press.
- Campbell, D. T., & Stanley, J. C. (1963). *Experimental and quasi-experimental designs for research*. Chicago: Rand McNally.
- Cohen, J. (1977). *Statistical power analysis for the behavioural sciences*. New York: John Wiley.
- Cohen, J. (1994). The earth is round ($p < .05$). *American Psychologist*, 49, 997–1003.
- Cohen, J. (1995). The earth is round ($p < .05$): Rejoinder. *American Psychologist*, 50, 1103.
- Condor, S. (1988). Race stereotypes and racist discourse. *Text*, 8, 69–89.
- Condor, S. (1990). Social stereotypes and social identity. In D. Abrams & M. Hogg (Eds.), *Social identity theory: Constructive and critical advances*. Hemel Hempstead: Harvester Wheatsheaf.
- Festinger, L., & Carlsmith, J. M. (1959). Cognitive consequences of forced compliance. *Journal of Abnormal and Social Psychology*, 58, 203–210.
- Fishman, J. A. (1956). An examination of the process and function of social stereotyping. *Journal of Social Psychology*, 43, 27–64.
- Fiske, S. T., & Leyens, J.-P. (1997). Let social psychology be faddish or, at least, heterogeneous. In C. McGarty & S. A. Haslam (Eds.), *The message of social psychology: Perspectives on mind in society* (pp. 92–112). Oxford: Blackwell.
- Gergen, K. J. (1997). Social psychology as social construction: The emerging vision. In C. McGarty & S. A. Haslam (Eds.), *The message of social psychology: Perspectives on mind in society* (pp. 113–128). Oxford: Blackwell.
- Haines, H., & Vaughan, G. M. (1979). Was 1898 a 'great date' in the history of experimental social psychology? *Journal of the History of the Behavioural Sciences*, 15, 323–332.
- Hammond, G. (1996). The objections to null hypothesis testing as a means of analysing psychological data. *Australian Journal of Psychology*, 46, 104–106.
- Hammond, K. R., Harvey, L. O., & Hastie, R. (1992). Making better use of scientific knowledge: Separating truth from justice. *Psychological Science*, 3, 80–87.
- Harkins, S. G. (1987). Social loafing and social facilitation. *Journal of Experimental Social Psychology*, 23, 1–18.
- Harré, R. (1997). Social life as rule-governed patterns of joint action. In C. McGarty & S. A. Haslam (Eds.), *The message of social psychology: Perspectives on mind in society* (pp. 129–145). Oxford: Blackwell.
- Haslam, S. A. (1997). Stereotyping and social influence: Foundations of stereotype consensus. In R. Spears, P. J. Oakes, N. Ellemers, & S. A. Haslam (Eds.), *The social psychology of stereotyping and group life* (pp. 119–143). Oxford: Blackwell.
- Haslam, S. A. (2001). *Psychology in organizations: The social identity approach*. London: Sage.
- Haslam, S. A., & McGarty, C. (1998). *Doing psychology: An introduction to research methodology and statistics*. London: Sage.
- Haslam, S. A., Turner, J. C., Oakes, P. J., McGarty, C., & Reynolds, K. J. (1998). The group as a basis for emergent stereotype consensus. *European Review of Social Psychology*, 9, 203–239.
- Hewstone, M. (1989). Changing stereotypes with disconfirming evidence. In D. Bar-Tal, C. F. Graumann, A. W. Kruglanski, & W. Stroebe (Eds.), *Stereotyping and prejudice: Changing conceptions* (pp. 207–223). New York/London: Springer-Verlag.
- Hewstone, M., Macrae, C. N., Griffiths, R., & Milne, A. B. (1994). Cognitive models of stereotype change: 5. Measurement, development, and consequences of subtyping. *Journal of Experimental Social Psychology*, 30, 505–526.
- Hopkins, N., Reicher, S., & Levine, M. (1997). On the parallels between social cognition and the 'new racism'. *British Journal of Social Psychology*, 36, 305–330.
- Jevons, F. R. (1973). *Science observed: Science as a social and intellectual activity*. London: Allen & Unwin.
- Judd, C. M., McClelland, G. H., & Culhane, S. E. (1995). Data analysis: Continuing issues in the everyday analysis of psychological data. *Annual Review of Psychology*, 46, 433–465.
- Judd, C. M., & Park, B. (1993). Definition and assessment of accuracy in social stereotypes. *Psychological Review*, 100, 109–128.

- Karau, S. J., & Williams, K. D. (1993). Social loafing: A meta-analytic review and theoretical integration. *Journal of Personality and Social Psychology*, *65*, 681–706.
- Katz, D., & Braly, K. (1933). Racial stereotypes of one hundred college students. *Journal of Abnormal and Social Psychology*, *30*, 175–193.
- Klineberg, O. (1950). The scientific study of national stereotypes. *International Social Science Bulletin*, *3*, 505–515.
- Kuhn, T. (1962). *The structure of scientific revolutions*. Chicago: University of Chicago Press.
- Mayo, E. (1933). *The human problems of industrial civilisation*. New York: Macmillan.
- McGarty, C. (1999). *Categorization in social psychology*. London: Sage.
- McGarty, C., & Haslam, S. A. (1997). Introduction and a short history of social psychology. In C. McGarty & S. A. Haslam (Eds.), *The message of social psychology: Perspectives on mind in society* (pp. 1–19). Oxford: Blackwell.
- McGuire, W. J. (1997). Going beyond the banalities of bubbapsychology: A perspectivist social psychology. In C. McGarty & S. A. Haslam (Eds.) *The message of social psychology: Perspectives on mind in society* (pp. 221–237). Oxford: Blackwell.
- Milgram, S. (1963). Behavioral study of obedience. *Journal of Abnormal and Social Psychology*, *67*, 371–378.
- Moscovici, S. (1976). *Social influence and social change*. London: Academic Press.
- Oakes, P. J., Haslam, S. A., & Reynolds, K. J. (1999). Social categorization and social context: Is stereotype change a matter of information or meaning? In D. Abrams & M. Hogg (Eds.), *Social identity and social cognition* (pp. 55–79). Oxford: Blackwell.
- Oakes, P. J., Haslam, S. A., & Turner, J. C. (1994). *Stereotyping and social reality*. Oxford: Blackwell.
- Oakes, P. J., & Reynolds, K. J. (1997). Asking the accuracy question: Is measurement the answer? In R. Spears, P. J. Oakes, N. Ellemers, & S. A. Haslam (Eds.), *The social psychology of stereotyping and group life* (pp. 51–71). Oxford: Blackwell.
- Parker, I., & Spears, R. (Eds.) (1997) *Psychology and society*. London: Pluto Press.
- Pfeffer, J. (1997). *New directions for organizational theory: Problems and prospects*. New York/Oxford: Oxford University Press.
- Pritchard, R. D. (1992). Organizational productivity. In M. D. Dunnette & L. M. Hough (Eds.) *Handbook of industrial and organizational psychology* (Vol. 3, pp. 443–471). Palo Alto, CA: Consulting Psychologists Press.
- Reicher, S. D., Hopkins, N., & Confor, S. (1997). Stereotype construction as a strategy of influence. In R. Spears, P. J. Oakes, N. Ellemers, & S. A. Haslam (Eds.), *The social psychology of stereotyping and group life* (pp. 94–118). Oxford: Blackwell.
- Roethlisberger, F. J., & Dickson, W. J. (1939). *Management and the worker: An account of a research program conducted by the Western Electric Company, Hawthorne Works, Chicago*. Cambridge, MA: Harvard University Press.
- Schuman, H. (1966). Social change and the validity of regional stereotypes in East Pakistan. *Sociometry*, *29*, 428–440.
- Smithson, M. (1993). Ignorance and science: Dilemmas, perspectives and prospects. *Knowledge: Creation, diffusion, utilization*, *15*, 133–156.
- Smithson, M. (1999). *Statistics with confidence*. London: Sage.
- Stangor, C. (2000). Volume overview. In C. Stangor (Ed.), *Stereotypes and prejudice: Key readings in social psychology* (pp. 1–19). Boston: Psychology Press.
- Svyantek, D. J., & Ekeberg, S. E. (1995). The earth is round (so we can probably get there from here). *American Psychologist*, *50*, 1101.
- Tajfel, H. (1972). Experiments in a vacuum. In J. Israel & H. Tajfel (Eds.), *The context of social psychology: A critical assessment* (pp. 69–119). London: Academic Press.
- Tajfel, H. (1977). Social psychology and social reality. *New Society*, *39*, 65–66.
- Tajfel, H. (1981). Social stereotypes and social groups. In J. C. Turner & H. Giles (Eds.), *Intergroup behaviour* (pp. 144–167). Oxford: Blackwell.
- Tajfel, H., Flament, C., Billig, M. G., & Bundy, R. F. (1971). Social categorization and intergroup behavior. *European Journal of Social Psychology*, *1*, 149–177.
- Tajfel, H., & Fraser, C. (1978). Social psychology as social science. In H. Tajfel & C. Fraser (Eds.), *Introducing social psychology: An analysis of individual reaction and response* (pp. 21–53). Harmondsworth: Penguin.

- Tannenbaum, A. S. (1966). *The social psychology of the work organization*. London: Tavistock.
- Thompson, P., & Warhurst, C. (Eds.) (1998). *Workplaces of the future*. Houndmills: Macmillan.
- Triandis, H. C., & Vassiliou, V. (1967). Frequency of contact and stereotyping. *Journal of Personality and Social Psychology*, 7, 316–328.
- Triplett, N. (1898). The dynamogenic factors in pacemaking and competition. *American Journal of Psychology*, 9, 507–533.
- Tukey, J. W. (1991). The philosophy of multiple comparisons. *Statistical Science*, 6, 100–116.
- Turner, J. C. (1981). Some considerations in generalizing experimental social psychology. In G. M. Stephenson & J. H. Davis (Eds.), *Progress in applied social psychology* (Vol. 1, pp. 3–34). Chichester: Wiley.
- Turner, J. C. (1985). Social categorization and the self-concept: A social cognitive theory of group behaviour. In E. J. Lawler (Ed.), *Advances in group processes* (Vol. 2, pp. 77–122). Greenwich, CT: JAI Press.
- Walker, I. (1997). *The long past, short history, and uncertain future of social psychology*. Unpublished manuscript, Murdoch University.
- Zawadzki, B. (1942). Limitations of the scapegoat theory of prejudice. *Journal of Abnormal and Social Psychology*, 43, 127–141.

Received 17 December 1998; revised version received 20 August 1999