British Journal of Psychology (1996), **87**, 431–445. *Printed in Great Britain*

Overconfidence, base rates and outcome positivity/negativity of predicted events

Briony D. Pulford and **Andrew M. Colman***
Department of Psychology, University of Leicester, Leicester LE1 7RH, UK

Overconfidence is said to occur when a person's confidence in a series of predictions exceeds the level of accuracy achieved. In this experiment, questionnaire items requiring predictions of personal life-events were selected according to their objective base rates of occurrence and their outcome positivity/negativity for participants in a pilot study. The 98 participants in the main experiment predicted whether they would experience each event within the next week and rated their confidence in their predictions. Predictions were compared with responses to a follow-up questionnaire a week later to determine accuracy. Significant overconfidence occurred, but it was greater for positive-outcome than negative-outcome items, and the results revealed a curvilinear relationship between base rates and overconfidence, with maximum overconfidence at intermediate base rate levels and underconfidence at both extremes of the base rate range. Subjects tended to overestimate base rates below 40 per cent and to underestimate higher base rates.

Recent evidence has shown that people tend to be overconfident in predicting both the behaviour of other people (Dunning, Griffin, Milojkovic & Ross, 1990) and their own behaviour (Vallone, Griffin, Lin & Ross, 1990). Across a wide range of social and self-predictions, subjective confidence has consistently exceeded the objective accuracy of predictions, especially when confidence was relatively high. This miscalibration of confidence and accuracy in judgements was first reported by Adams & Adams (1961) and has more recently been called the *overconfidence effect*. The findings associated with the effect suggest that the implicit psychological theories and other forms of intuitive psychology on which people base their social and self-judgements (Nisbett & Ross, 1980; Ross, 1977) may have less predictive validity than they generally assume, and that people's confidence estimates are especially likely to err when they are highly confident.

For any action that might be performed, or any other event that might be experienced by an individual, there is an objective base rate of occurrence defined by the relative frequency with which members of the relevant population experience the event in a specified time period. Dunning *et al.* (1990) found that, especially when individuals predicted that a person would behave atypically, these predictions tended to be inaccurate and also overconfident. By failing to give due weight to objective base rates and instead grounding their predictions on supposed dispositional factors internal to the actors, the participants made errors in predicting future behaviour and

* Requests for reprints.

displayed overconfidence in their predictions. If this interpretation is correct, then these errors in predicting *future* behaviour bear a family resemblance to the *overattribution bias* first reported by Jones & Harris (1967) and subsequently confirmed by a number of researchers (e.g. Gilbert & Jones, 1986; Johnson, Jemmott & Pettigrew, 1984; Ross, Amabile & Steinmetz, 1977; Webster, 1993) according to which people tend to overestimate the causal role of dispositional relative to situational factors in explaining an actor's *past* behaviour.

In the case of self-predictions, the individual has maximum knowledge of the target, yet prediction errors still occur, partly as a result of faulty inferences or construals of possible future events. Vallone *et al.* (1990) proposed that the overconfidence effect results from 'insufficient allowance for the possibility of this erroneous inference or misconstrual' (p. 583). Predictions of self-behaviour that go against the norm or base rate still tend to lead to low accuracy and overconfidence, even though dispositional knowledge about oneself is more accurate than dispositional knowledge about other people.

By definition, overconfidence arises either when confidence is unrealistically high for a given level of accuracy or, what amounts to virtually the same thing, when accuracy is unrealistically low for a given level of confidence. Can the phenomenon be explained entirely by a tendency for participants to make insufficient use of base rate information? If indiscriminate neglect of base rate information were the sole explanation, then what would be the relation between base rates and overconfidence? If individuals take full account of base rates, then at intermediate base rate levels, where prediction is relatively difficult because the prior probabilities of occurrence and non-occurrence of events are roughly equal, both accuracy and confidence should be relatively low. In regions of extremely low or extremely high base rates, where prediction is relatively easy because the prior probabilities of the event are close to either zero or unity, both accuracy and confidence should be relatively high if participants make full use of the base rate information. Thus if participants take full account of base rates, both their accuracy and their confidence levels should vary in parallel across the base rate range, and therefore no systematic relationship should exist between base rates and overconfidence. If, on the other hand, they neglect base rate information entirely, then neither accuracy nor confidence should be higher in the extreme than the intermediate base rate regions, and once again there should be no systematic covariation of base rates and overconfidence.

There is a third–more plausible—possibility, namely that base rates influence accuracy and confidence differentially, in which case base rates and overconfidence may turn out to be functionally related. The most natural assumptions would be (*a*) that people generally make at least some use of base rate information in formulating predictions, so that their predictions of events with extreme base rates tend to be relatively accurate compared to their predictions of events with intermediate base rates; and (*b*) that people generally take less account of base rates in modulating their feelings of confidence in their predictions. Although these assumptions have not been specifically tested, both are in line with closely related evidence in the field of decision making (e.g. Ajzen, 1977; Bar-Hillel, 1980; Tversky & Kahneman, 1980). If both assumptions are correct and applicable to situations in which the overconfidence effect occurs, then the effect should be greatest for events whose base rates of occurrence are in the intermediate range, and it should be relatively small for events with extreme base rates.

The overconfidence effect may thus be attributable at least partly to people's failure to make adequate use of base rate information in modulating the confidence of their predictions. Another judgemental bias that has been reported is a tendency for likelihood estimates of future events to be unrealistically optimistic, and there is evidence that this bias varies according to whether the prediction relates to positive or negative events (Weinstein, 1980). The principal aims of the experiment reported below are to examine the base rate–overconfidence relationship and to establish whether unrealistic optimism is reflected in confidence judgements as well as predictions. A subsidiary aim is to clarify the underlying mechanism of the overconfidence effect by determining whether overconfident subjects are able to estimate base rates accurately but are unable to make adequate use of them, or whether they are simply unable to estimate the base rates accurately in the first place. In the light of the well-established Pollyanna effect (Boucher & Osgood, 1969), according to which people tend to exaggerate the positive and minimize the negative aspects of their experience, a third aim of the experiment is to test the hypothesis that overconfidence is greater for events with positive, desirable outcomes than for negative-outcome events.

Method

Participants

Ninety-eight undergraduate students (29 males and 69 females) with a mean age of 20.65 years (range 18 43 years) participated in the experiment. All were first-year psychology undergraduates, who completed the experiment in return for course credits. They were unaware of the aims of the experiment and had not been exposed to any information about the overconfidence effect or issues of judgement calibration. Due to the week-long interval between the two testing sessions of the experiment, attrition of one participant occurred.

Design

A 2 × 12 repeated measures factorial design (positivity/negativity × base rate) was used. The dependent variables, both measured on percentage scales, were the accuracy and rated confidence of the predictions, and from these data levels of overconfidence were calculated as explained below. Participants' estimates of the base rates were also measured on a percentage scale.

Materials

A pilot study was carried out to select the questionnaire items to be used in the experiment. An initial pool of 170 candidate items, which included academic, social and personal events, was assembled largely from the records of three volunteers who kept diaries of all noteworthy events that occurred in their lives within a week, and additional candidate items were added intuitively by the experimenters with the aim of including relatively infrequent events that had not turned up in the diaries.

Two categories of events were constructed: those with outcomes that were assumed to be positive and those with outcomes that were assumed to be negative. The items were all rated for outcome positivity/negativity by the participants in the pilot study to establish empirically their true perceived positivity or negativity within the population being studied. Within these two categories (positive-outcome and negative-outcome), items were included with the aim of spanning the entire range of base rates, from low base rate events assumed to happen only very rarely within the target population of university students (e.g. 'break a bone in your body') to high base rate events that were assumed to be likely to happen to most students within a week (e.g. 'wash your hair'). Other items included 'fail to study sufficiently', 'smoke a cigarette', 'visit a doctor', 'have problems using a computer', 'have an argument', etc.

Fifty-nine undergraduate students (seven males and 52 females) with a mean age of 19.74 years participated in the pilot study. Full instructions were given at the top of the pilot questionnaire, and participants were instructed to ask an experimenter for help if they had any difficulty understanding them. These pilot study participants were asked to state, by circling either 'Yes' or 'No', whether each of the 170 events had occurred in their lives within the last week (e.g. 'In the last seven days did you use a computer?'). They were also asked to rate the outcome of each event as either positive (+), negative (−) or neutral (0). Positive-outcome events were described as events that they would like to experience and negative-outcome events were ones that participants would not like to experience; neutral events were ones about which participants felt neither positive nor negative.

Participants' responses to each item were scored 1 for a positive answer, 0 for a neutral answer and −1 for a negative answer. The total outcome positivity or negativity of each item was calculated by summing the participants' individual scores. The resulting outcome positivity/negativity scores ranged from −42 to +43, compared with a theoretically possible range (if all had responded identically to the most extreme items) of −59 to +59. Twelve of the items, classified as neutral (scoring from −10 to +10), were omitted from the main study because they were ambiguous inasmuch as there was insufficient consensus among students as to whether they were positive or negative. The base rate for each event was defined as the percentage of students in the pilot study who experienced the event in the critical week. The resulting list contained 158 items, which were rank-ordered according to their base rates.

The base rate scale was divided into 10 ranges representing 10 per cent intervals, and subdivided further to incorporate two extra ranges of 3 per cent at both extremes of the base rate range, because these were areas of special interest where fine detail was required. The final set of base rate ranges are shown in Table 1.

Table 1. Ranges and their corresponding base rates

Range	Base rate (%)	Range	Base rate (%)
1	0–3	7	50–60
2	3–10	8	60–70
3	10–20	9	70–80
4	20–30	10	80–90
5	30–40	11	90–97
6	40–50	12	97–100

Four positive-outcome and four negative-outcome items were selected out of the initial pool for each range, but only three negative-outcome items were available for range 10 and none for ranges 11 and 12. This was because, fortunately, there appeared to be few negative-outcome events that were experienced by a high proportion of the population within a given week. Thus a smaller pool of 87 items was finally selected which covered most of the base rate ranges and included 48 positive-outcome items and 39 negative-outcome items, and these items formed the main study questionnaire. The questionnaire was used twice in order to elicit predictions and then to determine actual outcomes: the items were presented in the same order, but instructions to participants varied, with different responses required according to which questionnaire was being completed.

Procedure

The questionnaire was completed by 98 participants who had not taken part in the pilot study. Participants were asked to make self-predictions by circling either 'Yes' or 'No' on the questionnaire according to whether they thought that they would experience each of the events within the next week. They were also asked to write down a number, using a half-range percentage scale (based on the scales used by Dunning *et al.*, 1990 and Vallone *et al.*, 1990) to indicate their degree of confidence that their judgements were correct. The anchors of the scale were 50 per cent (representing complete lack of confidence) at one extreme and 100 per cent (representing total confidence) at the other. The participants were instructed as follows:

> Your confidence estimate should correspond to the percentage of time you would expect to be right on judgements with that level of confidence. For example 90 per cent means that you are very sure that you are right in your prediction and that you'd expect to be right 9 times out of 10, or the probability that your answer is correct is .9. The 50 per cent level shows complete uncertainty and you would expect to be right about half of the time, due to chance. (Confidence estimates below 50 per cent are inappropriate as 40 per cent confidence in predicting 'Yes' is equivalent to saying 60 per cent confidence in 'No'.)

Participants were retested one week later when they were given another questionnaire, which included the same items as the first questionnaire, with instructions to state whether the event had occurred within the last week since filling in the first questionnaire. Participants were also asked to rate each item as positive, negative or neutral in outcome and to estimate how many other first-year students they thought would have experienced the event within the previous week.

Participants' self-predictions on the first questionnaire were then compared with their own answers on the second questionnaire to determine their accuracy for each item. A prediction that an event either would or would not occur within the specified time was defined as accurate if it was fulfilled in reality. The mean percentage confidence and accuracy were calculated for each item across subjects using the following equation:

$$\text{Overconfidence} = \text{Confidence} - \text{Accuracy}$$

The base rates were recalibrated using the follow-up data from the main study data alone, so that they would refer precisely to the week in which the participants were tested and about which they made their predictions. This resulted in many of the items shifting in base rate, and some moving across boundaries into different base rate ranges. One reason for this was that circumstances had changed slightly between the pilot study and the main study. In the pilot study, which was conducted near the beginning of the first semester, for example, only 32.20 per cent of students said that they would skip more than one lecture in the specified week, but by the time of the main study several weeks later, 71.13 per cent said they would do so, presumably because of changing attitudes towards lecture attendance and increased coursework commitments. The base rates were pertinent to the specific time span to which the judgements referred and to the specific target population under investigation. In general, base rates may apply to the future as a whole, but in this experiment they were limited to a specific week, after which the accuracy of the predictions for that week could be determined empirically.

Product moment correlation coefficients and t tests were calculated between participants' accuracy and confidence ratings across items and then between their estimates of item base rates and objective base rates. Items were then classified as positive or negative according to the outcomes of the specified events and were grouped into base rate ranges according to their objective frequencies of occurrence during the week in question. Differences between positive-outcome and negative-outcome items and differences across base rate ranges in confidence, accuracy, overconfidence and estimated base rates were then tested for significance with 2×12 analysis of variance.

Results

Across all items combined, subjects were accurate 82.10 per cent of the time, and their mean confidence was 84.42 per cent. These figures yielded a significant level of overconfidence of 2.32 per cent ($t(86) = 3.01$, $p < .001$). A curvilinear, inverted U-shaped relationship was found between base rates and overconfidence, as shown in Fig. 1.

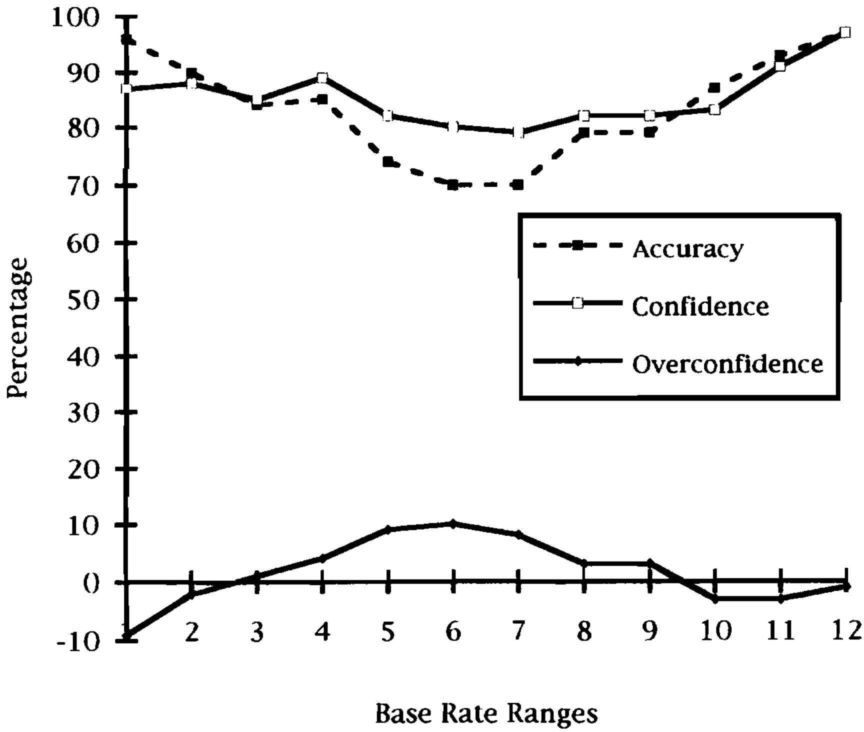

Figure 1. Mean accuracy, confidence and overconfidence across base rate ranges from low (range 1 = 0–3 per cent), through moderately common (range 6 = 40–50 per cent), to extremely common (range 12 = 97–100 per cent).

The effect of base rates on overconfidence was statistically significant ($F(11,65) = 8.58$, $p < .001$) and an *a posteriori* Tukey HSD test revealed that the extreme base rate ranges 1, 2, 10 and 11 were significantly lower in overconfidence than the central ranges 5, 6 and 7. Range 1 was also significantly lower than ranges 4 and 9 and range 3 was significantly lower than range 6. Range 12 did not differ significantly in overconfidence from the intermediate ranges, but this was probably due to the very small sample size of only two items within range 12. These results show that overconfidence in the extreme base rate ranges was significantly lower than in the middle ranges. In fact, at the extremes, *underconfidence* occurred, with accuracy higher than confidence, whereas at intermediate levels of base rates *overconfidence* occurred, peaking to a maximum of 10.15 per cent in the 40–50 per cent base rate range (range 6).

Figure 1 also shows the variations in accuracy and confidence across the base rate range, and these data disclose a possible reason for the inverted U-shaped relationship between base rates and overconfidence. The level of accuracy clearly followed a U-shaped curve across the base rate ranges, whereas participants' confidence in their judgements followed a much shallower U-shape, dropping much less in the intermediate base rate ranges. As a consequence, high levels of overconfidence occurred in the intermediate base rate ranges where confidence did not drop as low as accuracy. The data shown in Fig. 1 may appear to suggest that participants were reluctant to use the extremes of the scale when reporting their confidence. This does not appear to be the case, however, because of all the judgements made by the participants, 25.76 per cent were made with 90–99 per cent confidence and 29.80 per cent with 100 per cent (i.e. total) confidence.

Accuracy and confidence were positively and significantly correlated across items ($r(87) = .69$, $p < .001$), whereas accuracy and overconfidence were found to be significantly negatively correlated ($r(87) = -.76$, $p < .001$), with high levels of overconfidence being associated with low levels of accuracy. The correlation between confidence and overconfidence over items was non-significant. It is apparent, therefore, that it is variation in accuracy rather than variation in confidence that accounts for variation in overconfidence across base rates.

The effect on overconfidence of the second factor, outcome positivity/negativity, was also significant, with mean overconfidence higher for positive-outcome items ($M = 3.05$) than for negative-outcome items ($M = 1.42$; $F(1,65) = 12.39$, $p < .001$). There was no significant interaction between outcome positivity/negativity and range. The difference in accuracy between positive-outcome items ($M = 83.61$) and negative-outcome items ($M = 80.24$) was not significant. Confidence, however, was significantly higher for positive-outcome items ($M = 86.67$) than for negative-outcome items ($M = 81.66$; $F(1,65) = 7.25$, $p < .01$) (see Table 2).

Table 2. Differences between positive-outcome and negative-outcome events in base rates, predictions, accuracy, confidence and overconfidence of judgements

	Items			
	All	Positive	Negative	*F*
Predicted	51.59	56.06	46.10	1.88
Base rate	47.54	51.59	42.56	2.34
Confidence	84.42	86.67	81.66	7.25*
Accuracy	82.10	83.61	80.24	0.87
Overconfidence	2.32	3.05	1.42	12.39*
Estimated base rate	44.52	48.91	39.11	1.11
Directional accuracy	−3.02	−2.68	−3.45	0.53
Absolute accuracy	13.05	12.63	13.56	0.75

*$p < .01$.

Figure 2 shows that the inverted U-shaped relationship between base rates and overconfidence was present for both positive-outcome and negative-outcome items. The figure also shows that, for both types of items, overconfidence occurs only in the intermediate base rate ranges, with no overconfidence, but rather *underconfidence*, occurring in both low and high base rate ranges. Positive-outcome items elicited more overconfidence than negative-outcome items in all base rate ranges except range 6 (40–50 per cent base rates) where overconfidence was about equal (see Fig. 2).

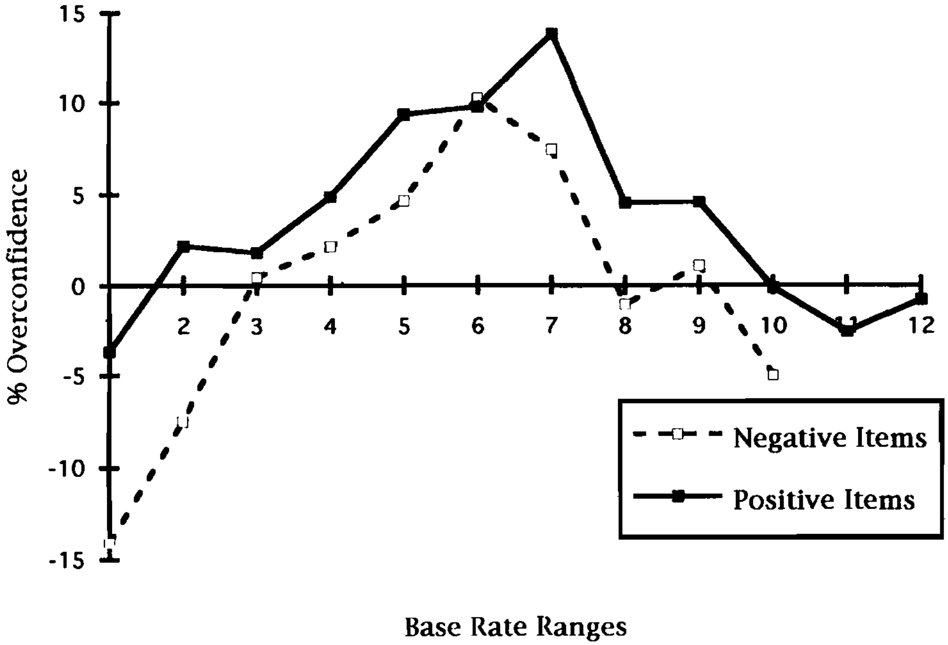

Figure 2. Mean overconfidence for positive-outcome and negative-outcome items across base rate ranges from extremely rare (range 1 = 0–3 per cent), through moderately common (range 6 = 40–50 per cent, range 7 = 50–60 per cent), to extremely common (range 12 = 97–100 per cent).

Across items, the numbers of participants predicting that each event would occur correlated highly with the numbers reporting that the events did in fact occur ($r(86) = .97$, $p < .001$). Over all items, participants predicted that 51.59 per cent of the specified events would occur. The actual number of events that did occur (the base rate) was 47.54 per cent, so participants were overpredicting that the events would occur by 4.05 per cent ($t(86) = 4.78$, $p < .001$). When this overprediction is broken down, no significant effects of outcome positivity/negativity can be seen: participants predicted that 56.06 per cent of positive-outcome events would occur, an overestimate of 4.47 per cent, and that 46.10 per cent of negative events would occur, which again was an overestimate of 3.54 per cent.

Participants' estimates of item base rates correlated highly with objective base rates ($r(87) = .88$, $p < .001$), and the difference between the mean estimated base rate (44.52) and the true mean base rate (47.54) was not significant. Thus, generally speaking, participants were able reasonably accurately to estimate the objective base rate of the events. The relationship between objective base rates and estimates of base rates is more complicated than this, however, as shown in Fig. 3.

Figure 3 shows that when the objective base rate of the item was less than 40 per cent (i.e. below base rate range 6), participants tended to overestimate its base rate by

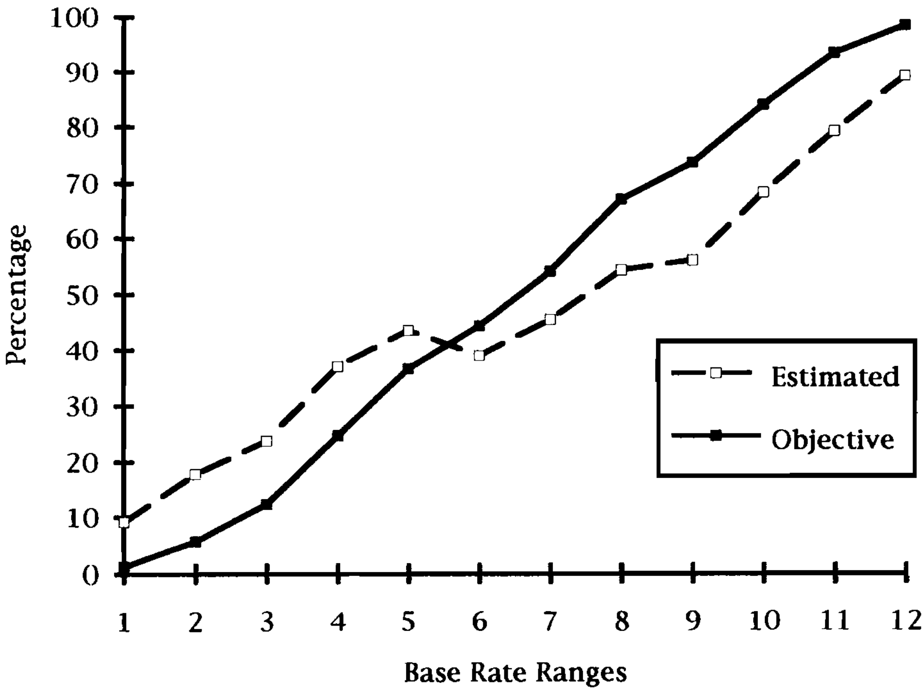

Figure 3. Objective and estimated base rates from low (range 1 = 0–3 per cent), through moderately common (range 6 = 40–50 per cent), to extremely common (range 12 = 97–100 per cent).

10.79 per cent on average, whereas for base rates above 40 per cent they tended to underestimate the base rate by 12.78 per cent on average ($z = 5.28$, $p < .001$). Outcome positivity/negativity had no significant effect on the participants' estimates of the base rate or the accuracy of their estimates. The objective base rates for positive-outcome and negative-outcome items were not significantly different.

The accuracy of base rate estimates was found by subtracting objective base rates from estimated base rates—this was a directional measure. There was no significant correlation between accuracy of base rate estimates and overconfidence. This indicates that it was not inaccuracies in estimating the base rate that produced overconfidence in an item. The accuracy of base rate estimates correlated significantly with the accuracy of predictions ($r(87) = .23$, $p < .05$), and with rated confidence of judgements ($r(87) = .31$, $p < .01$). Thus overestimates of the base rates were associated with higher accuracy and confidence in items. Taking an absolute deviation measure of the estimated base rate from the actual base rate reveals that more accurate estimates of base rates (i.e. lower discrepancies) are associated with significantly higher confidence ($r(87) = -.23$, $p < .05$), but there was no significant correlation between the absolute deviation and accuracy or overconfidence. The correlation between the objective base rates and the absolute accuracy of base rate estimates was also non-significant.

Table 3 clearly shows that there is no overall overconfidence (−2.12 per cent) for

Table 3. Predictions with and against base rate: confidence, accuracy and overconfidence

	Number of predictions	Mean % confidence	Mean % accuracy	Mean % overconfidence
With base rate	6409			
BR 0–50 (No)	3210	86.80	89.90	−3.10
BR 50–100 (Yes)	3199	84.27	85.42	−1.14
Mean		85.54	87.66	−2.12
Against base rate	2002			
BR 0–50 (Yes)	1138	76.72	46.37	30.36
BR 50–100 (No)	864	77.69	63.59	14.10
Mean		77.14	53.80	23.34

the 76.20 per cent of the judgements that were in line with the base rate. All of the overconfidence resulted from the 23.80 per cent of judgements that were made against the base rate, resulting in 23.34 per cent overconfidence. Participants did lower their confidence when making predictions against the base rate ($M = 77.14$ vs. $M = 85.54$) but evidently not sufficiently: their accuracy was only 53.80 per cent for predictions against base rate compared with 87.66 per cent for predictions in line with base rate.

The lowest level of accuracy ($M = 46.37$) resulted when participants predicted against the base rate that they would experience events when, objectively, most people would not in fact experience the event (i.e. when the base rate was less than 50 per cent). Conversely, when participants predicted that events would not happen to them when the base rate suggested that they probably would (i.e. when the base rate was more than 50 per cent), accuracy was higher ($M = 63.59$). Participants were approximately 77 per cent confident when making predictions against the base rate, and this results in the most erroneous (highly overconfident) predictions being for those events where the majority of the population do not experience an event but the participant predicted that it would happen to him/her—this resulted in 30.36 per cent overconfidence. Much less overconfidence (14.10 per cent) occurred when participants predicted that they would not experience events that most members of that population did in fact experience.

The sex of the participants had no discernible effect upon any of the variables measured. Although males had a mean level of overconfidence of 3.76 per cent and females had a lower level of 1.76 per cent, this difference was not significant.

Discussion and conclusions

Across all items combined, participants' confidence exceeded their accuracy, and this discrepancy produced a significant level of overconfidence in their judgements. These findings corroborate those of previous researchers (e.g. Dunning *et al.*, 1990; Vallone *et al.*, 1990) who have reported that confidence consistently exceeded the objective accuracy of predictions, especially when confidence was relatively high. These earlier results showed that predictions made with high levels of confidence tended to result in higher levels of overconfidence, but in the experiment reported in this article items about which people were highly confident were not necessarily the ones that resulted in high overconfidence, because the correlation between confidence and overconfidence was non-significant. This apparent inconsistency with earlier reported findings can be accounted for by the inclusion of items at the extremes of the base rate range, associated with very high accuracy and confidence, which produced underconfidence rather than overconfidence. Previous studies have found more overconfidence for predictions made with high confidence, because they did not include very predictable extreme base rate items; for example, Vallone *et al.* found 20.2 and 22.6 per cent overconfidence for judgements made with confidence over 90 per cent and with 100 per cent confidence respectively. As regards item positivity/negativity, the results reported in this article show that overconfidence was significantly greater for positive-outcome items than for negative-outcome items, as shown in Fig. 2.

The outcome positivity/negativity of the events being judged had no significant effects on the participants' predictions, their predictive accuracy, their base rate estimates or their accuracy of estimating base rates. Therefore, participants do not seem to have been overpredicting positive events and underpredicting negative ones, which the theory of unrealistic optimism (Weinstein, 1980) would suggest. What *is* affected by outcome positivity/negativity is the level of confidence and overconfidence in the judgements, which were both inflated for positive-outcome events. The results reported in this article are underestimates of the difference in overconfidence between positive-outcome and negative-outcome events, because there are no negative-outcome events at the top of the base rate range that would have lowered the mean overconfidence even further below the level for the positive items.

The difference between this study and the previous ones (e.g. Weinstein, 1980) that have reported unrealistic optimism may be due to the time scale sampled. In our experiment, the predicted events were to occur within one week, whereas in other studies the time scale was in terms of participants' whole future lives. In our experiment, participants overpredicted the occurrence of events for both positive-outcome and negative-outcome items, leading to unrealistic optimism of predictions observed for positive-outcome items only, and unrealistic *pessimism* for negative-outcome items. The level of overprediction for both types of items was relatively low, probably because there was a high level of reliable information available to the participants for making their predictions, thus reducing possible informational retrieval biases. Unrealistic optimism may be more likely to occur for predictions involving a long time scale, because of cognitive problems, and less likely for short-term, everyday judgements.

The unrealistic optimism found in the confidence and overconfidence of judgements involving positive-outcome events can probably be explained by motivational factors. Perhaps unrealistic optimism in *predictions* taps the informational/cognitive aspect of judgements and unrealistic optimism shown in *confidence* ratings reflects the motivational/emotional aspects. Weinstein (1980) did not find any effect of degree of outcome desirability on unrealistic optimism of predictions. This may be because in his research the events were all seen as either extremely desirable or extremely undesirable, so there may have been a ceiling effect. The events used in our experiment were much less emotive than Weinstein's and produced lower levels of unrealistic optimism for positive-outcome items and unrealistic pessimism for negative-outcome items. This suggests that there may be a motivational component operating in unrealistic optimism and that motivation was lower in our experiment because of the less emotive items used. Another consideration is that whereas in previous research participants compared their likelihood of experiencing each event with other people's likelihoods of experiencing those same events, in the experiment reported here, participants' self-reported predictions were compared with *objective* self-outcomes, which may provide information about the *absolute*, rather than the *relative*, level of unrealistic optimism.

A possible likely explanation for the significantly higher confidence for positive-outcome than negative-outcome items may lie in the Pollyanna effect (Boucher & Osgood, 1969), named after the unrealistically optimistic eponymous heroine of Eleanor Porter's 1913 novel and the films based on it. According to the Pollyanna effect, people tend to exaggerate the positive and minimize the negative aspects of their lives. It is hardly surprising in this light that people tend to feel more confident about the likelihood of positive than negative experiences, as clearly shown in Table 2, and one consequence of this is that overconfidence was greater for positive-outcome than negative-outcome items, as shown in Fig. 2.

Alternatively, Ginossar & Trope (1980) proposed that when more information has to be processed confidence drops, and the differences in overconfidence between positive-outcome and negative-outcome events could be due to differences in information loads for the two types of task. In our experiment positive-outcome and negative-outcome events elicited similar levels of overconfidence in the 40–50 per cent base rate range, and this may be due to similar levels of information being processed in this range where uncertainty is near to 50:50. Further empirical evidence will be needed to determine if there are different amounts of information processed in predictions for positive-outcome and negative-outcome events across the base rate range. The participants may have more confidence in their positive-outcome judgements because they feel they have more control over their occurrence in comparison to negative-outcome events, and the relation of confidence to perception of control also needs addressing in future research.

Perhaps the most important finding of this experiment was that overconfidence varied across the base rate range, which resulted in a curvilinear, inverted U-shaped relationship between base rates and overconfidence shown clearly in Figs 1 and 2. Overconfidence was greatest for intermediate base rate ranges 5 to 7, which spanned base rates between 30 and 60 per cent where there was maximal uncertainty as to whether or not the events would occur. This shows the hard–easy effect, that the greatest levels of overconfidence occur when accuracy is low due to uncertainty being high. Participants evidently expressed excessively high levels of confidence and failed to consider or allow for the levels of uncertainty of the situations in the middle base rate ranges. In extremely low base rate ranges 1 and 2 (corresponding to base rates between 0 and 10 per cent) and extremely high base rate ranges 10 to 12 (corresponding to base rates between 80 and 100 per cent), overconfidence was not observed, and in fact *underconfidence* occurred. This finding is in line with an earlier finding of Lichtenstein, Fischhoff & Phillips (1982), and is accounted for by the fact that very easy items generally fall at the extremes of the base rate range. It seems reasonable to conclude that the overconfidence effect is confined to events with intermediate base rates of occurrence. Extreme base rates do not appear to have been adequately investigated in previous research into the overconfidence effect.

These findings throw some light on the psychological mechanism underlying the overconfidence effect. They suggest that the effect may be explained by a tendency for participants to make insufficient use of base rate information in making self-ratings of confidence. The tendency of participants to ignore base rates when making self-predictions has been noted before, though in circumstances where such information is potentially less helpful (Osberg & Shrauger, 1986). Our findings do not suggest that the participants were oblivious of base rates, or that they were unable to estimate base rates or make any use of them. Furthermore, the overconfidence effect does not appear to be due to inaccuracies in participants' estimates of base rates, because there was no significant correlation between the accuracy of base rate estimates and overconfidence. Figure 3 shows a strong relationship between objective and estimated base rates ($r = .88$), and the mean difference between estimated and objective base rates was small and non-significant. Participants were thus able to estimate the objective base rate of the events reasonably accurately, although Fig. 3 shows clearly that they tended to overestimate low base rates (below 40 per cent) and to underestimate high base rates (above 40 per cent). Figure 1 shows that accuracy of predictions was highest in regions of extremely low and extremely high base rates, where base rates provide the most useful information as to whether or not an event will or will not occur, and lowest in intermediate base rate ranges, which shows that participants were probably using base rate information either directly or indirectly to inform their predictions.

Figure 1 also shows, however, that the participants took less account of base rates in modulating their ratings of confidence in their predictions. However, items for which they could accurately estimate base rates also generated higher confidence ratings. In intermediate base rate ranges where accuracy was lowest, confidence ratings were hardly lower than in extreme base rate regions where accuracy was significantly higher. The overconfidence effect is evidently explained by the gap between accuracy and confidence resulting from unrealistically high confidence in relation to the accuracy of difficult predictions for events with intermediate base rates of occurrence. The effect apparently disappears, and is even reversed, for events with extremely low or high base rates of occurrence, as shown in Figs 1 and 2.

The predictions against base rate reveal that denying that an event would occur to oneself, when it happens to most people, results in overconfidence, and is most common in the case of items that are negative in outcome. Much higher overconfidence occurs when people want an (positive) event to happen to them when the base rate is low, so it is unlikely to happen. This may indicate that self-presentation biases are operating, for example to protect self-esteem, or perhaps that there is a higher level of selective recall for information when events are positive in outcome.

Accuracy dropped in the middle base rate ranges where the statistical uncertainty of the judgements was highest. The lack of equal change of confidence estimates in the middle base rate ranges may be due to the effect of a general level of task confidence that is appropriate to a task but not flexible within that task. Having high confidence in a judgement or belief may make it more resistant to change, and this has obvious implications for health-related and risky forms of behaviour. Strategies for improving people's calibration of confidence may result in more effective attitude change and a less unrealistic perception of risky behaviours.

The principal aim of the experiment reported here was to investigate the relationship between base rates and overconfidence. The results show clearly that the relationship is strongly curvilinear, with overconfidence peaking at intermediate base rates around 40 to 60 per cent and decreasing for both lower and higher base rates. The experiment also sought to clarify the underlying mechanism of the overconfidence effect by determining whether people are aware of base rates but cannot make full use of them, or whether they are simply unaware of them. The results suggest that participants were probably aware of base rates, and made use of them directly or indirectly in formulating predictions, but that they did not take adequate account of them in adjusting their confidence ratings. It would be interesting to know whether these conclusions apply not only to self-predictions but also to social predictions.

What emerges from the results reported here is that the so-called overconfidence effect does not appear to be a unidirectional bias in human judgement across the full range of item base rates, as hitherto assumed, but a more subtle and complex phenomenon. It is evidently a bidirectional bias involving overconfidence for items with intermediate base rates of occurrence and underconfidence for items at both extremes of the base rate range, although in experiments using only items with intermediate base rates, it may appear to be a unidirectional overconfidence effect. Researchers investigating this phenomenon should in future take care to include items from across the entire spectrum of base rates.

Acknowledgements

Preparation of this article was supported by a research studentship awarded by the Economic and Social Research Council to the first author. The authors would like to thank Rachel ffield and Wendy Green for their valuable help in data collection and input.

References

Adams, J. K. & Adams, P. A. (1961). Realism of confidence judgements. *Psychological Review*, **68**, 33–45.

Ajzen, I. (1977). Intuitive theories of events and the effect of base rate information on prediction. *Journal of Personality and Social Psychology*, **35**, 303–314.

Bar-Hillel, M. (1980). The base-rate fallacy in probability judgments. *Acta Psychologica*, **44**, 211–233.

Boucher, J. & Osgood, C. E. (1969). The Pollyanna hypothesis. *Journal of Verbal Learning and Verbal Behaviour*, **8**, 1–8.

Dunning, D., Griffin, D. W., Milojkovic, J. D. & Ross, L. (1990). The overconfidence effect in social prediction. *Journal of Personality and Social Psychology*, **58**, 568–581.

Gilbert, D. T. & Jones, E. E. (1986). Perceiver-induced constraint: Interpretations of self-generated reality. *Journal of Personality and Social Psychology*, **50**, 269–280.

Ginossar, Z. & Trope, Y. (1980). The effects of base rates and individuating information on judgments about another person. *Journal of Experimental Social Psychology*, **16**, 228–242.

Johnson, J. T., Jemmott, J. B. & Pettigrew, T. F. (1984). Causal attribution and dispositional inference: Evidence of inconsistent judgments. *Journal of Experimental Social Psychology*, **20**, 567–585.

Jones, E. E. & Harris, V. A. (1967). The attribution of attitudes. *Journal of Experimental Social Psychology*, **3**, 1–24.

Lichtenstein, S., Fischhoff, B. & Phillips, L. D. (1982). Calibration of probabilities: The state of the art to 1980. In D. Kahneman, P. Slovic & A. Tversky (Eds), *Judgment Under Uncertainty: Heuristics and Biases*, pp. 306–334. Cambridge: Cambridge University Press.

Nisbett, R. E. & Ross, L. (1980). *Human Inference: Strategies and Shortcomings of Social Judgment.* Englewood Cliffs, NJ: Prentice-Hall.

Ross, L. (1977). The intuitive psychologist and his shortcomings: Distortions in the attribution process. *Advances in Experimental Social Psychology*, **10**, 173–220.

Ross, L. D., Amabile, T. M. & Steinmetz, J. L. (1977). Social roles, social control and biases in social-perception processes. *Journal of Personality and Social Psychology*, **35**, 485–494.

Osberg, T. M. & Shrauger, J. S. (1986). Self-prediction: Exploring the parameters of accuracy. *Journal of Personality and Social Psychology*, **51**, 1044–1057.

Tversky, A. & Kahneman, D. (1980). Causal schemas in judgments under uncertainty. In M. Fishbein (Ed.), *Progress in Social Psychology*, vol. 1, pp. 49–72. Hillsdale, NJ: Erlbaum.

Vallone, R. P., Griffin, D. W., Lin, S. & Ross, L. (1990). Overconfident prediction of future actions and outcomes by self and others. *Journal of Personality and Social Psychology*, **58**, 582–592.

Webster, D. M. (1993). Motivated augmentation and reduction of the overattribution bias. *Journal of Personality and Social Psychology*, **65**, 261–271.

Weinstein, N. D. (1980). Unrealistic optimism about future life events. *Journal of Personality and Social Psychology*, **39**, 806–820.

Received 10 January 1995; revised version received 21 April 1995