

# Communicating Likelihood and Managing Face: Can we Say it is Probable when we Know it to be Certain?

Jean-François Bonnefon (bonnefon@univ-tlse2.fr)

Centre National de la Recherche Scientifique

Laboratoire Travail et Cognition, 5 allées A. Machado 31058 Toulouse Cedex 9 France

Gaëlle Villejoubert (g.villejoubert@lubs.leeds.ac.uk)

Leeds University Business School, Maurice Keyworth Building, The University of Leeds, LS2 9JT United Kingdom

## Abstract

Different functions can be performed by probability phrases such as: ‘it is probable/possible/likely that  $x$ .’ Mainly, speakers may communicate a vague judgment about the likelihood of event  $x$ , or they may wish the recipient to focus on reasons for the occurrence of event  $x$ . We argue that there is another communicative function which has yet to be documented, namely, the face-management function. Such function consists in mitigating threat to the addressee when  $x$  is a criticism or an imposition. Data show that the phrases ‘possibly’ and ‘probably’ are indeed understood differently (have different membership functions) depending on whether they modify neutral or face-threatening contents. We consider the potential misunderstandings and judgmental mistakes that may arise from ambiguity about which function of verbal uncertainty is being performed.

## Functions of Verbal Uncertainty

Imagine a conversation between two very close friends, one of whom is complaining that she gets little attention from men. The other friend then replies: ‘It is probable that your bad breath makes men uncomfortable.’ What does she mean by ‘probable’?

From the onset of the study of the meaning verbal probabilities, researchers have proposed that expressions such as *it is probable that  $x$*  or *it is possible that  $x$*  are used to express, albeit vaguely, a judgment about the likelihood of the proposition or event they predicate (e.g., Lichtenstein & Newman, 1967; Wallsten et al., 1986). Thus, the sentence ‘it is probably going to rain tomorrow’ means that the speaker has some intermediate confidence in the prediction that it is going to rain, a confidence that could perhaps be expressed as a 70% chance, but definitely not as a 99% chance.

Starting from the assumption that probabilistic phrases are vague concepts in the sense of fuzzy set theory (Zadeh, 1965), Wallsten et al. (1986) and Rapoport, Wallsten, and Cox (1987) suggested that the meaning of a probability phrase could be represented by its membership function. The membership function of any given phrase assigns a number to each value on the probability line  $[0, 1]$  that represents its degree of membership in the concept defined by the phrase. Degree of membership is usually expressed as a real number from 0 to 1, such that memberships of 0 denote probabilities that are absolutely not in the concept, and memberships of 1 denote elements that are perfect exemplars of the con-

cept. All other values represent intermediate degrees of membership.

The peak of a membership function is the probability (or the mean of the probabilities) at which the function reaches its maximum value. A number of studies (reviewed in Budescu & Wallsten, 1995, and Wallsten & Budescu, 1995) have shown most probability phrases to be single peaked. For instance, the membership function of ‘possible’ increases monotonically until it peaks at around .7, then decreases monotonically. Only very specific phrases, such as ‘certain’ or ‘impossible’ monotonically increase or decrease all the way through the probability interval  $[0, 1]$ .

Membership functions indicate the extent to which each probability value may ‘typically’ be communicated by a given phrase. The possibility that a given phrase may translate into different probability values at different times is thus captured by this approach. Membership functions, however, are taken to portray the meaning of probability phrases independently of the context in which they are used. In other words, this approach precludes the possibility that the membership function of a given phrase would vary as a function of the context in which it is evaluated.

Teigen and Brun (1995, 1999) proposed that probability phrases can not only convey meanings of vague likelihood but also have communicative functions that are embedded in them. Such a communicative function, directionality, invites the hearer of the message to focus on the occurrence or the non-occurrence of the event being predicated, depending on the polarity of the phrase used by the speaker. For example, saying ‘it is unlikely that it will rain tomorrow’ rather than ‘there is a small chance that it will rain tomorrow’ focuses on the absence rather than the presence of rain. However, Budescu, Karelitz, & Wallsten (2003) showed that phrases’ directionality could be predicted from properties of their membership function. Consistent with the authors’ predictions, positive phrases were found to be negatively skewed and have peaks above 0.5 whereas negative phrases tended to be positively skewed and have peaks below 0.5. Moreover, the authors demonstrated that this relation was invariant across contexts.

Intuitively, however, neither of these interpretations of the meaning of verbal probability phrases seems to account for the meaning conveyed by the speaker who says: ‘It is probable that your bad breath makes men

uncomfortable.’ Indeed, it is believable that the speaker is quite confident if not certain that it is her friend’s bad breath that explains her romantic misfortunes. Similarly, directionality does not seem to play a role as the use of a negative phrase seems anomalous: ‘It is *unlikely* that your bad breath makes men uncomfortable’.

In this article, we wish to demonstrate that verbal probabilities phrases sometimes have another communicative function, more precisely a *face-management* function. Such a function will be at use when the speaker aims to avoid hurting somebody else’s feelings. Contrary to Budescu et al.’s (2003) argument that phrases membership functions are invariant across contexts, we will show that the phrases ‘probably’ and ‘possibly’ are understood very differently when applied to neutral contents and when applied to face-threatening contents (i.e., criticisms or impositions). We begin with a brief presentation of the face management function of probability phrases. We then argue that a given probability phrase will translate into quite different membership functions depending on its communicative function. Briefly, we predict that whilst the probability phrases ‘possibly’ and ‘probably’ have single peaked membership functions when they modify neutral contents, those membership functions will be monotonically increasing (or, at least, will peak much closer to 1) when they modify face-threatening acts. We report an experiment whose results support our predictions. Finally, we discuss how to extend our results and investigate situations wherein untoward consequences follow from some ambiguity about which function of verbal uncertainty is being performed. As for now, we will introduce the notion of a face-threatening act, and the associated pragmatic function of verbal uncertainty.

### Managing Face

The core assumption of Politeness Theory (Brown & Levinson, 1978/1987) is the universal want to maintain ‘face’. Face is decomposed into two components, *negative face*, the want to be free of action or from imposition, and *positive face*, the want to be approved of, liked, or admired. Any action or utterance that can be construed as a restriction or an imposition onto someone is a threat to negative face. Likewise, any behavior or utterance that can be construed as disapproval or criticism is a threat to positive face.

When an utterance amounts to a threat to positive or negative face (i.e., when it qualifies as a face-threatening act, henceforth an FTA), the speaker must decide on a politeness strategy to mitigate the threat. Politeness Theory describes a number of such possible strategies, and the considerations that drive the choice of a strategy for a given FTA in a given context. We do not need, however, to consider in further detail the issue of strategy choice, for the present report will only feature one specific strategy known as ‘hedging’. More precisely, among the numerous forms of hedging, we will only consider how a speaker might politely weaken a statement she knows to be true.<sup>1</sup>

<sup>1</sup>Hedges are not limited to probability phrases, and hedg-

For instance, a speaker might try to gently disagree with an addressee (a positive FTA) by saying ‘I think perhaps you’re wrong’, or ‘you’re possibly wrong here’. In other terms, the speaker will modify her positive FTA by a probability phrase. Likewise, a speaker might try to soften an imposition upon an addressee (a negative FTA) by saying ‘I guess you may have to pay for my lunch’ or ‘you’ll probably have to pay for my lunch’. Again, the speaker modifies her negative FTA by a probability phrase, for politeness concerns.

Note that the speaker is not, strictly speaking, modifying the likelihood of her disagreeing or imposing something upon the addressee. If the addressee recognizes the politeness strategy for what it is, he should not understand the probability phrase as communicating a judgment of likelihood, but simply as a politeness marker. Now if we assume that differences in the meaning of probability phrases translate into differences in membership functions, we should expect different membership functions for a given probability phrase, depending on whether it modifies an FTA or not.

### Different Membership Functions?

A number of studies established that the phrases ‘probably’ and ‘possibly’ have single peaked membership functions. We claim that this result only generalizes to neutral (i.e., not face-threatening) contents. When ‘probably’ and ‘possibly’ modify an FTA (a criticism or an imposition), they will perform a pragmatic, face-management function of communication. As a consequence, their prototypical probability value ought to be close to 1, resulting in a monotonically increasing membership function, or, alternatively, in a membership function peaking much closer to 1 than for neutral content.

### Method

**Participants** A total of 112 students at Leeds University Business School took part to the experiment; 49 were men, 57 were women, 6 did not provide gender information. Age ranged from 18 to 44 (median = 19). Most participants were management majors.

**Material and Procedure** Membership functions were elicited using the Multiple Stimuli Method introduced in Budescu et al. (2003). Participants were presented in turn with 8 statements of the form ‘[speaker] tells you that [modifier, content]’. The speaker was identified each time with a different first name, alternatively male and female. The modifier was a between-subject variable, either ‘possibly’ or ‘probably’. The type of content was a within-subject variable: Four contents were neutral and the otherfour were face-threatening (two criticisms and two impositions). The four neutral contents were taken from Budescu et al. (2003) and adapted to a British setting, for instance by speaking of Scotland rather than Wisconsin, and of rugby rather than basketball. See Table 1 for a complete list of neutral contents and FTAs.

ing as a strategy is not limited to the preservation of the addressee’s face, but can perform a vast number of functions (Channell, 1994; Jucker, Smith, & Lüdge, 2003).

Table 1: Neutral statements and face-threatening acts (FTAs) used as experimental material.

Content type	Sentence
Neutral	In the near future the wages in some areas of the country will increase slightly
	Next week Scotland will experience a major snow storm
	In the next few years rugby teams from the south-west will dominate the sport
	I will find an excellent job immediately after I graduate from university
FTAS	You will have to drive me to the airport
	You have made a big mistake
	You will have to pay for the cinema tickets
	Your bad breath makes people uncomfortable

Alongside with each of the eight statements, participants were presented with 10 probabilities (introduced in the form of percentages, from 10% up to 100%) and had to judge to what extent the speaker may have had this probability in mind when uttering the statement. This judgment was expressed on a 10-point scale anchored at ‘absolutely not’ to the left and ‘absolutely’ to the right. For example, after having read that ‘Eileen tells you: You have probably made a big mistake’, participants were asked the following question: ‘Can Eileen think the probability that you have made a big mistake is...’ followed by the 10 possibilities 10%, 20%, up to 100%, and had to provide a judgment for each of the 10 possibilities.<sup>2</sup> The order of statements was counterbalanced across questionnaires. Questionnaires were filled out during class. Participation was optional and anonymous, but participants had the option to provide an email address so as to take part in a prize draw. Two participants won a prize of £25 in vouchers. The experiment was conducted in English.

## Results and discussion

Mean membership of each probability from 0.1 to 1.0 was computed by averaging membership judgments across participants. Both for ‘possibly’ and ‘probably’, membership functions were remarkably similar for the four neutral contents on the one hand, and for the four FTAs on the other hand. Figures 1 and 2 respectively depict the membership function of ‘probable’ and ‘possible’ as a function of content.

The ‘peak’ of a membership function was computed by (a) averaging for each participant the probability values which received the highest membership ratings, then (b) averaging across participants the values so obtained – or, alternatively, considering the median of the values so obtained. (Note that some small discrepancy is not uncommon between the numerical computation of the peak and its graphical representation: Thus the values of Table 2 do not exactly match what appear to be the peaks of the membership function in Figures 1 and 2.) With respect to neutral contents, we obtain classic, single peaked

<sup>2</sup>Unlike Budescu et al. (2003), we did not ask participants about the membership of probability 0. We were concerned that asking up front a rather infelicitous question might compromise the credibility of the task.

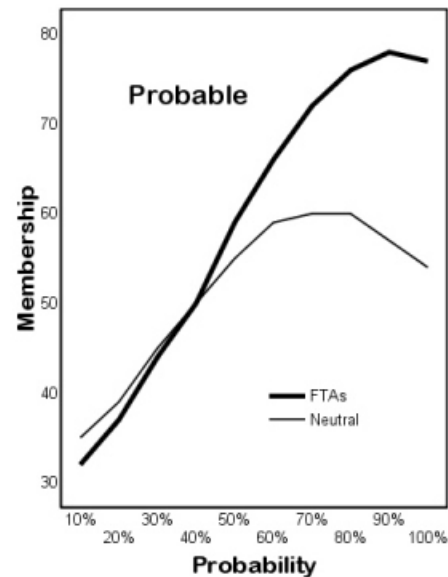


Figure 1: Membership function of ‘probable’ for neutral content and for face-threatening acts (FTA).

membership functions for ‘possible’ and ‘probable’. The two membership functions peak at a little below .6 and .7, respectively (see Table 2). The membership functions with respect to FTAs are clearly different, even though they do not support the strong version of our hypothesis. The membership function of ‘possible’ is clearly single peaked rather than monotonically increasing, even for FTAs. The membership function of ‘probable’ is nearly monotonically increasing, but does feature a small decrease at the very end of the probability interval.

The weaker version of our hypothesis is, however, clearly supported by the data, as membership functions peak much closer to 1 for FTAs than for neutral contents (see Table 2). This is especially true if one considers the median value for the distribution of peak across participants, rather than the mean value which in this case appears to be a little too sensitive to outliers.

The membership function of ‘probable’ peaks reliably higher for FTAs than for neutral contents,  $t = 6.0$  ( $n =$




Figure 2: Membership function of ‘possible’ for neutral content and for face-threatening acts (FTA).

Table 2: Mean and median peaks of the membership functions of “possible” and “probable” depending on the content they modify (neutral vs. face-threatening).

	Mean ( <i>SD</i> )	Median
Possible		
Neutral	.58 (.26)	.55
FTA	.70 (.25)	.74
Probable		
Neutral	.67 (.23)	.66
FTA	.80 (.19)	.85

53),  $p < .001$ ; the 95% confidence interval for that difference is [.08, .16] and Cohen’s  $d$  is .48, a moderate effect size. Likewise, the membership function of ‘possible’ peaks reliably higher for FTAs than for neutral contents,  $t = 5.2$  ( $n = 59$ ),  $p < .001$ ; the 95% confidence interval for that difference is [.08, .18] and Cohen’s  $d$  is .61, again a moderate effect size.

The median peak of the membership functions differs widely as a function of content, with a gap of nearly .20 between the prototypical probability of FTAs and the prototypical probability of neutral contents (both for ‘possibly’ and ‘probably’). This difference is statistically reliable,  $Z = 4.7$  for ‘possible’ and  $Z = 4.5$  for ‘probable’,  $p < .001$  in both cases.

Although data do not provide full support to the stronger version of our hypothesis (especially since the membership function of ‘possible’ does not seem to be monotonic even when the phrase modifies an FTA), they clearly validate its weaker version. The meaning of ‘probable’ and ‘possible’, as assessed through

their membership functions, is different depending on whether these probability phrases modify neutral or face-threatening content. The prototypical probability associated to a probability phrase is much higher when this phrase modifies face-threatening content than when it modifies neutral content. Moreover, this result was obtained using a *within-subject design*, that is, at the individual rather than group level.

## General Discussion

The preliminary investigation we have just reported lends credence to the idea that people can and do distinguish between the likelihood-communication function of verbal uncertainty, and its pragmatic, face-management function. Now that our initial assumption has been experimentally grounded, we intend to take this research in two complementary directions. First, we need to extend our initial result to other probability phrases, other hedges than probability phrases, and to a greater variety of FTAs. Second, we intend to explore the limits of people’s ability to distinguish between the two functions, and to investigate the types of mistakes that can arise from confusion between the two functions. We will now consider in turn these two aspects of our research program.

## Extensions

**Other phrases** As of now, we have only considered the phrases ‘possible’ and ‘probable’. There seems to be no reason why our analysis should not apply to some other common probability phrases such as ‘fair chance’ or ‘likely’, or variations such as ‘very probable’ or ‘somewhat likely’. Modal verbs such as ‘I guess’, ‘I suppose’, or ‘I think’ would also fall within the scope of our analysis. Indeed, Brown and Levinson (1978/1987) lists a number of probability phrases that may be used to soften or weaken a FTA, not to mention prosodic and kinesic cues to uncertainty such as raised eyebrow, hands thrown up, *umms* and *errs*, or even the use of pitch.

**Other FTAs** Our approach to FTAs has been rather crude insofar as we have only considered two kinds of threats, both directed towards the addressee’s face. Other threats to the addressee’s face include suggestions, reminders, and dares (negative FTAs), or bad news, non-sequiturs, and the raising of divisive topics (positive FTAs). In addition, threats to the speaker’s own face could be worth considering, such as expressions of gratitude, unwilling offers, and responses to faux-pas (negative FTAs), or apologies, confessions, or acceptance of compliments (positive FTAs).

Aside from the intrinsic value of generalizing our initial findings, the consideration of other FTAs would allow for some control over a potential bias of our experimental material. All the FTAs we have considered include a direct reference to the addressee (they all began by ‘You’ or ‘Your’), whilst none of the neutral contents do. Using the various FTAs we have just introduced will eliminate the need for systematic reference to the addressee.

**Other hedges** Just as probability phrases can modify the likelihood of a proposition, other phrases can modify other aspects of information. For example, the phrase ‘some’ modifies quantity the same way the phrase ‘probably’ modifies likelihood: Saying that ‘some vegetarians eschew milk’ conveys the information that not all do, the same way as saying that ‘Nicole will probably not have milk’ conveys the information that she might after all.<sup>3</sup> One may thus wonder whether the polite, face-management function of likelihood modifiers may have its counterpart with respect to quantity modifiers.

It is apparently the case that quantity modifiers can be understood as mere politeness markers when they modify an FTA. Imagine that you have been invited to give a seminar in a university you are visiting. After the seminar, you ask one of the local students what people generally thought of your talk, and her answer is: ‘Some thought you were terribly long.’ Does the use of ‘some’ here exclude the possibility that everybody felt that way? It rather seems that ‘some’ is used for face-management purpose in lieu of quantity modification purpose. The possibility still obtains that the whole audience thought you were too long, but that the student is too polite not to hedge her answer.

## Mistakes and Confusion

The results we have reported establish that people can distinguish between the likelihood communication function and the face management function of verbal uncertainty, and that they interpret probability phrases accordingly: More precisely, upon hearing that some event  $x$  is ‘probable’ or ‘possible’, they tend to think of  $x$  as having a higher numerical probability if it is an FTA rather than a neutral content. We are not claiming, however, that people are perfectly accurate in their assessment of which function a given probability phrase is meant to perform. On the contrary, we reckon that some situations can make it difficult to judge which function is being performed. People may thus mistake genuine uncertainty for politeness, or politeness for genuine uncertainty. We will now consider in turn two examples of potential confusion between face management and likelihood communication.

**Mistaking uncertainty for politeness** Sometimes, a speaker will want to modify an FTA with a probability phrase because she is genuinely uncertain about its likelihood. If the addressee understands the probability phrase to be a mere politeness marker, he will think of the FTA as having a higher probability than what the speaker had in mind. This may help explain a puzzling phenomenon reported in Weber and Hilton (1990):

<sup>3</sup>Both the inferences from ‘probably’ to ‘not certainly’ and from ‘some’ to ‘not all’ are considered *scalar implicatures* by pragmaticists (for a general neo-Gricean approach to scalar implicatures, see Levinson, 2000; for an experimental take at a Relevance theoretical approach of the implicature from ‘some’ to ‘not all’, see Bott & Noveck, in press). Conversational pragmatics thus pervades even what we have deemed the likelihood communication function of verbal uncertainty, just as it pervades most reasoning and decision making processes (see Hilton, 1995, for a review).

When a probability phrase modifies a medical diagnosis, its prototypical interpretation is positively correlated with the severity of the condition. For example, addressees will think of a much higher probability when told they ‘probably’ have cancer compared to when told they ‘probably’ have a common cold. Now, announcing bad news is an FTA according to Politeness Theory, and all the more threatening when the news is worse. Thus, diagnosing a common cold does not require a politeness strategy, but the need for such a strategy increases with the severity of the diagnosed condition. Hence, the more severe the condition is, the more likely the addressee is to consider the probability phrase as a politeness marker rather than a true likelihood modifier – and when an addressee does so, he will automatically think of the condition as having higher probability than otherwise, thus the result of Weber and Hilton (1990).

**Mistaking politeness for uncertainty** Just as numerical probabilities can convey a misleading impression of precision, verbal probabilities may convey a misleading impression of uncertainty. Sometimes, an addressee will not realize that an FTA is being performed, and will thus mistakenly think of a probability phrase as communicating genuine uncertainty. Consider the situation wherein a subordinate has to correct a superior. Correcting a mistake is intrinsically an FTA, and the need for politeness is made even more crucial by the fact the addressee has greater social status than the speaker. Thus, however certain the speaker is that the addressee has made a mistake, she will hedge her statement so as not to be rude. Now the addressee might understand the probability phrase as implying that the speaker is genuinely unsure whether the mistake is real, and thus fail to question and revise his initial beliefs. Being surrounded by people who systematically use probability phrases when correcting your mistakes is therefore likely to lead you to have unwarranted confidence in the accuracy of your own judgments and knowledge, a widely studied phenomenon known as ‘overconfidence’ (e.g., Klayman, Soll, Gonzalez-Vallejo, & Barlas, 1999). Indeed, Klayman and Burt (1998) report that overconfidence is prominent in ‘constrained’ social networks, i.e., small networks with a leader in a central, coordinating position, and weak connections to outsiders. Correcting the mistakes of the person at the center of such a network would arguably call for politeness strategies such as hedging. If the politeness strategy is mistaken for genuine uncertainty about whether the leader is truly wrong, overconfidence will quickly pervade the network.

## Conclusion

We have argued for the existence of two distinct functions of verbal uncertainty, which translate to different numerical probabilities for a given probability phrase. Whilst people appear to be sensitive to which function is being performed, the possibility obtains in some contexts for confusion and misunderstandings to arise. This risk is all the higher in constrained, hierarchized networks within which people make massive use of verbal uncertainty, have to assess the possibility of severe events, and

are prone to have divergent interpretations of the relevant information they work on. From this perspective, we are tempted to think of intelligence agencies as providing us with our worst-case scenario.

### Acknowledgments

Part of this research was conducted while Gaëlle Villejoubert was an invited lecturer at the Laboratoire Dynamiques Sociales et Vie Politique, Université de Toulouse-Le Mirail, France. We thank Denis Hilton and Ben Slugoski for helpful discussions about probabilities and pragmatics. We also thank Steven Sloman for helpful comments on an earlier version of this report.

### References

- Bott, L., & Noveck, I. A. (in press). Some utterances are underinformative: The onset and time course of scalar inference. *Journal of Memory and Language*.
- Brown, P., & Levinson, S. C. (1978/1987). *Politeness: Some universals in language usage*. Cambridge: Cambridge University Press.
- Budescu, D. V., & Wallsten, T. S. (1995). Processing linguistic probabilities: General principles and empirical evidence. In J. Busemeyer, D. L. Medin, & R. Hastie (Eds.), *Decision making from a cognitive perspective*. San Diego, CA: Academic Press.
- Budescu, D. V., Karelitz, T. M., & Wallsten, T. S. (2003). Determining the directionality of probability words from their membership functions. *Journal of Behavioral Decision Making*, 16, 159–180.
- Channell, J. M. (1994). *Vague Language*. Oxford: Oxford University Press.
- Hilton, D. J. (1995). The social context of reasoning: Conversational inference and rational judgement. *Psychological Bulletin*, 107, 248–271.
- Jucker, A. H., Smith, S. W., & Lüdtge, T. (2003). Interactive aspects of vagueness in conversations. *Journal of Pragmatics*, 35, 1737–1769.
- Karelitz, T. M., & Budescu, D. V. (2004). You say ‘probable’ and I say ‘likely’: Improving interpersonal communication with verbal probability phrases. *Journal of Experimental Psychology: Applied*, 10, 25–41.
- Klayman, J., & Burt, R. S. (1998). *Individual differences in confidence and experiences in social networks*. Working paper, Center for Decision Research, Graduate School of Business, University of Chicago.
- Klayman, J., Soll, J. B., Gonzalez-Vallejo, C., & Barlas, S. (1999). Overconfidence: It depends on how, what, and whom you ask. *Organizational Behavior and Human Decision Processes*, 79, 216–247.
- Lichtenstein, S. & Newman, J. R. (1967). Empirical scaling of common verbal phrases associated with numerical probabilities. *Psychonomic Science*, 9, 563–564.
- Levinson, S. C. (2000). *Presumptive meaning: The theory of generalized conversational implicature*. Cambridge, MA: MIT Press.
- Rapoport, A., Wallsten, T. S., & Cox, J. A. (1987). Direct and indirect scaling of membership functions of probability phrases. *Mathematical Modeling*, 9, 397–417.
- Teigen, K. H., & Brun, W. (1995). Yes, but it is uncertain: direction and communicative intention of verbal probabilistic terms. *Acta Psychologica*, 88, 233–258.
- Teigen, K. H., & Brun, W. (1999). The directionality of verbal probability expressions: effects on decisions, predictions, and probabilistic reasoning. *Organizational Behavior and Human Decision Processes*, 80, 155–190.
- Wallsten, T. S., & Budescu, D. V. (1995). A review of human linguistic probability processing: general principles and empirical evidence. *The Knowledge Engineering Review*, 10, 43–62.
- Wallsten, T. S., Budescu, D., Rapoport, A., Zwick, R., & Forsyth, B. (1986). Measuring the vague meanings of probability terms. *Journal of Experimental Psychology: General*, 115, 348–365.
- Weber, E. U., & Hilton, D. J. (1990). Contextual effects in the interpretation of probability words: perceived base-rate and severity of events. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 781–789.
- Zadeh, L. A. (1965). Fuzzy sets. *Information and Control*, 8, 338–353.