Title
Investigating the Difference between Surprise and Probability Judgments

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 31(31)

ISSN
1069-7977

Authors
Meguire, Phil
Maguire, Rebecca

Publication Date
2009

Peer reviewed

Investigating the Difference between Surprise and Probability Judgments

Phil Maguire (pmaguire@cs.nuim.ie)
Department of Computer Science, NUI Maynooth
Co.Kildare, Ireland

Rebecca Maguire (rebecca.maguire@dbs.ie)
Department of Psychology, Dublin Business School
34/35 South William Street, Dublin 2, Ireland

Abstract

Surprise is often defined in terms of disconfirmed expectations, whereby the surprisingness of an event is thought to be dependent on the degree to which that event contrasts with a more likely, or expected, outcome. We propose that surprise is more accurately modelled as a manifestation of an ongoing sense-making process. Specifically, the level of surprise experienced depends on the extent to which an event necessitates representational updating. This sense-making view predicts that differences in subjective probability and surprise arise because of differences in representational specificity rather than differences between an expectation and an outcome. We describe two experiments which support this hypothesis. The results of Experiment 1 demonstrate that generalised representations can allow subjectively low probability outcomes to be integrated without eliciting high levels of surprise, thus providing an explanation for the difference between the two measures. The results of Experiment 2 reveal that the level of contrast between expectation and outcome is not correlated with the difference between probability and surprise. The implications for models of surprise are discussed.

Keywords: Surprise, probability, likelihood judgments, expectation, representation, reasoning.

Introduction

Surprise is a familiar experience to us all, whether induced by a noise in the dark, or by an unexpected twist in a murder mystery. Due to its pervasiveness, surprise has been a topic of interest for researchers in psychology and its cognate disciplines for quite some time (e.g. Darwin, 1872). This research has shown that, as well as being one of the most basic and universal of human emotions, surprise has many important cognitive ramifications (Fisk, 2002; Meyer, Reisenzein & Schützwohl, 1997; Ortony & Partridge, 1987; Schützwohl, 1998; Teigen & Keren, 2003). For example, a surprising event, as well as giving rise to a 'feeling of surprise' at a subjective and physiological level, usually results in an interruption to ongoing activities and an increased focusing of attention on the event in question (e.g. Schützwohl & Reisenzein, 1999). As such, one hypothesis is that surprise plays a key role in learning and prediction: it interrupts activity to focus one's attention on why the surprising event occurred in the first place, so that a similar event may be predicted and avoided in the future (Darwin, 1872; Meyer et al. 1997).

In spite of the considerable research that has been conducted on the subjective and behavioural ramifications of surprise and the important role it plays in difference contexts, it is not clear how and when a person becomes surprised. In this article, we attempt to shed light on this issue by investigating the factors which cause an event to be perceived as surprising. In particular, we seek to explain the difference between probability and surprise judgments. Why do some unlikely events elicit surprise, while others do not?

Surprise as Probability

A review of the literature reveals that the prevailing definition of surprise relates it directly to expectation (e.g. Meyer et al, 1997; Ortony & Partridge, 1987; Teigen & Keren, 2003). Indeed, this view corresponds to people's own naïve understanding of the phenomenon (e.g. Bartsch & Estes, 1997, found that both children and adults conceptualise surprise in terms of expectation). Theoretically, expectation is formalised in terms of probabilities, where an unexpected outcome is considered to be a low probability event, and vice-versa (Teigen & Keren, 2003). If we relate this to surprise, then low probability events should lead to a feeling of surprise, while high probability events should not.

While there has been some empirical support for this view (e.g. Fisk, 2002; Itti & Baldi, 2006; Reisenzein, 2000), the intuitive relationship between probability and surprise does not always hold. Most of the events that occur in everyday life are quite unlikely based on prior knowledge, yet their occurrence does not always lead to surprise. Teigen and Keren (2003) carried out a number of experiments which illustrated a divergence between probability and surprise. Participants rated both the probability of a certain event and also how surprised they would be if the event were to occur. In one study, for instance, participants were presented with a scenario that described Erik, an athlete who was competing in a 5,000m race. One set of participants were told that, with two laps to go, all the athletes in the race – including Erik – were running together in a large group (multiple alternatives condition). Another set of participants were told that Erik was in second place, lagging behind a lead athlete with the rest of the athletes far behind (single alternative condition). Both groups were then asked to rate the probability of Erik winning the race, and how surprised they would be if he won. While participants in the single alternative condition (where Erik was in second place) correctly rated the probability of Erik winning the race as higher than those in the multiple alternative condition (where all the athletes were in one group), they also rated this possibility as being more surprising.

To explain this effect, Teigen and Keren (2003) proposed the contrast hypothesis of surprise. This theory holds that surprise is governed by the relative probabilities of alternative events, rather than by the absolute probability of the observed outcome. When there are multiple alternatives to an outcome (i.e. when any of the athletes, including Erik, has a chance of winning the race), participants should be less surprised at Erik winning than if there is just one likely alternative outcome (i.e. when only one other athlete – the lead runner – is likely to win). According to Teigen and Keren (2003), this is because the 'contrast' between the observed and the expected outcome is greater in the latter version of the scenario.

Teigen and Keren's (2003) contrast hypothesis has greater scope than the probability-based view, since it explains why many events, which have a low probability of occurring, do not lead to surprise. However, the theory only applies to situations in which explicit expectations are formed. This is a significant limitation as, intuitively, most surprise reactions do not contradict prior expectations (e.g. a brick coming through the window). In order to subsume these alternative forms of surprises into a single comprehensive theory, a more general explanation is required.

Surprise as Sense-making

Kahneman and Miller (1986) originally proposed that surprise reflects a person's success, or more appropriately their failure, to make sense of an event. In line with this view, Maguire, Maguire and Keane (2007) proposed that the experience of surprise reflects a representation updating process. Maintaining a current and valid representation of the environment is of utmost importance to any organism in order to allow it to act appropriately; allowing it to diverge from reality can have serious consequences. Maguire et al.'s integration hypothesis proposes that surprise occurs when a coherent representation 'breaks down' in light of a discrepant stimulus. In such cases it makes sense to focus attention immediately on the event so that appropriate action can be taken as soon as possible

People are constantly updating their representations in very minor ways. For example, a person's attention tends to be directed towards information which is least congruent with their representation: Itti and Baldi (2006) found that 84% of gaze shifts were directed towards locations that were more surprising when participants were shown television and video games. The best strategy for incorporating a discrepant event is to direct additional cognitive resources towards it and to sample additional information from the environment. For that reason, the emotional state of surprise is generally accompanied by physiological arousal, as well as distinctive changes in facial expression, such as eye-widening and the opening of the mouth. Darwin (1872) originally proposed that these changes prepare an organism to react. Susskind et al. (2008) suggested that the facial expression associated with surprise has evolved to enhance the intake of sensory information. They found that participants with wide-open eyes detected peripheral objects more quickly and performed side-to-side eye movements faster. The nasal cavity was also enlarged, enhancing the absorption of odours and allowing participants to take in more air with each breath without exerting any extra effort.

Once the physiological surprise response has subsided, the urge to understand a discrepant event persists. Imagine finding a gorilla in your car. At first you would be taken aback, experiencing the physiological changes associated with a surprise reaction. After calling the zoo to have the gorilla removed, this initial emotional response would subside, yet the urge to reconcile this bizarre event with your representation of reality would persist nonetheless (how on earth did the gorilla get into the car?) We maintain that this form of 'cognitive surprise' is driven by the same conditions which give rise to the more visceral form of surprise, namely the need to maintain a valid representation of reality. If both forms of surprise are manifestations of the same underlying phenomenon, then they should be explained by a single unified theory.

In this article we argue that the integration hypothesis encompasses the predictions of Teigen and Keren's contrast hypothesis as well as accounting for other forms of surprise that do not involve explicit expectations. Importantly, the integration hypothesis also provides a strong theoretical motivation for the phenomenon of surprise, as opposed to explaining it in terms of other measures such as probability.

Experiment 1

If people insisted on understanding the causal factors giving rise to all events in their environment, then every subjectively low probability event would elicit surprise. However, representations involving this level of detail are not required. Some events will be inconsequential to the interests of the individual and thus can be ignored. Other events are simply not amenable to explanation because the causal factors are extremely convoluted. Consequently, much of the information in a representation may be generalised in terms of frequencies rather than in terms of precise explanatory factors. For example, rather than scrupulously monitoring and modelling the atmospheric conditions which give rise to precipitation, most people will simply accept that it rains sporadically. Similarly, in a lottery draw, people will accept that a set of unpredictable random numbers will be drawn, rather than for example, furiously trying to explain why the number 36 happened to emerge. As a result of these generalisations, events can occur which, while recognised as having been relatively unlikely, do not require representational updating. The integration hypothesis therefore predicts that differences in surprise and probability judgments arise because of differences in representational specificity.

In the following experiment we investigated the validity of this premise. Participants were asked to provide judgments for four different representations of a weather system. The descriptions were varied according to specificity and also according to the extent to which they supported the outcome. The aim of the experiment was to investigate whether the specificity of a representation affects the level of surprise experienced for subjectively low probability events.

Method

Participants 84 undergraduate students from NUI Maynooth participated voluntarily in this experiment. All were native English speakers.

Materials The four weather representations generated were general-supportive ('it rains five days a week'), general-unsupportive ('it rains one day a week'), specific-supportive ('a cold front approaching from the west will lead to overcast, unsettled weather over the next few days') and specific-unsupportive ('an approaching area of high pressure will bring clear, sunny conditions over the next few days').

Design We used a two-way repeated-measures model. The independent variables of specificity and support were repeated by participant, with probability and surprise judgments as the dependent variables.

Procedure For all scenarios, participants were asked to provide both surprise and probability judgments for the possibility of rain the following day. Surprise ratings were provided on a 7-point scale (7 being the most surprising), while probability was rated in terms of a percentage (100% reflecting certainty). The order of presentation of the scenarios was randomised between participants, as was the order in which they rated surprise and probability.

Results and Discussion

Some of the participants failed to reason probabilistically. For example, in the scenario "it rains one day a week", the probability of rain on a given day must be 14%. However, some participants provided much higher probability ratings, indicating that they were relying on sense-making rather than on frequency information for their judgments. In other words, they were confusing mathematical probability with surprise. Given that the aim of the experiment was to elucidate the difference between probability and surprise, these participants could not be included in the study. Accordingly, we eliminated the responses of any participant who rated the probability for the abstract-unsupportive scenario as higher than 20%. This removed a total of 23 participants, 12 of whom had rated surprise first and 11 of whom had rated probability first. The average probability rating provided by these participants for the abstract-unsupportive condition was 61%, varying from 30% (4 participants) up to 90% (3 participants). The extent of this logical error indicates that people are prone to relying on representation-fit for making likelihood judgments, even when explicit frequency information is available.

The average probability ratings are provided in Table 1. Both the general-unsupportive and the specific-unsupportive scenarios were rated as similarly improbable (15% and 16% respectively), yet the specific-unsupportive scenario was rated as twice as surprising as the general-unsupportive scenario (5.2 and 2.6 respectively). We conducted repeated measures ANOVAs examining the relationship between specificity, probability and surprise. Probability ratings were not affected by representational specificity: The specificity X supportiveness interaction was not significant, $F(1,58) = .241$, $p = .626$, $MS_e = 179.621$. As expected, there was a significant main effect of supportiveness, $F(1,58) = 2360.673$, $p < .001$, $MS_e = 97.946$, though no significant effect of specificity, $F(1,58) = .1.019$, $p = .317$, $MS_e = 174.693$. The surprise ratings displayed a different pattern of results. In this case, there was a strong interaction between specificity and supportiveness, $F(1,58) = 70.188$, $p < .001$, $MS_e = 1.235$. There was a significant main effect of supportiveness, $F(1,58) = 186.47$, $p < .001$, $MS_e = 1.744$, as well as a significant main effect of specificity, $F(1,58) = 100.97$, $p < .001$, $MS_e = 1.115$, with the general-unsupportive scenario being rated as far less surprising than the specific-unsupportive scenario.

In order to better analyse the effect of specificity on the difference between surprise and probability ratings, we converted both to a single scale and subtracted one from the other, yielding a total difference score. As shown in Table 1, surprise and probability ratings diverged markedly for the general-unsupportive scenario (59%) but were relatively consistent for the other three scenarios (16%, 11% and 15%). We computed a two-way repeated measures ANOVA on the difference values. Again, there was a significant interaction between specificity and supportiveness, $F(1,58) = 83.669$, $p < .001$, $MS_e = 392.886$, with a significant main

Table 1
Mean probability and surprise ratings for Experiment 1

Order of ratings	General-Supportive		General-Unsupportive		Specific-Supportive		Specific-Unsupportive	
	Probability	Surprise	Probability	Surprise	Probability	Surprise	Probability	Surprise
Probability first	77%	1.3	15%	2.6	83%	1.6	17%	5.1
Surprise first	76%	1.6	14%	2.6	76%	1.6	15%	5.2
Mean	77%	1.5	15%	2.6	79%	1.6	16%	5.2
Difference	15.9%		58.7%		10.5%		14.6%	

effect of supportiveness, $F(1,58) = 54.000$, $p < .001$, $MS_e = 514.823$, as well as a significant main effect of specificity, $F(1,58) = 44.111$, $p < .001$, $MS_e = 406.838$. Pairwise comparisons using Bonferroni adjustments revealed that there was a significant difference between the general-unsupportive condition and the other three conditions (all *p*s < .0083) but no other significant differences within these three.

These results illustrate clearly that generalised representations can lead to lower levels of surprise for subjectively low probability outcomes. Although participants acknowledged that it was unlikely to rain the following day in the general-unsupportive scenario (15%), they would not have been very surprised if it did rain (2.6 out of 7). This observation is in line with the integration hypothesis, which predicts that surprise should be lower when low probability events can be easily integrated without requiring representational updating. For example, the occurrence of rain on a particular day does not undermine the representation that it rains one day each week.

The observed pattern of results suggests that differences in surprise and probability can be linked to the specificity of the representations on which such judgments are based. The more specific a representation, the less likely it is to be compatible with low probability events, causing surprise and probability ratings to converge. In idiosyncratic situations which are not amenable to generalisation, people's representations are likely to be over-fitted and less able to accommodate subjectively low probability outcomes. In this case, anything that deviates from expectation will require a fundamental re-evaluation of the representation. For example, Reisenzein (2000) asked participants to rate confidence and surprise for answers on a multiple-choice test and found a very strong correlation between these measures ($r = -0.78$). The likelihood of being correct about a particular multiple-choice question is difficult to generalise. If you are confident that you know the answer to a particular question and are subsequently shown to be incorrect (a subjectively low probability event), then there is a need to update to your beliefs in order to prevent subsequent errors of judgment. Unlike the situation involving the weather, a generalisation in this case is unacceptable.

Integration or Anticipation?

The pattern of results observed in Experiment 1 could potentially be accommodated by the contrast hypothesis. For example, it could feasibly be claimed that the general-unsupportive representation did not contradict any expectation, while the specific-supportive representation contradicted the expectation that there would be clear, sunny weather. Because expectations and anticipatory processes are by definition based on one's current representation of reality, events which contrast with an explicit expectation will necessarily be events which are difficult to accommodate. Typically, the greater the contrast, the more difficult the integration process. As a result, the contrast and integration hypotheses make similar predictions for representations which are specific enough to set up an expectation.

A significant limitation of Teigen and Keren's (2003) contrast hypothesis is that the range of potential outcome events in a given situation can rarely be divided up in terms of a discrete number of competing alternatives. As a result, most instances of surprise cannot be explained in terms of contrast. For example, if you are sitting on the couch watching television and a brick comes through the window, it is difficult to construe this event as being in contrast with the expectation that a brick would not come through the window. People do not usually sit in their living room thinking about bricks, yet they would certainly be surprised if they saw one coming towards them. According to the integration hypothesis, the range of possible outcomes is so great that events are typically evaluated after they have occurred, as part of a sense-making process, rather than being predicted beforehand. The potential for events to be judged as surprising or unsurprising is usually *implicit* to a representation. In other words, people do not always have well-formed expectations about what is going to happen next. What they *do* have is a representation which can be used to make sense of the events that happen to unfold. In other words, reality does the hard work of figuring out what happens next and people do the easier work of trying to make sense of it.

If the integration hypothesis is correct, then it should be possible for events to simultaneously violate expectations *and* be judged as unsurprising, provided some effective way of rationalising those events is established. For example, although it might be surprising to find that a trailing runner wins a race (as in Teigen and Keren's example), it should seem *less* surprising when a convincing explanation is provided (e.g. the leader stumbles). Maguire and Keane (2006) investigated this possibility, creating 16 scenarios which instantiated an explicit expectation and then analysing surprise ratings provided for a set of different outcomes. In the Confirm condition the outcome confirmed the expectation set up by the representation while in the Disconfirm condition the outcome disconfirmed that expectation. In the Disconfirm-Enable condition, the disconfirming outcome was paired with another enabling event which facilitated the overall integration of the conjunctive outcome with the representation. For example, one of the scenarios described Anna setting her radio alarm clock for 7am in preparation for an important job interview the next morning. In the Disconfirm condition, participants were asked to rate surprise for the outcome that the alarm clock failed to ring. In the Disconfirm-Enable condition, they were asked to rate surprise for the outcome that there was a power-cut during the night and the alarm clock failed to ring.

Maguire and Keane (2006) found that the Disconfirm-Enable outcomes were rated as significantly less surprising than those in the Disconfirm condition. In other words, when participants were provided with a reason for why an unexpected event might occur, their surprise was lower than when the disconfirming event was presented on its own. These findings undermine the contrast hypothesis, as they demonstrate that the same unexpected event is not always judged as equally surprising in different contexts. Instead, surprise for an unexpected event is mitigated when a means for rationalising that event is made available.

Maguire and Keane's (2006) findings could potentially be reconciled with the contrast hypothesis. One could make the case that having an explanation for a disconfirming event lowers the perceived level of contrast between that event and the expected outcome because the outcome seems more likely. Tverksy and Kahneman (1983) demonstrated that a conjunction of associated propositions are often rated as more probable than either proposition in isolation. Thus, knowing that an alarm clock failed to ring because of a power failure may reduce the level of contrast with the expectation that it should have rung at the appropriate time. The key difference between the contrast hypothesis and the integration hypothesis centres on whether or not people develop expectations against which subsequent outcomes are contrasted. While the contrast hypothesis posits that expectations are a key factor in the experience of surprise, the integration hypothesis claims that surprise can be modelled based on the outcome event alone. In the following experiment we investigated whether or not surprise ratings are associated with differences in contrast.

Experiment 2

One limitation of Teigen and Keren's (2003) study was that they did not provide any specific measurements of contrast. In each experiment, a single scenario was presented to participants involving a pair of conditions which were assumed to reflect high and low levels of contrast. Although Teigen and Keren reported significant differences, no measure was provided of the overall correlation between contrast and surprise. We addressed this lacuna by deriving levels of contrast for Maguire and Keane's 16 scenarios and comparing them against the corresponding surprise ratings. Teigen and Keren stated that the "surprise associated with an outcome is determined by the relative, rather than absolute probabilities involved" (p. 58). Accordingly, contrast was calculated by obtaining probability ratings for the confirming and disconfirming outcomes and subtracting them.

The contrast hypothesis predicts that the level of contrast should be closely correlated with surprise ratings. On the other hand, the integration hypothesis maintains that surprise is based solely on the ease with which an event can be integrated. As demonstrated in Experiment 1, subjectively low probability outcomes are less easily integrated with representations which are specific enough to instantiate an explicit expectation. In such situations, probability and surprise ratings tend to converge. Accordingly, the integration hypothesis suggests that probability and surprise ratings should be closely matched for the scenarios under investigation.

Method

Participants 120 undergraduate students from NUI Maynooth voluntarily took part in the experiment. All were native English speakers.

Materials The 16 event sequences generated by Maguire et al. (2006) were used, with the three conditions Confirm, Disconfirm and Disconfirm-Enable.

Design The three conditions were counterbalanced across three lists of scenarios. Each participant was given one of these lists which contained the 16 scenarios paired with one of the three endings.

Procedure Participants were randomly assigned to judge either probability or surprise. In the probability condition, the scenario body was followed by the question *"What is the probability that: X ?"*, where X referred to the event, or series of events, corresponding to one of the three possible outcomes. In the surprise condition, the question was *"How surprised would you be if: X ?"*. As before, surprise ratings were provided on a 7-point scale, while probability was rated in terms of a percentage. The scenarios were presented in a different random order to each participant.

Results and Discussion

The average probability ratings were 75.8%, 14.6% and 20.5% for the Confirm, Disconfirm and Disconfirm-Enable conditions respectively and the average surprise ratings were 1.9, 4.9 and 4.5. The Disconfirm-Enable condition was rated as more probable than the Disconfirm condition for 13 of the 16 scenarios. There was a significant difference between these conditions, indicating a robust conjunction fallacy effect, $F_1(1,15) = 5.980$, $p = .027$, $MS_e = 47.165$, (see Tversky & Kahneman, 1983). The average level of contrast for the Disconfirm scenarios (61.2%) was significantly greater than the average level of contrast for the Disconfirm-Enable condition (55.3%), $F_1(1,15) = 5.872$, $p = .029$, $MS_e = 48.031$. This observation lends support to the idea that Maguire and Keane's (2006) findings can be explained in terms of contrast.

We computed the degree of contrast for both conditions by subtracting the probability ratings from those of the corresponding confirming scenarios. There was no significant correlation between contrast and the surprise ratings, $r = .113$, $p = .537$. However, as predicted by the integration hypothesis, there was a significant correlation between the probability and surprise ratings, $r = -.418$, $p = .017$. All but one of the 16 scenarios displayed the same direction of difference between the Disconfirm and Disconfirm-Enable conditions for both surprise and probability ratings, indicating a close relationship.

These results indicate that differences in contrast are not associated with differences in surprise: low probability events can be just as surprising when a scenario does not support a clear expectation. For example, the lowest probability ratings for any confirming outcome were provided for the scenario where Sarah calls to her parents' house and knocks on the front door. Because either one of her parents must open the door, the probability of the most likely outcome cannot exceed 50% (assuming no bias towards either parent). Although this substantially lowers the potential contrast with any alternative outcome, it does not necessarily lower the surprise for a low probability event: a stranger opening the door is just as surprising, regardless of the probability of the most likely possible alternative.

One might claim that the expectation in this case is that either one of Sarah's parents will answer the door and that the level of contrast should be computed on this basis. In order to make this argument, one would have to concede that an expectation can maintain some degree of ambiguity. However, this is exactly what our proposal of a generalised representation is intended to reflect: people allow a certain level of flexibility in their representations of the world so that a broad range of potential outcomes can be accommodated. Because they do not seek to model every last detail, people's representations are primarily used for making sense of rather than for predicting events.

Conclusion

Experiment 1 demonstrated that differences in representational specificity lead to differences between surprise and probability judgments. Experiment 2 demonstrated that the level of contrast between an expectation and an outcome is not correlated with the difference between surprise and probability ratings. These results provide converging evidence in support of the integration hypothesis while undermining the contrast hypothesis. Because expectation-supporting representations tend to be specific, Teigen and Keren's experimental conditions varied considerably in specificity. Thus, it seems likely that the effect observed by Teigen and Keren (2003) was most likely due to differences in representational specificity rather than differences in contrast per se.

The fact that surprise is a relatively infrequent phenomenon demonstrates that people maintain extremely accurate representations that adapt readily to unfolding events. However, the assumption of the contrast hypothesis that events are continuously predicted by these representations trivialises the complexity of the environment and overstates the level of detail that can be represented. Because real world events are often so unpredictable, generating explicit expectations is seldom practical or even feasible. An unfortunate individual adopting the expectation-making strategy would live their life in a state of perpetual surprise.

Accordingly, we have proposed that people maintain representations which, although not sufficiently detailed to make explicit predictions, can be relied on to make sense of events. A generalised representation allows much of the environment's variability to be captured in a succinct manner, while ignoring irrelevant details. Because low probability events can be congruent with a generalised representation, they do not necessarily require representational updating and hence do not always elicit surprise. In conclusion, we have provided converging evidence that differences between probability and surprise arise not because of contrasts between outcomes and expectations, but because representations can be generalised to facilitate the integration of subjectively low probability events.

References

Bartsch, K. & Estes, D. (1997). Children's and adults' everyday talk about surprise, *British Journal of Developmental Psychology, 15*, 461-475.

Darwin, C. (1872). *The Expression of Emotion in Man and Animals.* London: J. Murray.

Fisk, J. E. (2002). Judgments under uncertainty: Representativeness or potential surprise? *British Journal of Psychology, 93,* 431-449.

Itti, L. & Baldi, P. (2006). Bayesian surprise attracts human attention. *Advances in Neural Information Processing Systems, 19,* 1-8.

Kahneman D, & Miller DT. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review, 93*, 136-53.

Maguire, R. & Keane, M.T. (2006). Surprise: disconfirmed expectations or representation-fit? *Proceedings of the Twenty-Eight Annual Conference of the Cognitive Science Society*.

Maguire, R., Maguire, P. & Keane, M. T. (2007). Understanding surprise: Can less likely events be less surprising? *Proceedings of the Second European Cognitive Science Conference*.

Meyer, W.U., Reisenzein, R. & Schützwohl, A. (1997). Towards a process analysis of emotions: The case of surprise. *Motivation and Emotion, 21(3),* 251-274.

Ortony, A., & Partridge, D. (1987). Surprisingness and expectation failure: What's the difference? In *Proceedings of the tenth international joint conference on artificial intelligence.* Milan, Italy.

Reisenzein, R. (2000). Exploring the strength of association between components of emotion syndromes: the case of surprise. *Cognition and Emotion, 10,* 241-277.

Schützwohl, A. & Reisenzein, R. (1999). Children's and adults' reactions to a schema-discrepant event: A developmental analysis of surprise. *International Journal of Behavioral Development, 23,* 37-62.

Susskind, M. J., Lee, D. H., Cusi, A., Feiman, R., Grabski, W. & Anderson, A. (2008). Expressing fear enhances sensory acquisition. *Nature Neuroscience*, *11*, 843-850.

Teigen, K.H., & Keren, G. (2003). Surprises: Low probabilities or high contrasts? *Cognition, 87*, 55-71.

Tversky, A., & Kahneman, D. (1983). Extensional vs. intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 293-315.