

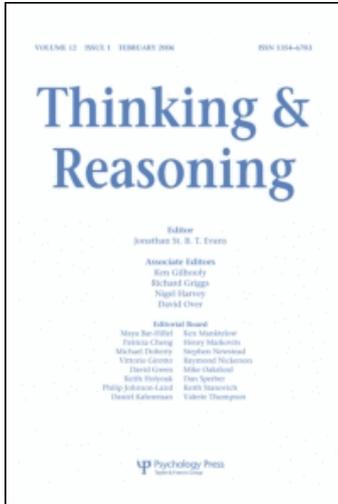
This article was downloaded by: [Princeton University]

On: 15 November 2010

Access details: Access Details: [subscription number 917269152]

Publisher Psychology Press

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Thinking & Reasoning

Publication details, including instructions for authors and subscription information:

<http://www.informaworld.com/smpp/title~content=t713685607>

Updating beliefs in light of uncertain evidence: Descriptive assessment of Jeffrey's rule

Jiaying Zhao^a; Daniel Osherson^a

^a Princeton University, NJ, USA

First published on: 13 October 2010

To cite this Article Zhao, Jiaying and Osherson, Daniel(2010) 'Updating beliefs in light of uncertain evidence: Descriptive assessment of Jeffrey's rule', *Thinking & Reasoning*, First published on: 13 October 2010 (iFirst)

To link to this Article: DOI: 10.1080/13546783.2010.521695

URL: <http://dx.doi.org/10.1080/13546783.2010.521695>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf>

This article may be used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

Updating beliefs in light of uncertain evidence: Descriptive assessment of Jeffrey's rule

Jiaying Zhao and Daniel Osherson
Princeton University, NJ, USA

Jeffrey (1983) proposed a generalisation of conditioning as a means of updating probability distributions when new evidence drives no event to certainty. His rule requires the stability of certain conditional probabilities through time. We tested this assumption (“invariance”) from the psychological point of view. In Experiment 1 participants offered probability estimates for events in Jeffrey’s candlelight example. Two further scenarios were investigated in Experiment 2, one in which invariance seems justified, the other in which it does not. Results were in rough conformity to Jeffrey’s (1983) principle.

Keywords: Belief updating; Jeffrey’s rule; Reasoning.

Consider an idealised agent whose beliefs are represented by a (subjective) probability distribution Pr_1 over an outcome space Ω . Let the event $B \subseteq \Omega$ be such that $Pr_1(B) > 0$ and suppose that experience intervenes to convince the agent that B is certainly true. What probability distribution Pr_2 should represent the agent’s new state of belief? The Bayesian answer (Hacking, 2001, Ch. 15) identifies Pr_2 with the result of conditioning Pr_1 on B , that is:

(1) SIMPLE UPDATING: For all events $A \subseteq \Omega$, $Pr_2(A) = Pr_1(A|B)$.

It is easy to verify that simple updating defines a genuine probability distribution Pr_2 , and that $Pr_2(B) = 1$.

Correspondence should be address to Jiaying Zhao, Department of Psychology, Princeton University, Green Hall, Princeton, NJ 08540, USA. E-mail: jiayingz@princeton.edu

We thank Nick Chater, David Over, Mike Oaksford, Steven Sloman, and an anonymous reviewer for very helpful comments. The research reported here was facilitated by a grant from the Henry Luce Foundation.

To illustrate, suppose you think the probability of the parade being cancelled assuming that it rains is .6. That is, $Pr_1(\text{cancelled} \mid \text{rain}) = .6$. You subsequently learn that, in fact, it's raining. Then (1) recommends that your new distribution Pr_2 result from conditionalising Pr_1 on rain; in particular, $Pr_2(\text{cancelled}) = Pr_1(\text{cancelled} \mid \text{rain}) = .6$.

Much can be said in favour of simple updating from the normative perspective. For example, it follows from compelling axioms on belief change (Gärdenfors, 1988, §5.2), and its violation exposes the agent to sure-loss betting contracts (Harman, 1999, §4.12). Simple updating has also been examined from the psychological perspective with focus on the use of Bayes' Theorem to compute conditional probability (see Oaksford & Chater, 2007; Stanovich, 2010, Ch. 3).

Simple updating is not always suited to represent the impact of new information. In particular, Jeffrey (1983, §11.1) notes that the passage of experience need not raise the probability of any event to one. He gives the example of examining cloth by faint candlelight. The cloth's potential colours might correspond to different events over the sample space Ω but none may become certain as a result of the examination. Nor is it feasible to augment Ω to include visual sensations, with the idea of setting one of them to unity. Such sensations are too difficult to express and individuate. Instead, says Jeffrey (pp. 165–166), “the best we can do is to describe, not the quality of the visual experience itself, but rather its effects on the observer,” for example, that the probability of blue has shifted to .75 from its original value.

To fill in the rest of Pr_2 after experience has set the value of $Pr_2(B)$, Jeffrey relies on the law of total probability. Let $G \subseteq \Omega$ be given, and suppose that $0 < Pr_2(B) < 1$. Then:

$$(2) \quad Pr_2(G) = Pr_2(G \mid B)Pr_2(B) + Pr_2(G \mid \overline{B})Pr_2(\overline{B}).$$

That is, $Pr_2(G)$ is the weighted sum of $Pr_2(G \mid B)$ and $Pr_2(G \mid \overline{B})$ where the weights are the probabilities according to Pr_2 of the respective conditioning events, B and \overline{B} . (The overbar represents negation.) If experience has not influenced the conditional probability of G given B nor that of G given \overline{B} then *invariance* is said to hold (Jeffrey, 2004, §3.2).¹ That is:

$$(3) \quad \begin{aligned} \text{INVARIANCE FOR } G, B: Pr_2(G \mid B) \\ = Pr_1(G \mid B) \text{ and } Pr_2(G \mid \overline{B}) = Pr_1(G \mid \overline{B}) \end{aligned}$$

¹Many authors, including Oaksford and Chater (2007) and Over and Hadjichristidis (2009) use the term “rigidity” instead of “invariance”.

Substituting (3) into (2) yields:

$$(4) \text{ GENERALISED UPDATING: } Pr_2(G) \\ = Pr_1(G|B)Pr_2(B) + Pr_1(G|\bar{B})Pr_2(\bar{B}).$$

Equation (4) is known as “Jeffrey’s rule”. It shows how change in the probability of B is propagated to G , without experience directly affecting G .²

Returning to the earlier example, suppose you learn not that it is raining for sure, but that the probability that it is raining is .8. That is, $Pr_2(\text{rain}) = .8$. You still believe (invariance) that the conditional probability of cancellation given rain is the same as before, i.e., $Pr_2(\text{cancelled} | \text{rain}) = Pr_1(\text{cancelled} | \text{rain}) = .6$. And you also believe (invariance again) that the probability of cancellation in the absence of rain is the same as before, say .1, so $Pr_2(\text{cancelled} | \bar{\text{rain}}) = Pr_1(\text{cancelled} | \bar{\text{rain}}) = .1$. Then Jeffrey’s rule (4) implies that your new probability for cancellation is given by:

$$Pr_2(\text{cancelled}) = Pr_1(\text{cancelled} | \text{rain})Pr_2(\text{rain}) + Pr_1(\text{cancelled} | \bar{\text{rain}})Pr_2(\bar{\text{rain}}) \\ = (.6 \times .8) + (.1 \times .2) = .5.$$

Assuming invariance, it can be demonstrated that (4) defines a genuine probability distribution Pr_2 , and that $Pr_2(B)$ is set to its new value (namely, in the example above, .8). Also, in the special case of $Pr_2(B) = 1$, Pr_2 agrees with the result of applying conditionalisation via rule (1). Moreover, Williams (1980) proved that Pr_2 as given by (4) is the closest distribution to Pr_2 that yields the new probability of B , where “closeness” is measured by cross-entropy with respect to Pr_1 . The normative status of Jeffrey’s rule has nonetheless been questioned because successive uses produce distinct distributions depending on the order in which events are considered (Döring, 1999). In our view, such doubts disappear on closer inspection of the evidential weight of probability judgements (Osherson, 2002; Wagner, 2002). It is important to observe that invariance (3) is not normatively justified in every situation. Sometimes the conditional probabilities ought to be shifted by experience. For example, let G represent vigorous growth of a potted plant, and let B represent the decision to place it in the bedroom. Then noticing sunshine in the bedroom would increase not just the probability of B but also the probability of G given B (non-invariance). In contrast, an

²It is straightforward to extend (4) to finer partitions, in place of the binary partition B, \bar{B} . We stick to the binary case in this paper.

impression of emptiness in the bedroom may increase the probability of B without a change in the probability of G given B , yielding invariance.

To decide whether invariance is warranted in a given situation, we rely on an observation due to Pearl (1988, §2.3.3). (A mathematical treatment is provided by Chan & Darwiche, 2005.) Given the experience e that intervenes between times 1 and 2, we expect invariance to hold if at time 1, G is conditionally independent of e given B , that is:

$$(5) \quad \text{Conditional independence of } G \text{ from } e \text{ given } B : Pr_1(G | B, e) \\ = Pr_1(G | B).$$

To see that invariance depends on (5), observe that $Pr_2(G | B) = Pr_1(G | B, e)$ since e is what transpires between times 1 and 2, and the latter term equals $Pr_1(G | B)$ if and only if (5) holds—in which case $Pr_2(G | B) = Pr_1(G | B)$ by the transitivity of equality. Thus, invariance and the conditional independence expressed by (5) are equivalent. As Pearl notes, this provides insight into Jeffrey's rule inasmuch as conditional independence often has a qualitative character, discovered through reasoning about causal relations. In what follows we rely on (5) to decide whether invariance is normatively sustained. In the example above, G is independent of the impression of emptiness in the bedroom given the decision B to place the plant in the bedroom (growth does not depend on such impressions once the plant is installed in the bedroom). In contrast, G is not independent of sunshine in the bedroom given B inasmuch as sunshine promotes growth in any room.

We may apply Pearl's criterion to an illuminating example raised by Oaksford and Chater (2007) in their comprehensive theory of reasoning. Suppose that for you, $Pr(\text{the car starts} | \text{the key is turned}) = x$. If you now learn that the car did not start, should invariance hold? That is, should the foregoing conditional probability still be x ? Since the car starting seems not to be conditionally independent of the car failing to start, given that the key is turned, there appears to be no normative case for invariance. Indeed, as predicted by Oaksford and Chater, reasoners often reduce their estimate of the conditional probability in this example, upon learning that the car did not start.

Jeffrey's rule is increasingly discussed in the context of psychological theories of reasoning, for example, in Hahn and Oaksford's (2007) Bayesian analysis of reasoning fallacies like circularity, and in Oaksford and Chater's (2007) theory of inference. It may also be relevant to the interpretation of conditionals (Oberauer & Wilhelm, 2003) insofar as conditionals are understood as expressing conditional probability (Adams, 1991). As suggested by Over and Hadjichristidis (2009), it may thus be useful to directly address the descriptive adequacy of Jeffrey's rule. Such is the focus of the present paper. We investigate whether invariance is honoured in judgement when it is

mandated normatively. In the first experiment undergraduates offer probabilities for events in Jeffrey's candle example. Two further scenarios are then investigated, one in which invariance seems justified, the other in which it does not. The results show modest but non-negligible deviations from invariance where it seems justified normatively.

EXPERIMENT 1

Participants

A total of 70 undergraduates from Princeton University participated in exchange for partial course credit (45 female, mean age 19.8 yrs, $SD = 2.1$).

Materials

We simulated Jeffrey's candle example by having participants examine coloured paper cards with a dim flashlight. There were 12 blue cards and 38 purple cards. Each card was marked with either a hippopotamus or a giraffe on one side. Table 1 summarises the respective frequencies of cards. We chose *giraffe* and *blue* as the categories G and B evoked in the Introduction.

Procedure

A total of 40 participants performed the *experimental* condition and 30 performed the *control* condition. Each participant was interviewed individually. The purpose of the control condition was to assess the impact of being asked to evaluate the same probabilities a second time. In both conditions each card was first shown to the participant (both sides), and the cards were shuffled. The experimenter then turned away from the participant, drew one card from the shuffled deck, and put it in her pocket. The draw appeared to be random but in fact was guaranteed across participants to deliver equal numbers of blue and purple cards (for statistical purposes). The participant was informed that the card was randomly chosen, and then answered the following questions about the card, via computer interface (the order was randomised for each participant):

TABLE 1
Number of cards used in Experiment 1

	<i>Blue cards</i>	<i>Purple cards</i>
Giraffe	8	10
Hippo	4	28

PROBABILITY QUESTIONS:

$Pr(G)$	What's the probability that there is a giraffe on the card?
$Pr(B)$	What's the probability that the card is blue?
$Pr(G B)$	What's the probability that there is a giraffe on the card assuming that the card is blue?
$Pr(G \overline{B})$	What's the probability that there is a giraffe on the card assuming that the card is purple?
$Pr(B G)$	What's the probability that the card is blue assuming that there is a giraffe on the card?
$Pr(B \overline{G})$	What's the probability that the card is blue assuming that there is a hippo on the card?

The estimates $Pr(B | G)$ and $Pr(B | \overline{G})$ served as a contrast with $Pr(G | B)$ and $Pr(G | \overline{B})$. Given our procedure (see below), the former estimates were not expected to be invariant across the flashlight experience, whereas the latter were.

The participant was then informed that s/he would briefly see the card under dim light. It was explained that the animal side of the card would not be revealed and the participant would only see the colour. The experimenter then turned off the lights in the room, moved the card from her pocket to the table, and illuminated it via dim flashlight for about one second. The card was then returned to the experimenter's pocket, and the participant answered the same set of questions shown above (in a different random order). Since participants had to give their estimates twice to the same questions, we informed them that they were free to provide the same estimate or a different estimate the second time around.

In the control condition the procedure was the same except that the flashlight was applied to the chosen card immediately after its draw. Participants in the control condition thus answered the questions shown above just once, after briefly seeing the card under the dim light.

Results

Average responses. We separately analysed results for participants who were exposed to blue cards (Blue group) and those exposed to purple cards (Purple group). In the experimental condition we use Pr_1 to refer to probability estimates before the light experience and Pr_2 for estimates after the experience. Average estimates are shown in Table 2. For example, the fourth row shows the average pre-flashlight estimates of the experimental group assigned a purple card at the outset.

Control versus experimental conditions. Since experimental participants had to provide estimates twice to the same questions, we compared their second estimates to those of the control group. Independent-sample *t*-tests for each of the six questions (after exposure to the flashlight for both

TABLE 2
Objective probabilities in the cards and average subjective estimates in Experiment 1

	$Pr(G)$	$Pr(B)$	$Pr(G B)$	$Pr(G \bar{B})$	$Pr(B G)$	$Pr(B \bar{G})$
Objective	0.36	0.24	0.67	0.26	0.44	0.13
Blue Pr_1	0.38 (0.08)	0.35 (0.16)	0.52 (0.24)	0.35 (0.18)	0.54 (0.24)	0.34 (0.25)
Blue Pr_2	0.48 (0.17)	0.74 (0.26)	0.54 (0.23)	0.30 (0.20)	0.59 (0.23)	0.51 (0.34)
Purple Pr_1	0.35 (0.11)	0.38 (0.16)	0.61 (0.22)	0.26 (0.12)	0.53 (0.21)	0.32 (0.21)
Purple Pr_2	0.33 (0.17)	0.08 (0.10)	0.46 (0.29)	0.28 (0.13)	0.30 (0.28)	0.27 (0.28)
Blue control	0.44 (0.13)	0.71 (0.16)	0.58 (0.18)	0.33 (0.20)	0.57 (0.15)	0.35 (0.20)
Purple control	0.31 (0.16)	0.10 (0.10)	0.58 (0.22)	0.32 (0.24)	0.41 (0.17)	0.27 (0.23)

Standard deviations in parentheses. The row headed Blue Pr_1 gives the averages for participants assigned a blue card, prior to seeing its colour under dim light. The row headed Blue Pr_2 gives the averages for the same group after the light. The row headed Blue control shows averages for the Blue control group after the light (this group made no estimates prior to the light). Purple is interpreted similarly.

groups) revealed no reliable differences between Pr_2 and the control estimates. Thus experimental participants seem not to have been influenced by having to evaluate the same probabilities twice.

Analysis of the Experimental Blue group. We report the Experimental Blue and Experimental Purple participants separately, starting with Blue. As a manipulation check we first determined whether $Pr(B)$ increased after participants saw the blue card under dim light. As expected, $Pr_2(B)$ was reliably larger than $Pr_1(B)$: paired $t(19) = 5.8$, $p < .01$, Wilcoxon $p < .01$. (Respective medians were .30 and .80.)

We now consider invariance, and note that it is normatively endorsed by the conditional independence test discussed in the introduction (since the flashlight does not affect the frequencies of the four types of card). To see whether invariance holds as described in (3), we compared $Pr_1(G | B)$ to $Pr_2(G | B)$ via paired t -test and found no reliable difference, $t(19) = 0.7$, $p = .52$. Nonetheless, of the 20 Blue participants, 11 offered a different $Pr_2(G | B)$ estimate from $Pr_1(G | B)$. Of those 11, 8 participants gave a higher $Pr_2(G | B)$ estimate than $Pr_1(G | B)$. $Pr_1(G | \bar{B})$ was likewise close to $Pr_2(G | \bar{B})$, paired $t(19) = 0.43$, $p = .39$. But 8 out of 20 participants gave a different $Pr_2(G | \bar{B})$ estimate from $Pr_1(G | \bar{B})$ and of those 8 only 1 gave a higher $Pr_2(G | \bar{B})$ estimate than $Pr_1(G | \bar{B})$.

To more precisely quantify violation of invariance, for each participant we calculated the absolute movement between his/her two estimates as a percentage of the original estimate, as follows:

$$(6) \text{ Invariance violation} = \frac{|Pr_2(G | B) - Pr_1(G | B)|}{Pr_1(G | B)}$$

We compared these invariance violations to the movement of the converse probability (i.e., blue given giraffe), computed via:

$$(7) \quad \text{Converse movement} = \frac{|Pr_2(B|G) - Pr_1(B|G)|}{Pr_1(B|G)}$$

Following the normative analysis based on conditional independence (see the introduction), invariance is not expected for converse movement. This is because the flashlight does provide additional information about the colour of a card even if its animal is known. Consistent with the normative difference, the means for invariance violations and converse movements were 23.5% and 82.3%, respectively. This difference is reliable by paired *t*-test, $t(19) = 2.3$, $p < .05$, and Wilcoxon test, $p < .05$. Regarding converse movement, of the 20 Blue participants, 15 offered a different $Pr_2(B|G)$ estimate from $Pr_1(B|G)$ (binomial $p < .05$). We also computed the converse movement for $Pr(B|\bar{G})$ which normatively need not be invariant. Indeed, the mean movement was 97.2%.

For each participant we also computed invariance violation for $Pr(G|\bar{B})$ via the following counterpart to (6):

$$(8) \quad \text{Invariance violation for } \bar{B} = \frac{|Pr_2(G|\bar{B}) - Pr_1(G|\bar{B})|}{Pr_1(G|\bar{B})}$$

The mean invariance violation of $Pr(G|\bar{B})$ was 15.7%, not reliably different from the 23.5% violation of $Pr(G|B)$ reported above, paired $t(19) = 0.7$, $p = .47$.

From the results above invariance seems to hold, albeit somewhat approximately. We therefore asked about its use in updating the probability of *G*. Specifically, for each participant, we computed the value of $Pr(G)$ via the law of total probability (2)—namely: $Pr_2(G) = Pr_2(G|B)Pr_2(B) + Pr_2(G|\bar{B})Pr_2(\bar{B})$ —relying on the participant's estimates for the quantities at the right of the equality. We call this value the total probability of *G*, or $Pr_{\text{total}}(G)$ for short. Likewise, for each participant we computed the value of $Pr(G)$ via Jeffrey's rule (4)—namely: $Pr_2(G) = Pr_1(G|B)Pr_2(B) + Pr_1(G|\bar{B})Pr_2(\bar{B})$. We call this value the Jeffrey probability of *G*, or $Pr_{\text{Jeff}}(G)$ for short. The latter estimates were compared to the participant's direct evaluation of $Pr_2(G)$ via absolute difference:

$$(9) \quad \text{total error} = |Pr_2(G) - Pr_{\text{total}}(G)|$$

$$\text{Jeffrey error} = |Pr_2(G) - Pr_{\text{Jeff}}(G)|$$

The means for total and Jeffrey error were .09 and .11, respectively, not reliably different via paired t -test, $t(19) = 1.2$, $p = .24$, or Wilcoxon test, $p = .34$.

Jeffrey's rule thus seems to yield values that are close to the subjective estimates of $Pr_2(G)$. At the same time, participants appeared to focus more on the first Jeffrey term $Pr_1(G|B)Pr_2(B)$ in (4) than the second $Pr_1(G|\bar{B})Pr_2(\bar{B})$. This was revealed by a regression of $Pr_2(G)$ on the two terms. Only the first proved significant: $t = 2.1$, $p < .05$, coefficient = 0.40 versus $t = -0.5$, $p = .36$, coefficient = -0.44 .³ As for the law of total probability, participants likewise seemed to focus more on the first term $Pr_2(G|B)Pr_2(B)$ in (2) than the second $Pr_2(G|\bar{B})Pr_2(\bar{B})$: $t = 2.4$, $p < .05$, coefficient = 0.47 versus $t = -0.2$, $p = .83$, coefficient = -0.11 .

Analysis of the Experimental Purple group. We first performed a manipulation check to determine whether $Pr(B)$ decreased after participants saw the purple card under dim light. As expected, $Pr_2(B)$ was reliably less than $Pr_1(B)$, paired $t(19) = 6.7$, $p < .01$, Wilcoxon $p < .01$. (Respective medians were .38 and .05.)

To see whether invariance holds as described in (3), we compared $Pr_1(G|B)$ to $Pr_2(G|B)$ via paired t -test and found no reliable difference, $t(19) = 1.8$, $p = .08$. Of the 20 Purple participants, 11 offered a different $Pr_2(G|B)$ estimate from $Pr_1(G|B)$. Of those 11, 7 participants gave a lower $Pr_2(G|B)$ estimate than $Pr_1(G|B)$. $Pr_1(G|\bar{B})$ was also found to be close to $Pr_2(G|\bar{B})$, paired $t(19) = 1.1$, $p = .29$. Out of 20 participants, 10 gave different estimates for $Pr_2(G|\bar{B})$ and $Pr_1(G|\bar{B})$, with 4 of the 10 giving a lower estimate of $Pr_2(G|\bar{B})$ compared to $Pr_1(G|\bar{B})$.

Just as for the Blue group, we computed invariance violations of $Pr(G|B)$ using (6) and converse movement of $Pr(B|G)$ using (7). The means for invariance violations and converse movements were 42.5% and 63.7%, respectively. This difference is reliable by paired t -test, $t(19) = 2.1$, $p < .05$, and Wilcoxon test, $p < .05$. The mean invariance violation for $Pr(G|\bar{B})$ via (8) was 29.2%, not reliably smaller than the 42.5% violation of $Pr(G|B)$, paired $t(19) = 0.8$, $p = .43$. Regarding converse movement, of the 20 Purple participants, 17 offered a different $Pr_2(B|G)$ estimate from $Pr_1(B|G)$, binomial $p < .01$.

Once again invariance seems to hold, although only approximately. So for each participant we computed the value of $Pr(G)$ via the law of total probability (2), again denoting this value by $Pr_{\text{total}}(G)$. Likewise, for each participant we computed the value of $Pr(G)$ via Jeffrey's rule (4), denoting this value by $Pr_{\text{Jeff}}(G)$. The latter estimates were compared to $Pr_2(G)$ via the absolute differences shown in (9). The means for total and Jeffrey error were

³We thank Steve Sloman for suggesting this analysis.

.05 and .10, respectively, not reliably different via paired t -test, $t(19) = 1.4$, $p = .22$, or Wilcoxon test, $p = .25$.

In these Purple participants (just as for Blue), Jeffrey's rule yields values that are close to the subjective estimates of $Pr_2(G)$. Regressing $Pr_2(G)$ on the two Jeffrey terms reveals more focus on the second compared to the first, in contrast to the Blue group. The second term, $Pr_1(G | \bar{B})Pr_2(\bar{B})$, yields $t = 2.3$, $p < .05$, coefficient = 0.8, whereas the first term, $Pr_1(G | B)Pr_2(B)$, yields $t = 0.1$, $p = .89$, coefficient = 0.09. Likewise for the law of total probability: participants seemed to focus more on the second term $Pr_2(G | \bar{B})Pr_2(\bar{B})$ in (2) than the first $Pr_2(G | B)Pr_2(B)$: $t = 5.7$, $p < .01$, coefficient = 1.11 versus $t = 1.4$, $p = .19$, coefficient = 0.67.

Combining Blue and Purple groups. For an overall assessment of invariance violation and converse movement we collapsed the Blue and Purple groups. Of the 40 experimental participants, 2 showed more than 100% invariance violation—in the sense of (6)—and another three showed more than 100% converse movement—in the sense of (7). Excluding these five outliers, we examined the distributions of percent movements in $Pr(G | B)$ and in $Pr(B | G)$; see Figure 1. The average invariance violation was 0.18 and the average converse movement was 0.45. Thus the light affected $Pr(B | G)$ more than it did $Pr(G | B)$, paired $t(34) = 3.2$, $p < .01$. This contrast is consistent with the normative case for invariance in the second case but not the first.

Discussion of Experiment 1

In the procedure of Experiment 1 respect for the invariance principle (3) seems normatively mandated, inasmuch as experience with the light

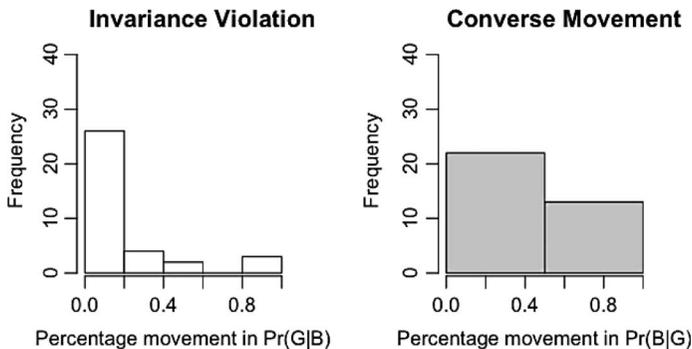


Figure 1. Distributions of invariance violation ($N = 38$) and converse movement ($N = 37$) in Experiment 1 (combining Blue and Purple groups) for shifts within 100%.

provides no further information about G once it is granted that the card is blue. In other words, $Pr_2(G | B) = Pr_1(G | B, I) = Pr_1(G | B)$, where I is the experience provided by the light (as discussed in the Introduction). A majority of participants, in contrast, gave different estimates for $Pr_2(G | B)$ compared to $Pr_1(G | B)$ after gaining new information about colour via the light. Indeed, collapsing over the 40 participants (blue and purple groups combined), 22 were non-invariant for $Pr(G | B)$ and 18 for $Pr(G | \bar{B})$.

On the other hand, the violations of invariance were relatively small in extent. As a percentage of $Pr_1(G | B)$, the absolute difference between $Pr_2(G | B)$ and $Pr_1(G | B)$ averaged 33.0% across all 40 participants. This difference is not trivial but nonetheless reliably smaller than the absolute percentage difference (averaging 73.0%) for the converse probabilities $Pr_2(B | G)$ and $Pr_1(B | G)$. Normatively, invariance is not expected with respect to $Pr(B | G)$. When used to estimate $Pr_2(G)$ via the law of total probability, we saw that $Pr_1(G | B)$ could be substituted for $Pr_2(G | B)$ with little loss of accuracy—total versus Jeffrey error, as in (9). This provides another indication of the relative modesty of invariance violations in Experiment 1. Similar remarks apply to $Pr_1(G | \bar{B})$ compared to $Pr_2(G | \bar{B})$.

In comparison with the mitigated respect for invariance in Experiment 1, simple conditioning (1) was honoured near perfectly in a separate procedure that we carried out with 20 additional participants. Specifically, 10 participants viewed all the cards, just as before, then estimated $Pr_1(G | B)$. Next a random card was drawn (contrived to be blue), and its colour was revealed by viewing it in daylight (so that the blueness becomes certain). Finally the participant was invited to estimate $Pr_2(G)$. The same procedure was applied to another 10 participants who viewed purple. For the Blue group simple conditioning requires $Pr_2(G) = Pr_1(G | B)$ whereas it requires $Pr_2(G) = Pr_1(G | \bar{B})$ for the Purple group. In fact these equations held almost perfectly. For Blue participants the mean of the $Pr_2(G)$ estimates was 0.59 ($SD = 0.17$) while it was 0.56 ($SD = 0.16$) for $Pr_1(G | B)$ [$t(9) = 1.2$, $p > .05$]. Likewise, for Purple participants the means of $Pr_2(G)$ and $Pr_1(G | \bar{B})$ were 0.27 ($SD = 0.13$) and 0.24 ($SD = 0.13$), respectively, $t(9) = 0.8$, $p > .05$.

EXPERIMENT 2

Candlelight in Jeffrey's example produces impressions of colour that are difficult to describe; the same is likely true in the flashlight procedure of Experiment 1. In contrast, our second experiment was designed so that the passage of experience yields events that are more easily characterised.

We asked a new group of participants to estimate probabilities for events in two different scenarios, one in which invariance is justified, the other in which it is not. The scenarios involved a lottery and the ultimatum game, respectively.

Participants

A total of 100 undergraduates from Princeton University participated in exchange for partial course credit (58 female, mean age 19.5 yrs, $SD = 1.1$). None had served in Experiment 1.

Materials and procedure

A total of 50 participants served in the *experimental* condition, and another 50 in the *control* condition. As in Experiment 1, the purpose of the control condition was to assess the impact of being asked to evaluate the same probabilities a second time. Each participant in both conditions was presented with the lottery and the ultimatum game scenarios (the order was counter-balanced).

Lottery scenario

In this scenario we substitute C (*buying a car*) for G , and W (*winning the lottery*) for B . The following description was presented on the computer screen for each participant:

Imagine that a randomly chosen adult (call him Mr X) in New Jersey has just purchased the *Jersey Cash 5* lottery for this week. In this lottery, there are 5 numbers to be drawn, each from 1 to 40. Each number is drawn from the bowl and then put aside. The lottery jackpot is \$240,000 which will be shared by players who have all 5 winning numbers (the order of the numbers doesn't matter). The numbers on Mr X's lottery ticket are 12 17 24 32 39.

In the experimental condition the participant answered the following questions before any lottery numbers were drawn (the order was randomised for each participant):

LOTTERY PROBABILITY QUESTIONS:

- $Pr(C)$ What's the probability that Mr X will buy a new car in the next two years?
 $Pr(W)$ What's the probability that Mr X will win the jackpot?
 $Pr(C | W)$ What's the probability that Mr X will buy a new car in the next two years assuming that he wins the jackpot?
 $Pr(C | \overline{W})$ What's the probability that Mr X will buy a new car in the next two years assuming that he does NOT win the jackpot?

The participant was then presented with the following additional information.

It's the night of the lottery, and the numbers are being drawn. Mr X becomes excited because the first four draws are 32, 12, 24, and 17. In other words, the first four numbers drawn match the numbers on his ticket.

The participant then answered the same set of questions shown above (in a different random order) prior to the last number being drawn. Since participants had to give their estimates twice to the same questions, we informed them that they were free to provide the same estimate or a different estimate the second time around. In the control condition participants saw the description of the lottery immediately followed by the results of the first four numbers, and then answered the questions shown above just once.

In the lottery scenario knowing the results of the first four draws provides no further information about C once it is granted that Mr X wins the lottery. In other words, $Pr_2(C | W) = Pr_1(C | W, f) = Pr_1(C | W)$, where f is knowledge of the results of the first four draws. For this reason, invariance seems justified.

Ultimatum game scenario

Here we use A (*accepting the offer*) and O (*offering at least \$4*) to replace G and B . The following description was presented on the computer screen for each participant:

Imagine that two undergraduate students are randomly chosen from Princeton University to play a game. The game works as follows. The two students are given the opportunity to split \$10. One student is the proposer and the other is the responder. The proposer makes an offer as to how \$10 should be split between the two. The responder can either accept or reject this offer. If the responder accepts the offer, the money is split as proposed, but if the responder rejects the offer, then neither of them receives anything. The students have just finished the first trial of the game.

In the experimental condition the participant was informed that the two students were about to play the second trial. The participant then answered the following questions (randomly ordered) about the second trial *prior to learning the outcome of the first trial*:

ULTIMATUM GAME PROBABILITY QUESTIONS:

- $Pr(A)$ What's the probability that the responder will accept the offer from the proposer in the second trial?
- $Pr(O)$ What's the probability that the proposer will offer AT LEAST \$4 to the responder in the second trial?
- $Pr(A | O)$ What's the probability that the responder will accept the offer assuming that the proposer offers AT LEAST \$4 in the second trial?
- $Pr(A | \bar{O})$ What's the probability that the responder will accept the offer assuming that the proposer offers LESS THAN \$4 in the second trial?

The participant was then presented with the following additional information about the scenario:

Now you learn that in the first trial the responder rejected the proposer's offer and neither of them received anything. They are about to play the second trial.

The participant then answered the same set of questions shown above (in a different random order) about the second trial of the game. Again we informed participants that they were free to provide the same estimate or a different estimate the second time around. In the control condition participants saw the description of the ultimatum game immediately followed by the outcome of the first trial, and then answered the questions shown above just once.

In this scenario invariance is not normatively required because the outcome of the first trial suggests that the responder is sensitive to the fairness of offers. This could affect the second trial in various ways (e.g., by causing the proposer to raise the offer). Thus $Pr_2(A | O) = Pr_1(A | O, t) \neq Pr_1(A | O)$, where t is knowledge of the outcome of the first trial.

Results

Average responses. In each scenario we use Pr_1 to refer to probability estimates before the experience and Pr_2 for estimates after the experience. Average probabilities are shown in Table 3.

Control versus experimental conditions. Since experimental participants had to provide estimates twice to the same questions, we compared their second estimates to those of the control group. For each scenario considered separately, independent-sample t -tests for each of the four questions

TABLE 3
Average estimates in Experiment 2

<i>lottery</i>	$Pr(C)$	$Pr(W)$	$Pr(C W)$	$Pr(C \bar{W})$
Pr_1	0.29 (0.22)	0.01 (0.01)	0.77 (0.22)	0.21 (0.14)
Pr_2	0.33 (0.25)	0.03 (0.02)	0.76 (0.20)	0.20 (0.15)
Control	0.31 (0.19)	0.04 (0.04)	0.70 (0.29)	0.29 (0.20)
<i>ultimatum</i>	$Pr(A)$	$Pr(O)$	$Pr(A O)$	$Pr(A \bar{O})$
Pr_1	0.65 (0.22)	0.61 (0.29)	0.75 (0.19)	0.44 (0.30)
Pr_2	0.68 (0.22)	0.74 (0.24)	0.72 (0.20)	0.38 (0.30)
Control	0.62 (0.25)	0.63 (0.28)	0.70 (0.25)	0.36 (0.29)

Standard deviations in parentheses.

revealed no reliable differences between Pr_2 and the control estimates. Thus experimental participants seem not to have been influenced by having to evaluate the same probabilities twice.

Analysis of the lottery scenario. As a manipulation check, we first determined whether $Pr(W)$ increased after participants saw the results of the first four draws. As expected, $Pr_2(W)$ was reliably larger than $Pr_1(W)$: paired $t(49) = 7.7$, $p < .01$, Wilcoxon $p < .01$. Specifically, the median probability of winning moved from .7% to 2.5%. (Notice that, as a group, the participants' initial estimates were far too high whereas the second estimates were more accurate.)

To see whether invariance holds, we compared $Pr_1(C | W)$ to $Pr_2(C | W)$ via paired t -test and found no reliable difference, $t(49) = 0.89$, $p = .37$. Of the 50 participants, only 12 offered different values for $Pr_2(C | W)$ versus $Pr_1(C | W)$. $Pr_1(C | \bar{W})$ was likewise found to be close to $Pr_2(C | \bar{W})$, paired $t(49) = 1.0$, $p = .34$. Out of 50 participants, 15 gave a different $Pr_2(C | \bar{W})$ estimate from $Pr_1(C | \bar{W})$.

For a more precise quantification, for each participant we calculated invariance violation via the following:

$$\text{Invariance violation} = \frac{|Pr_2(C | W) - Pr_1(C | W)|}{Pr_1(C | W)}$$

The average invariance violation was only 1.42% with a median violation of 0. The invariance violation for $Pr(C | \bar{W})$ was also computed. The mean violation was 1.33% with a median of 0. Since we obtained no estimate for $Pr(W | C)$, converse movement was not computed.

From the results above, invariance seems to hold rather well. For each participant we therefore computed $Pr_{\text{total}}(C)$ via (2) and $Pr_{\text{Jeff}}(C)$ via (4), with G and B substituted by C and W . These estimates were compared to the participant's direct evaluation of $Pr_2(C)$ via absolute difference:

$$\text{total error} = |Pr_2(C) - Pr_{\text{total}}(C)|$$

$$\text{Jeffrey error} = |Pr_2(C) - Pr_{\text{Jeff}}(C)|$$

The means for total and Jeffrey error were .15 and .13, respectively, not reliably different via paired t -test, $t(49) = 0.9$, $p = .36$, or Wilcoxon test, $p = .74$. Jeffrey's rule thus provides a reasonably accurate approximation of $Pr_2(C)$.

Regression of $Pr_2(C)$ on $Pr_1(C | W)Pr_2(W)$ and $Pr_1(C | \bar{W})Pr_2(\bar{W})$ suggests that participants focused more on the second term than the first. The second yields $t = 4.4$, $p < .01$, coefficient = 1.0, whereas the first

produces $t = 0.1$, $p = .87$, coefficient = -0.1 . The same results are seen from the regression of $Pr_2(C)$ on $Pr_2(C | W)Pr_2(W)$ and $Pr_2(C | \overline{W})Pr_2(\overline{W})$, which again suggests that participants focused more on the second term than the first. The second term yields $t = 4.3$, $p < .01$, coefficient = 0.9 whereas the first produces $t = 0.24$, $p = .81$, coefficient = 0.1 .

Analysis of the ultimatum game scenario. As a manipulation check we first determined whether $Pr(O)$ increased after the reported rejection in the preceding trial. As expected, $Pr_2(O)$ was reliably larger than $Pr_1(O)$: paired $t(49) = 4.2$, $p < .01$, Wilcoxon $p < .01$. (Medians 60% and 80%, respectively.)

We found that $Pr_1(A | O)$ was close to $Pr_2(A | O)$, paired $t(49) = 0.3$, $p = .75$. But these averages mask movement in different directions by the majority of participants. Indeed, 33 of the 50 participants offered a different $Pr_2(A | O)$ estimate from $Pr_1(A | O)$ ($p < .05$ via binomial test). Of the 33, 19 offered a higher $Pr_2(A | O)$ than $Pr_1(A | O)$ and 14 offered a higher $Pr_1(A | O)$ than $Pr_2(A | O)$. $Pr_1(A | \overline{O})$ was likewise found to be close to $Pr_2(A | \overline{O})$, paired $t(49) = 1.3$, $p = .20$. However, out of 50 participants, 38 gave a different $Pr_2(A | \overline{O})$ estimate from $Pr_1(A | \overline{O})$ ($p < .01$ via binomial test). Of the 38, 14 offered a higher $Pr_2(A | \overline{O})$ than $Pr_1(A | \overline{O})$ and 24 offered a higher $Pr_1(A | \overline{O})$ than $Pr_2(A | \overline{O})$.

As in the lottery scenario, for each participant we calculated invariance violation via:

$$\text{Invariance violation} = \frac{|Pr_2(A | O) - Pr_1(A | O)|}{Pr_1(A | O)}$$

The average invariance violation was 18.71% with a median of 12%. As expected, this violation was reliably greater than that in the lottery case (1.42%), paired $t(49) = 3.7$, $p < .01$. For $Pr(A | \overline{O})$ the mean invariance violation was 17.38%, reliably greater than $Pr(C | \overline{W})$ in the lottery case (1.33%), paired $t(49) = 2.9$, $p < .01$. The means of total and Jeffrey error were .13 and .16 respectively, which shows an opposite trend to those in the lottery scenario (total error = .15 and Jeffrey error = .13). However, the difference is not reliable via paired t -test, $t(49) = 1.01$, $p = .32$. Overall, as expected on normative grounds, invariance was violated to a greater extent here than in the lottery scenario.

Discussion of Experiment 2

In the lottery scenario invariance held for a majority of participants, and the absolute difference between $Pr_2(C | W)$ and $Pr_1(C | W)$ as a percentage of

$Pr_1(C | W)$ was quite small. When used to estimate $Pr_2(C)$ via the law of total probability, we saw that $Pr_1(C | W)$ could be substituted for $Pr_2(C | W)$ with little loss of accuracy. In contrast to Experiment 1, regression revealed greater impact of $Pr_1(C | \overline{W})Pr_2(\overline{W})$ compared to $Pr_1(C | W)Pr_2(W)$ when using Jeffrey's rule to predict $Pr_2(C)$ in the lottery scenario. The difference between the two experiments might result from implicit causal analysis of the variables in play (as discussed in Over, Hadjichristidis, Evans, Handley, & Sloman, 2007). It might also be that participants focus more on the Jeffrey term, which is associated with the highest unconditional probability.⁴ Consistent with this hypothesis, participants focused more on $Pr_1(C | \overline{W})Pr_2(\overline{W})$, and $Pr_2(\overline{W})$ was greater than $Pr_2(W)$. In the flashlight study the Blue participants focused more on $Pr_1(G | B)Pr_2(B)$, and $Pr_2(B)$ was greater than $Pr_2(\overline{B})$; in contrast, the Purple participants focused more on $Pr_1(G | \overline{B})Pr_2(\overline{B})$, and $Pr_2(\overline{B})$ was greater than $Pr_2(B)$.

In the ultimatum scenario a majority of participants gave different estimates for $Pr_2(A | O)$ compared to $Pr_1(A | O)$ after learning the outcome of the first trial. Thus invariance seems not to hold for the ultimatum scenario, as it need not on normative grounds.⁵

Finally we note that the mean invariance "violation" in the ultimatum scenario is less than the violation in Experiment 1 (18.45% versus 27.73%, averaging over all participants and relevant questions in the respective experiments). This might seem strange, inasmuch as invariance was normatively expected for the flashlight procedure but not the ultimatum game. Does the comparison highlight the non-normative character of the judgements exhibited in Experiment 1? In fact any such comparison is hazardous. Although invariance may not be normatively mandated in a given situation, there might be other reasons to believe in the relative stability of $Pr(A | O)$ and $Pr(A | \overline{O})$ through time. For example, it might be believed that the responder aims for consistency in her successive decisions. In any event, 71 out of 100 participants showed conditional probability movement in the ultimatum scenario but only 40 out of 80 did in Experiment 1.

GENERAL DISCUSSION

In Experiment 1 experience with the light changed the probability that the chosen card was blue, but had only mild (although non-negligible) impact on the probability of the giraffe *given that* the card was blue. That is,

⁴We thank an anonymous reviewer for this observation.

⁵Without giving details, we note that Experiment 2 was repeated with 330 participants recruited over the internet via *Amazon Turk*. The results were in close agreement with those reported here.

$Pr_2(G | B) \approx Pr_1(G | B)$ as well as $Pr_2(G | \bar{B}) \approx Pr_1(G | \bar{B})$. These results conform to Jeffrey's (1983) invariance requirement for updating a distribution on the basis of events whose probabilities are modified without reaching certainty. As a consequence, the updated probability $Pr_2(G)$ was equally well predicted from the law of total probability on the basis of $Pr_1(G | B)$ and $Pr_1(G | \bar{B})$ versus $Pr_2(G | B)$ and $Pr_2(G | \bar{B})$.

The invariance documented in Experiment 1, moreover, was selective inasmuch as greater movement was seen between the converse probabilities $Pr_1(B | G)$ and $Pr_2(B | G)$ than between $Pr_1(G | B)$ and $Pr_2(G | B)$. The difference in movement makes normative sense because the giraffe is conditionally independent of the light given the colour of the card, whereas the colour of the card is not conditionally independent of the light given the giraffe. Experiment 1 thus provides evidence that the participants were sensitive to the normative appeal of Jeffrey's rule, distinguishing (at least partially) between situations where it legitimately applies and where it does not.

The same conclusion is suggested by the results of Experiment 2. Only one of the two scenarios—involving the state lottery rather than the Ultimatum game—gave grounds for invariance, and participants honoured the principle more in the lottery context. In the latter setting, $Pr_2(C)$ (the revised probability of a car purchase) was predicted equally well from the law of total probability based on $Pr_1(C | W)$ and $Pr_1(C | \bar{W})$, as it was from $Pr_2(C | W)$ and $Pr_2(C | \bar{W})$.

Although the experiments suggest that many people (tacitly) respect Jeffrey's rule, the fact remains that a majority of Experiment 1 participants (22 of 40) changed their estimate of $Pr(G | B)$ between times 1 and 2. Almost half of the participants (18 of 40) did so for $Pr(G | \bar{B})$. In percentage terms these shifts were sizeable in Experiment 1 (averaging around 33%) although much smaller in Experiment 2 (less than 2%). The psychology of updating is incomplete without an explanation of why invariance is not respected scrupulously in settings where it seems to be required normatively.

The mere fact of evaluating the same probabilities twice might explain some of the violation of invariance. However, the results of our control conditions suggest that this effect was minor. Recall that control participants responded just once, in the phase 2 setting (e.g., after the light), yet produced estimates that were not reliably different from those gathered in phase 2 of the experimental condition. Another source of invariance violation might be illicit conversion of conditional probability statements, e.g., evaluating $Pr(B | G)$ in place of the requested $Pr(G | B)$. This explanation is consistent with studies that highlight such conversion (as in Dawes, Mirels, Gold, & Donahue, 1993), but inconsistent with our own examination of conditional probability judgements (Zhao, Shah, & Osherson, 2009) in which little conversion was observed.

A third possibility is that the flashlight procedure drew attention to the colour dimension of the stimulus, reminding the participant of its predictive value. This realisation might have been translated into more extreme conditional probabilities (higher for blue, lower for purple). The less-vivid experience in the lottery scenario (merely being told about the first four numbers) would have had less impact, explaining the difference between the two experiments. As an alternative to vivacity, the greater impact of the flashlight might be related to the ineffable character of sensory impressions (as stressed by Jeffrey, 1983, §11.1); the event of matching the first four lottery numbers, in contrast, is more tangible. Of course, more data are needed to test hypotheses such as these.

Additional studies are needed to examine individual differences in respect for Jeffrey's rule, e.g., based on cognitive ability, as in Evans, Handley, Neilens, and Over (2007). Finally, and perhaps most fundamentally, we note that our study leaves untouched the psychological basis of conditional independence, which underlies invariance. Its explanation is no doubt connected to the perception of causal relations among events (or their absence), as discussed in Sloman and Lagnado (2004). Moreover, questions about conditional independence lead to questions about the nature of conditional probability (Zhao et al., 2009) and its relation to conditional constructions in natural language (Over et al., 2007). Gaining even small insight into these issues would count as great progress in the psychology of reasoning.

Manuscript received 30 June 2010

Revised manuscript received 25 August 2010

First published online day/month/year

REFERENCES

- Adams, E. W. (1991). On the meaning of conditionals. *Philosophical Topics*, *XV*, 5–21.
- Chan, H., & Darwiche, A. (2005). On the revision of probabilistic beliefs using uncertain evidence. *Artificial Intelligence*, *163*, 67–90.
- Dawes, R., Mirels, H. L., Gold, E., & Donahue, E. (1993). Equating inverse probabilities in implicit personality judgements. *Psychological Science*, *4*, 396–400.
- Döring, F. (1999). Why Bayesian psychology is incomplete. *Philosophy of Science*, *66*, 379–389.
- Evans, J. S. B. T., Handley, S. J., Neilens, H., & Over, D. E. (2007). Thinking about conditionals: A study of individual differences. *Memory and Cognition*, *35*, 1759–1771.
- Gärdenfors, P. (1988). *Knowledge in flux: Modeling the dynamics of epistemic states*. Cambridge, MA: MIT Press.
- Hacking, I. (2001). *An introduction to probability and inductive logic*. Cambridge, UK: Cambridge University Press.
- Hahn, U., & Oaksford, M. (2007). The rationality of informal argumentation: A Bayesian approach to reasoning fallacies. *Psychological Review*, *114*, 704–732.
- Harman, G. (1999). *Reasoning, meaning and mind*. Oxford, UK: Oxford University Press.

- Jeffrey, R. C. (1983). *The logic of decision* (2nd ed.). Chicago, IL: The University of Chicago Press.
- Jeffrey, R. C. (2004). *Subjective probability: The real thing*. Cambridge, UK: Cambridge University Press.
- Oaksford, M., & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford, UK: Oxford University Press.
- Oberauer, K., & Wilhelm, O. (2003). The meaning(s) of conditionals: Conditional probabilities, mental models and personal utilities. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *29*, 688–693.
- Osherson, D. (2002). *Order dependence and Jeffrey conditionalisation*. Available via <http://www.princeton.edu/osherson/papers/jeff3.pdf>
- Over, D., & Hadjichristidis, C. (2009). Uncertain premises and Jeffrey's rule. *Behavioral and Brain Sciences*, *32*, 97–98.
- Over, D., Hadjichristidis, C., Evans, J. S. B. T., Handley, S. J., & Sloman, S. A. (2007). The probability of causal conditionals. *Cognitive Psychology*, *54*, 62–97.
- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems*. San Mateo, CA: Morgan Kaufmann.
- Sloman, S. A., & Lagnado, D. (2004). Causal invariance in reasoning and learning. In B. Ross (Ed.), *The psychology of learning and motivation* (pp. 287–325). San Diego, CA: Academic Press.
- Stanovich, K. (2010). *Decision making and rationality in the modern world*. Oxford, UK: Oxford University Press.
- Wagner, C. G. (2002). Probability kinematics and commutativity. *Philosophy of Science*, *69*, 266–278.
- Williams, P. M. (1980). Bayesian conditionalisation and the principle of minimum information. *British Journal for the Philosophy of Science*, *31*, 131–144.
- Zhao, J., Shah, A., & Osherson, D. (2009). On the provenance of judgements of conditional probability. *Cognition*, *113*, 26–36.