

Bayesian Inference Underlies the Contraction Bias in Delayed Comparison Tasks

Paymon Ashourian^{1*}, Yonatan Loewenstein^{2*}

1 Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, Massachusetts, United States of America, **2** Department of Neurobiology, Edmond and Lily Safra Center for Brain Sciences, The Interdisciplinary Center for Neural Computation and Center for the Study of Rationality, The Hebrew University of Jerusalem, Jerusalem, Israel

Abstract

Delayed comparison tasks are widely used in the study of working memory and perception in psychology and neuroscience. It has long been known, however, that decisions in these tasks are biased. When the two stimuli in a delayed comparison trial are small in magnitude, subjects tend to report that the first stimulus is larger than the second stimulus. In contrast, subjects tend to report that the second stimulus is larger than the first when the stimuli are relatively large. Here we study the computational principles underlying this bias, also known as the contraction bias. We propose that the contraction bias results from a Bayesian computation in which a noisy representation of a magnitude is combined with a-priori information about the distribution of magnitudes to optimize performance. We test our hypothesis on choice behavior in a visual delayed comparison experiment by studying the effect of (i) changing the prior distribution and (ii) changing the uncertainty in the memorized stimulus. We show that choice behavior in both manipulations is consistent with the performance of an observer who uses a Bayesian inference in order to improve performance. Moreover, our results suggest that the contraction bias arises during memory retrieval/decision making and not during memory encoding. These results support the notion that the contraction bias illusion can be understood as resulting from optimality considerations.

Citation: Ashourian P, Loewenstein Y (2011) Bayesian Inference Underlies the Contraction Bias in Delayed Comparison Tasks. PLoS ONE 6(5): e19551. doi:10.1371/journal.pone.0019551

Editor: Adrian G. Dyer, Monash University, Australia

Received: December 9, 2010; **Accepted:** April 5, 2011; **Published:** May 12, 2011

Copyright: © 2011 Ashourian, Loewenstein. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Funding: This research was supported by grants from the National Institute for Psychobiology in Israel - founded by The Charles E. Smith Family - and from the Israel Science Foundation (grant No. 868/08). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors have declared that no competing interests exist.

* E-mail: paymon@mit.edu (PA); yonatan@huji.ac.il (YL)

Introduction

Comparing magnitudes of two temporally separated stimuli is one of the fundamental tools of experimental psychology and neuroscience. Interestingly, choice behavior in these experiments reveals a fundamental bias: when the first stimulus is small, subjects tend to overestimate it, whereas when it is large, they tend to underestimate it. The first account of this bias, known as the contraction bias, was published a century ago by Harry Levi Hollingworth who later became one of the pioneers of applied psychology. Hollingworth presented subjects with square cards of various sizes for a brief period of time and asked them to memorize their sizes [1]. Each card presentation was followed by a short delay, after which the subjects selected a matching card from a set of probe cards. Surprisingly, Hollingworth observed that subjects tended to choose a probe card that was too large when the memorized card was small compared to the other cards used in the experiment, whereas the opposite behavior, i.e. picking too small a probe card, was observed when the memorized card was relatively large. This illusion has been demonstrated numerous times since Hollingworth's publication for a variety of analog magnitudes in the visual, auditory, and somatosensory modalities [1–7, for review see 8].

The customary explanation for the contraction bias is that the perceived magnitude of a stimulus is a weighted combination of its veridical magnitude and a reference magnitude, such as an

average of all contextually relevant stimuli, that serves as an anchor [3,9, but see 10]. Thus in Hollingworth's experiments and others [1–5] the anchor is thought to make a larger contribution to the subjective magnitude of the memorized stimulus than to the subjective magnitude of the probe stimulus. As a result, the memorized stimulus is biased towards the anchor more than the probe stimulus, which results in the overestimation of small memorized stimuli and the underestimation of large memorized stimuli. This explanation, however, is at best partial since there is no consensus on the choice of the contextually relevant stimuli that comprise the anchor, or on the relative weights of the physical and reference magnitudes. Moreover, it is not clear why the weight applied to the memorized stimulus should be different from the weight applied to the probe stimulus. Finally, the computational principles underlying this bias remain unknown. In order to address these questions we explored whether the contraction bias can be understood as resulting from optimality considerations.

There is a growing body of literature suggesting that the brain utilizes Bayes' rule to optimally combine information from different sources [11–18]. In particular, the application of Bayes' rule has been demonstrated in slant perception [19], sensorimotor learning [20], speed estimation [18], time estimation and interval timing [21], motion perception [22], and integration of information from different sensory modalities [12,23]. In addition, it has been suggested that Bayesian inference underlies the effect of categories on behavior in reconstruction tasks [24]. Therefore, we

hypothesized that the contraction bias in delayed comparison tasks results from a Bayesian inference in which noisy representations of stimuli are combined with knowledge about the a-priori distribution of magnitudes in order to optimize performance. Intuitively, such an inference should lead to the contraction bias because the perception of extreme magnitudes of the first stimulus, which are unlikely given unimodal prior distributions, will be biased toward the ‘center’ of the prior distribution.

In order to test this hypothesis, we conducted an experiment in which we instructed subjects to memorize the length of a bar presented on a computer screen and then compare this memorized length to the length of a probe bar. We show that contraction bias depends on the prior distribution of bar lengths, and increasing the uncertainty in the memory of bar lengths enhances the contraction bias, both of which are consistent with the Bayesian hypothesis.

When within a trial does the Bayesian computation take place? Is the encoded memory biased or does the prior information bias the result of the length comparison? By manipulating uncertainty in the memory of bar lengths *after* memory encoding and measuring the magnitude of the contraction bias we demonstrate that prior information is introduced during memory retrieval/decision making rather than when the first stimulus is encoded in memory.

Some of the findings presented here have appeared previously in abstract form [25].

Results

Example of Contraction Bias

In the standard task (Figure 1A), subjects viewed a horizontal bar (L_1) for 1 sec and were instructed to memorize its length. After a delay of 1 sec, during which screen remained blank, they viewed a probe bar (L_2). The probe bar remained visible on the screen until subjects reported which of the two bars was longer by pressing dedicated keys on the keyboard. The first bar, L_1 , was drawn from a uniform distribution in the logarithmic scale between 150 and 600 pixels. The difference in length between L_1 and L_2 varied between -30% and $+30\%$. Both bars were presented at random locations on the screen and no feedback was provided to the subjects on performance on individual trials (See Materials and Methods). We quantified the proficiency of individual subjects on the delayed comparison task by measuring psychometric curves that depict the percentage of ‘ $L_1 > L_2$ ’ responses as a function of the difference between the memorized and probe stimuli. The average psychometric curve of the subjects ($n = 9$) is plotted in Figure 1B, showing that accuracy improved as the absolute difference between the lengths of L_1 and L_2 increased.

Our purpose is to quantify the contraction bias in these experiments. Previous studies have demonstrated a contraction bias in delayed comparison tasks by showing that the *pattern* of errors made by subjects depends on the magnitude of the memorized stimulus. When the memorized stimulus is small, subjects tend to make more errors in trials in which the probe stimulus is larger than the memorized stimulus, compared to trials in which the probe is smaller than the memorized stimulus. The opposite behavior is observed when the magnitude of the memorized stimulus is large [1–5]. However, these errors only provide a qualitative measure of the bias because the number of errors depends on the relative difficulty of the task, i.e., the difference between the two stimuli (L_1 and L_2 in our experiments) relative to the width of the psychometric curve. We used a different approach to overcome this limitation: Unbeknownst to the subjects, we included a subset of trials in which the lengths of the two bars were identical. We term these trials “impossible

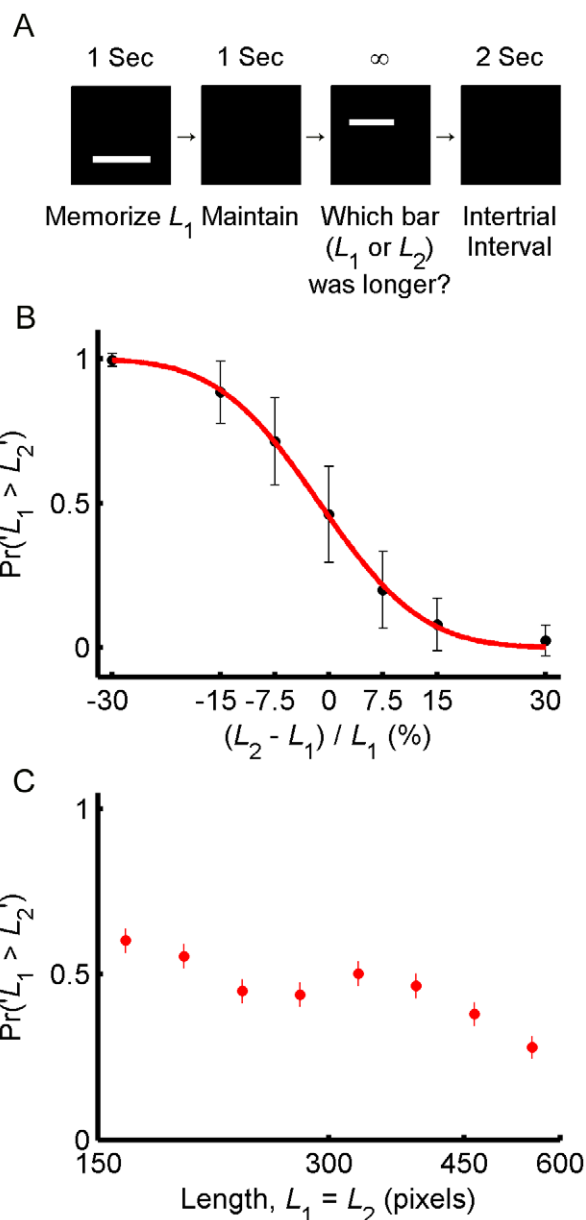


Figure 1. The delayed comparison task and subjects' performance. **A**, The standard task. Subjects viewed a horizontal bar (L_1) on a computer screen for 1 sec and memorized its length. After a delay period of 1 sec, during which the screen remained blank, the subjects viewed a second bar (L_2) and were instructed to report which of the two bars was longer. The second bar, L_2 remained visible until subjects made a response. The difference in length between L_1 and L_2 varied between -30% and $+30\%$. Unbeknownst to the subjects, on roughly 50% of the trials, the lengths of the first and second bars were equal ($L_1 = L_2$). **B**, The average psychometric curve of 9 subjects. The abscissa corresponds to the difference between the two bar lengths, $(L_2 - L_1) / L_1$ and the ordinate corresponds to the fraction of trials in which subjects chose L_1 as longer than L_2 . Error bars depict standard error of the mean (SEM). Line is a least-square fit of an error function: $\Pr(L_1 > L_2) = \frac{1}{2} \cdot \left(1 - \operatorname{erf} \left(\frac{(L_2 - L_1) / L_1 - \mu}{\sqrt{2}\sigma} \right) \right)$ where $\mu = -0.013$ and $\sigma = 0.11$. **C**, Average response curve of 9 subjects. Fraction of times in which subjects reported ‘ $L_1 > L_2$ ’ on the impossible trials are plotted as a function of bar length. Subjects overestimated the magnitude of the memorized L_1 bar when it was relatively small and underestimated L_1 when it was relatively long, consistent with the contraction bias. Each data point corresponds to 21 impossible trials per subject. Error bars depict SEM. doi:10.1371/journal.pone.0019551.g001

trials” because there is no correct answer to the question “which bar (L_1 or L_2) was longer?”. Impossible trials are well suited for the analysis of the contraction bias because performance on these trials is independent of the proficiency of individual subjects in distinguishing the difference in the length of the two bars.

The average response curve of 9 subjects is depicted in Figure 1C, where we plot the percentage of trials in which the subjects reported that ‘ $L_1 > L_2$ ’ as a function of the length of L_1 ($L_1 = L_2$). Note that despite the fact that L_1 and L_2 were identical on these trials, subjects reported that L_1 was longer than L_2 on roughly 60% of the shortest trials (left-most point in Figure 1C) whereas they reported L_1 was longer than L_2 only in 28% of the longest trials (right-most point in Figure 1C). The slope of the regression line fitted to the impossible trials was significantly smaller than zero (mean slope = -0.28 , 95% bootstrap confidence interval (CI) = $[-0.36, -0.21]$, see Materials and Methods for procedure), indicating that subjects were more likely to report ‘ $L_1 > L_2$ ’ for shorter L_1 bars as compared to longer L_1 bars, thus exhibiting the contraction bias.

Bayesian Inference and Contraction Bias

Our aim is to account for the contraction bias in a Bayesian framework of decision making. In order to see how the contraction bias emerges from Bayesian inference, we consider a control region in the brain, such as the prefrontal cortex [26], that is presented with the neural representations of L_1 and L_2 and has to decide which of the two bars is longer. We assume that: (1) the control region knows that the neural representations of L_1 and L_2 are noisy, e.g. due to noise in the sensory pathway. Moreover the representation of L_1 is noisier than that of L_2 because L_1 has to be stored in memory, a process that may contribute additional noise to the representation of L_1 ; (2) the control region has information about the marginal distribution of bar lengths. This distribution can be approximated based on the history of the experiment; (3) the control region utilizes Bayes’ rule and combines the noisy representations of L_1 and L_2 with knowledge about the prior distribution in order to construct the posterior distributions for the two bar lengths. These posteriors are then used to minimize error in judgment. A formal description of this process appears in the Materials and Methods section. To illustrate how the contraction bias could emerge from such a Bayesian computation, we consider the following three examples:

- (1) **L_1 is unknown, L_2 is known.** Consider a hypothetical subject who forgets the length of L_1 , but has no ambiguity about the length of L_2 , i.e., the neural representation of L_1 is infinitely noisy whereas there is no noise in the neural representation of L_2 . In this case, the posterior of L_1 is the prior distribution. In contrast, the prior distribution makes no contribution to the posterior of L_2 . Therefore, the optimal strategy would be to report ‘ $L_1 > L_2$ ’ in trials where L_2 is smaller than the median of the prior distribution and to report ‘ $L_1 < L_2$ ’ in trials in which L_2 is larger than the median. Therefore, in the impossible trials in which $L_1 = L_2$, the subject would report ‘ $L_1 > L_2$ ’ if L_1 is smaller than the median of the prior distribution and would report ‘ $L_1 < L_2$ ’ otherwise, as depicted in Figure 2A. This response pattern is consistent with the contraction bias because it appears as though the subject is overestimating relatively small L_1 and underestimating relatively large L_1 .
- (2) **L_1 and L_2 are equally uncertain.** Consider a case where the estimated uncertainties in the representations of L_1 and L_2 are equal. This would be true if the only uncertainty in the representations of L_1 and L_2 results from sensory noise, and

memory storage does not add any additional noise to the representation of L_1 . In the impossible trials in which the two bars are physically identical, the contribution of the prior distribution to the posteriors of L_1 and L_2 is equal. Symmetry considerations indicate that the subject would report that ‘ $L_1 > L_2$ ’ at chance level for all bar lengths, i.e., there is no contraction bias, as depicted in Figure 2B.

- (3) **L_1 is less certain than L_2 .** In intermediate cases where the level of uncertainty in L_1 is larger than that of L_2 , for example, as a result of added noise due to memory storage, we expect the resultant response curve to reside between the response curves of Figures 3A and 3B, resulting in a smooth decrease in

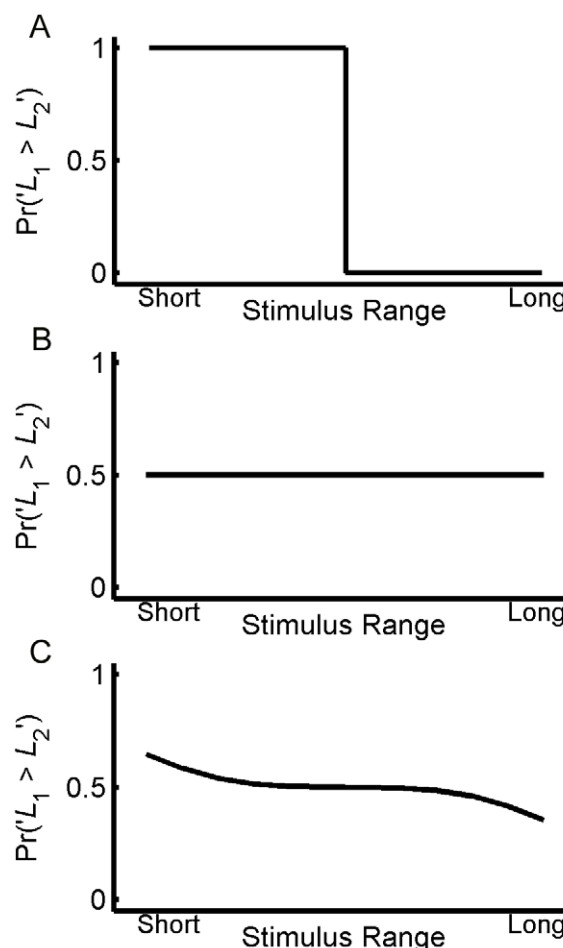


Figure 2. Bayesian inference and the contraction bias. **A**, Response curve of a Bayesian model with infinite noise in the representation of L_1 (no memory of L_1) and no noise in representation of L_2 . The model reports ‘ $L_1 > L_2$ ’ in trials where L_2 is smaller than the median of the prior distribution, and ‘ $L_1 < L_2$ ’ in trials in which L_2 is larger than the median. This behavior arises because the posterior distribution of L_1 is the same as the prior distribution of the bar lengths, whereas the posterior distribution of L_2 is not influenced by the prior at all. **B**, Response curve of the Bayesian model with equal noise in the representations of L_1 and of L_2 . The contribution of the prior to the posteriors of L_1 and L_2 is identical because the levels of noise in the two representations are equal. Thus, in trials where $L_1 = L_2$, the model reports ‘ $L_1 > L_2$ ’ at chance level independently of the length of the bars. **C**, Response curve of the Bayesian model assuming that there is more noise in the representation of L_1 than in representation of L_2 . Response curve is a combination of **A** and **B**. doi:10.1371/journal.pone.0019551.g002

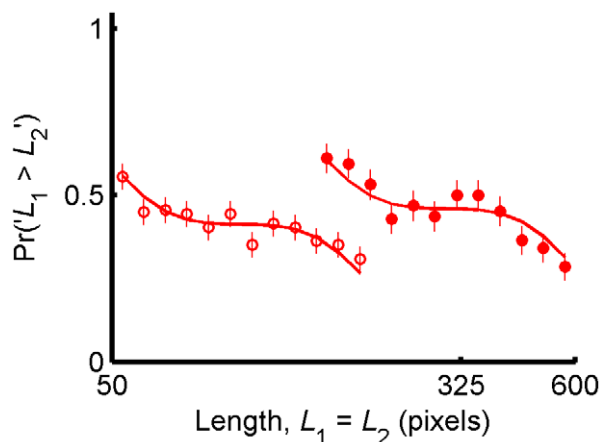


Figure 3. Effect of the prior on the response curve. Assuming that noise is independent of bar length, the model predicts that the shape of the response curve is independent of the physical range of the stimuli. Thus, a lateral shift in the prior would result in a lateral shift in the response curve. Two groups of subjects completed the task in Figure 1A for two overlapping uniform priors, 50 to 200 pixels (open circles) and 150 to 600 pixels (filled circles). The response curve in the impossible trials was not significantly different between the two groups. Each data point corresponds to 14 impossible trials per subject. Error bars depict SEM. Lines are the best fit of the Bayesian model, see Materials and Methods.
doi:10.1371/journal.pone.0019551.g003

the fraction of trials in which L_1 is reported to be larger than L_2 as a function of the lengths of L_1 and L_2 (Figure 2C).

Model Predictions and Behavioral Results

Effect of Changing the Prior. If the contraction bias results from Bayesian inference, then changing the prior distribution is expected to change the response curve. In particular, assuming that noise is independent of the length of the bars, a translational shift in the prior distribution would result in an equal translational shift in the response curve without changing its shape. To test this prediction, we asked a new group of naïve subjects to participate in the experiment of Figure 1A, in which L_1 was drawn from a new uniform distribution in the logarithmic scale between 50 and 200 pixels ($n=10$). Similar to the first experiment, all lengths were presented in logarithmic scale to satisfy the assumption of independence of noise and bar length. We compared the responses of this group to the original group who saw stimuli that were drawn from a uniform distribution in the logarithmic scale between 150 to 600 pixels. The accuracy of subjects in the trials in which $L_1 \neq L_2$ (non-impossible trials) was indistinguishable between the two groups ($83\% \pm 2\%$ for 50–200; $85\% \pm 1\%$ for 150–600; $t_{17} = 0.69$, $p = 0.49$, two-tailed), supporting the assumption that the level of noise in the neural representation of the bars is independent of bar length in these ranges. Response curves for the two groups in the impossible trials are depicted in Figure 3 (50–200 pixels: open circles; 150–600 pixels: filled circles). The slope of a regression line fitted to the response curve for the 50–200 pixel group was significantly smaller than zero (mean slope = -0.20 , 95% bootstrap CI = $[-0.27, -0.13]$). Response curve slopes were not significantly different between the group who saw 50–200 pixel lines and the original group who saw 150–600 pixel lines (average difference = -0.08 ; 95% bootstrap CI = $[-0.19, +0.02]$). Thus, as predicted by the Bayesian hypothesis, a translational shift in the prior distribution resulted in a translational shift in the response curve.

Qualitatively, the shape of the response curve does not seem linear. Rather, the slope of the curve is more negative for bar lengths at the high and low ends of the spectrum. The non-linear response curve is consistent with our Bayesian model whose two parameters, the level of noise in the representation of the two bars, were chosen to minimize the fit mean square error. The resultant parameters indicate that the uncertainty in the representation of L_1 is 30% higher than the uncertainty in L_2 (Materials and Methods). Moreover, the Bayesian model is qualitatively similar to the experimental results, supporting our hypothesis that the contraction bias results from Bayesian inference.

Effect of Noise. According to the Bayesian hypothesis, the contraction bias emerges because the contribution of the prior distribution to the posterior of the first bar is larger than the contribution of prior to the posterior of the second bar. This asymmetry results from the fact that the uncertainty in the representation of the memorized bar, L_1 , is larger than that of the probe bar, L_2 . The larger the asymmetry in the contribution of the prior to the posteriors of the two bars, the more pronounced the contraction bias should be. Therefore, increasing the level of noise in the representation of L_1 is expected to enhance this asymmetry and thus enhance the contraction bias.

To test this prediction, we modified the task of Figure 1A and added a distracting memory task between the presentations of the two bars in randomly selected half of the trials: 500 msec after the disappearance of L_1 , four different colors were flashed on the screen in a random order for 400 msec each (Figure 4A). Then, subjects were instructed to recall the n 'th presented color, where n was a number between 1 and 4, appearing on the screen immediately after the last color. Following the answer, L_2 was presented and subjects were instructed to compare it with L_1 as before. On average, subjects correctly recalled the color in 96% of the trials indicating that the color task was not disregarded.

The distracting task was designed to interfere with the memory of the first bar. In the Bayesian framework, it was intended to add 'noise' to the representation of L_1 . As predicted, accuracy of performance on the bar comparison task, measured in the trials in which $L_1 \neq L_2$, was lower on the trials interrupted by the secondary task compared to the performance on trials that were not interrupted by the secondary task (Figure 4B). This decrease in performance was significant both in the easiest trials in which the difference between L_1 and L_2 was $\pm 30\%$ ($p = 3 \times 10^{-7}$, two-tailed t -test), the intermediate trials in which the difference between L_1 and L_2 was $\pm 15\%$ ($p = 4 \times 10^{-4}$, two-tailed t -test) and the most difficult trials in which the difference between L_1 and L_2 was $\pm 7.5\%$ ($p = 6 \times 10^{-4}$, two-tailed t -test).

In order to characterize the effect of the secondary task on the contraction bias, we compared the response curve of subjects in the impossible trials with interference from the secondary task (open circles in Figure 4C) with the response curve of the same subjects in the impossible trials with no interference from the secondary task (filled circles in Figure 4C). The slope of the linear fit to the response curve in trials devoid of the secondary task was -0.19 (95% bootstrap CI = $[-0.32, -0.07]$). In contrast, the slope of the linear fit to the response curves in trials with the secondary task was -0.63 (95% bootstrap CI = $[-0.75, -0.53]$), which was significantly more negative than the slope in trials without the secondary task (average difference = -0.45 ; 95% bootstrap CI = $[-0.61, -0.28]$). These results suggest that, as predicted by the Bayesian model, an increase in internal noise, which manifests as a decrease in behavioral accuracy, is associated with an increase in the level of contraction bias, which manifests as an increase in the magnitude of the slope of the response curve. To further test this hypothesis, we examined the accuracy of performance on an

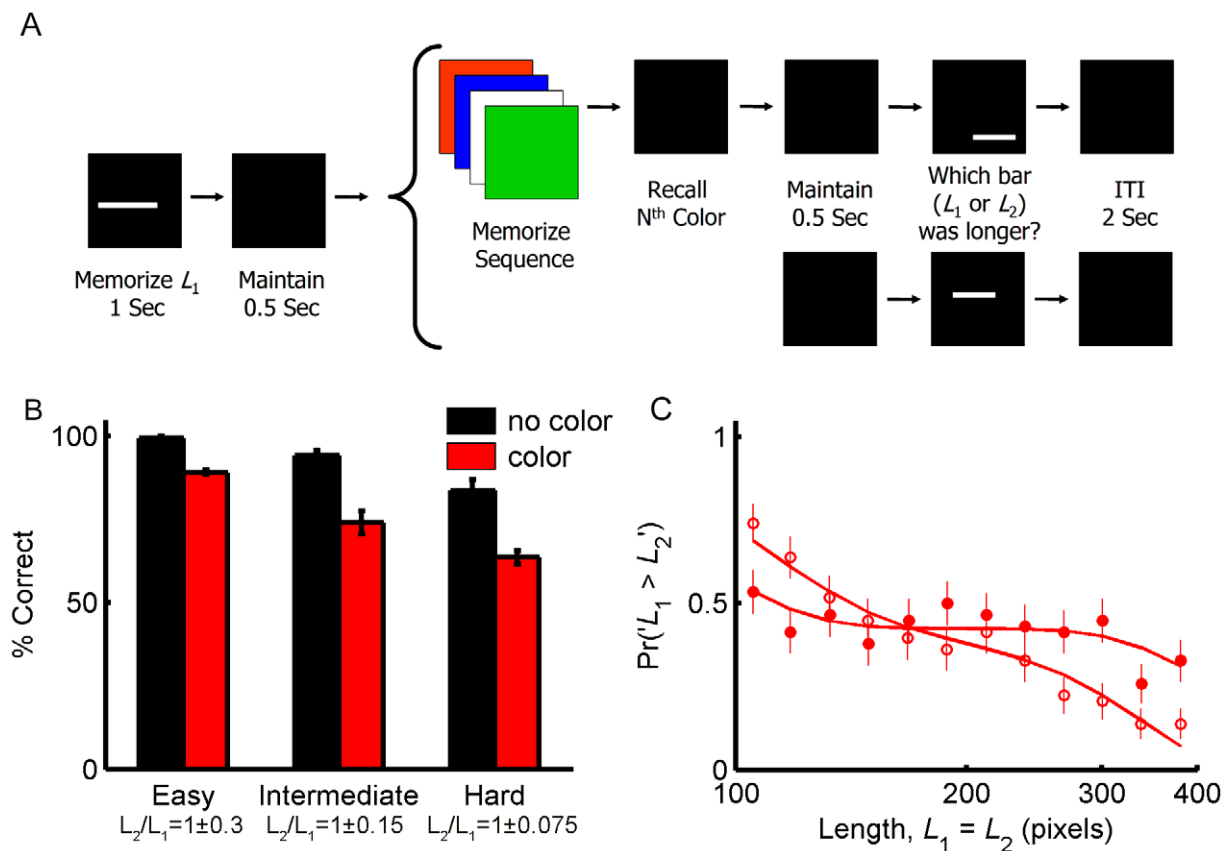


Figure 4. Effect of noise on the response curve. **A**, Subjects performed a modified experiment where a secondary task had to be performed between the presentations of the two bars on randomly selected 50% of the trials. Top row depicts sequence of events in trials with interference: a sequence of 4 colors was presented on the screen 500 msec after the presentation of L_1 . Each color was presented for 400 msec and subjects were instructed to memorize the sequence. 400 msec after the disappearance of the last color, a number from 1 to 4 appeared on the screen. Subjects were instructed to recall the color that corresponded to the number. **B**, Percentage correct in bar length comparison in the standard (black) and modified (red) trials. The ability to memorize the length of L_1 was impaired in the modified trials compared to the standard unperturbed trials, in both the easy ($\pm 30\%$, left), intermediate ($\pm 15\%$, center) and hard ($\pm 7.5\%$, right) trials. These results suggest that the secondary task increased uncertainty in the memory of the length of L_1 . **C**, Response curve in the standard (filled circles) and modified trials (open circles). The larger slope of the response curve on the modified trials compared to the standard trials suggests that the secondary task caused an enhancement of the contraction bias. Each data point corresponds to 6 impossible trials per subject. Error bars depict SEM. Lines are the best fit of the Bayesian model, see Materials and Methods.

doi:10.1371/journal.pone.0019551.g004

individual basis by fitting a psychometric curve (cumulative Gaussian function similar to Figure 1B) to the responses of each subject, once in trials with and then in trials without the distracting task, and estimating the width of the psychometric curve (σ) in each case. Next, we calculated the correlation between the slope of the linear fit to the response curve of each subject (i.e. the magnitude of contraction bias) and their respective psychometric σ . The correlation coefficient between values of σ and the slopes of the response curves was -0.74 ($p = 0.0002$, two-tailed), supporting the assertion that a decrease in performance is associated with an increase in the magnitude of the contraction bias.

Discussion

We examined the hypothesis that the contraction bias in delayed comparison tasks results from a Bayesian inference in which information about the prior distribution is combined with noisy measurement in order to optimize performance. This hypothesis makes two predictions: a translational shift in the prior distribution is expected to result in a similar translational shift in

the bias curve, and increasing noise in memory is expected to increase reliance on prior knowledge and thus increase the bias. Our results are consistent with both predictions, suggesting that the contraction bias results from a Bayesian inference.

Within a single trial, when does information about the prior distribution combine with the sensory measurement? One possibility is that it takes place during the encoding of L_1 . In this case, the encoded memory of L_1 is already biased in the direction of the prior distribution. Another possibility is that the memory of L_1 is unbiased and the Bayesian computation takes place at the comparison stage, when the encoded L_1 is compared with L_2 . To address this question we again considered the choice behavior of subjects in the experiment with the interfering task. We found that in this experiment, the slope of the response curve was more negative in trials with interference from the secondary task, compared to the standard trials (Figure 4C). In other words, more weight was given to the prior distribution in trials interrupted by the secondary task. Recall that trials containing this task were randomly intermixed with trials that did not contain interference. Therefore, at the time of encoding of L_1 (up to 0.5 sec after the end of the presentation of L_1) the subjects could not know whether

they would be presented with an interfering task and therefore could not know what weight to give to the prior distribution. Therefore, if the computation had taken place at the time of the encoding of L_1 , we would have observed no difference in the slope of the response curve between the two conditions. Therefore, the Bayesian computation necessarily took place *after* the interfering task, at the time of L_1 retrieval or later, when L_1 and L_2 were compared.

How do subjects learn the prior distribution? In order to address this question, we compared the level of contraction bias, as measured by the slope of the response curve, in the first 20 impossible trials to the slope in the last 20 impossible trials for subjects who completed the experiment in Figure 1A where the bar lengths were drawn from the 150–600 and 50–200 ranges. We found no statistical difference in these slopes (−0.29 for the first 20 trials; −0.28 for the last 20 trials; average difference = −0.01; 95% bootstrap CI for the difference in slopes, [−0.29 0.27]). These results indicate that the contraction bias emerges within a small number of trials, suggesting that the prior distribution of bar lengths in the experiment is estimated using a small number of trials.

In this study we examined the effect of a translational shift in the prior, but we did not alter the shape of the prior distribution. Previous studies have shown that subjects are sensitive to the shape of the prior distribution in category and sensorimotor learning [20,24]. Consistent with these results, changing the shape of the prior distribution in our model changes the shape of the response curve. The extent to which the shape of the prior distribution can be learned and utilized in Bayesian reasoning, however, awaits future studies.

Contraction bias in delayed comparison tasks is a common cognitive illusion observed in many different modalities and under different experimental conditions [1–8]. In this paper we provide a normative interpretation of this bias, supported by an experiment in visual domain. Our results are consistent with a growing body of literature showing that the brain utilizes close-to-optimal computational strategies.

Materials and Methods

Ethics Statement

All subjects gave written informed consent using methods approved by the Massachusetts Institute of Technology Committee on the Use of Humans as Experimental Subjects.

Subjects

Subjects were undergraduate and graduate students from the Massachusetts Institute of Technology. All subjects had normal or corrected-to-normal vision and no subjects took part in more than one of the experiments. Each subject received \$10 plus 1 cent for every correct trial in the experiment for a session lasting less than an hour.

Stimuli

Stimuli were white horizontal bars on a black background displayed on a 17" computer screen with a resolution of 1024×768. All bars were 3 pixels wide.

Procedure

Subjects sat approximately 60 cm from a computer screen in a dimly lit room. Each subject completed 400 to 600 trials in one hour and received feedback on their overall performance after every 20 trials. No other feedback was provided. In the standard task, each trial started with the presentation of a L_1 at a random

location on the screen for 1 sec. After a delay period of 1 sec, during which screen remained blank, L_2 appeared at another random location on the screen. L_2 remained visible until the subjects pressed one of two keys indicating which bar was longer. The difference in length between L_1 and L_2 varied between −30% and +30%. Unbeknownst to the subjects, in roughly 50% of the trials, the lengths of the first and second bars were equal ($L_1 = L_2$). Subjects did not receive feedback on performance on individual trials. Each trial was followed by a 2 sec intertrial interval during which the screen remained blank. Two distinct groups of subjects completed the standard task. One group ($n = 9$) saw L_1 bars chosen uniformly in the logarithmic scale from the [50, 200] pixel interval, while the other group ($n = 10$) saw bars chosen from the [150, 600] pixel interval.

The modified task was identical to the standard task with two exceptions: (1) L_1 bars were chosen uniformly in the logarithmic scale from the [100, 400] pixel interval; (2) subjects completed a distracting memory task between the presentation of L_1 and L_2 in a randomly selected 50% of the trials: 500 msec after L_1 disappeared, a random sequence of four colors (red, blue, white, and green) were displayed on the screen for 400 msec each. 400 msec after the disappearance of the last color, a number from 1 to 4 appeared in yellow on the screen. Subjects were instructed to recall the color that corresponded to the number and press one of four dedicated keys to indicate this color. L_2 appeared 500 msec after subjects made their color choice.

A Bayesian Model of Contraction Bias

According to our Bayesian hypothesis, the contraction bias emerges because subjects use Bayes' law to combine noisy information about the lengths of the bars with knowledge about the prior information in order to optimize performance. In this section we formalize this intuition.

In accordance with Weber's law, the lengths of the bars are measured in logarithmic scale. Let L_i and R_i be the logarithm of the length of bar i and its neural representation, respectively. We assume that this representation is noisy such that $R_i = L_i + z_i$ where z_i is drawn from a zero-mean Gaussian distribution with variance σ_i^2 , $z_i \sim \mathcal{N}(0, \sigma_i^2)$. This is illustrated in Figure 5A where we plot the probability of a neural representation R_i for a given representation of bar length $L_i = L_i^*$, also known as a likelihood function and denoted as $\Pr[R_i|L_i = L_i^*]$. We assume that the prior distribution of bar lengths, $\Pr[L_i]$, is uniform (Figure 5B). Bayes' rule provides a method for combining information about the prior distribution with the noisy neural representation, in order to compute the posterior distribution, $\Pr[L_i|R_i]$ (Figure 5C). According to Bayes' rule

$$\Pr[L_i|R_i] = \frac{\Pr[R_i|L_i] \cdot \Pr[L_i]}{\Pr[R_i]} \tag{1}$$

$$\text{where } \Pr[R_i] = \int_{-\infty}^{\infty} \Pr[R_i|L_i] \cdot \Pr[L_i] dL_i.$$

Given a pair of neural representations, (R_1, R_2) , of the lengths of the first and second bars, the probability that the first bar is longer than the second bar is given by

$$\Pr[L_1 > L_2 | R_1, R_2] = \int_{-\infty}^{\infty} \Pr[L_1 | R_1] \int_{-\infty}^{L_1} \Pr[L_2 | R_2] dL_2 dL_1 \tag{2}$$

This is illustrated in Figure 5D where we use a color scale to plot $\Pr[L_1 > L_2 | R_1, R_2]$ for different values of R_1 and R_2 . The black line

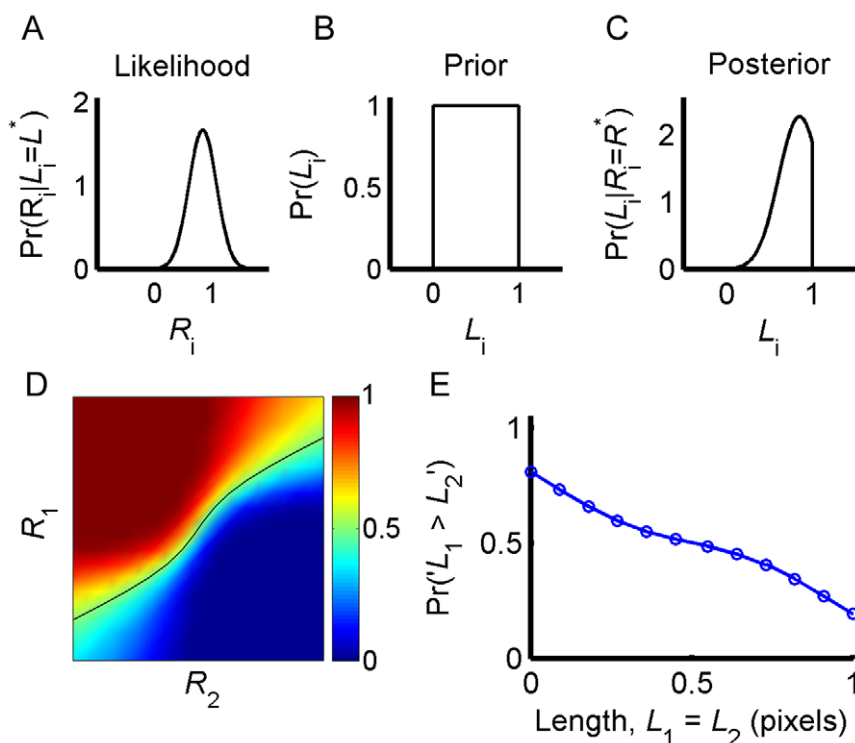


Figure 5. Ideal decision maker solution to the task in Figure 1A. **A**, The likelihood of a representation R_i given a particular length (here $L_i=0.85$, $\sigma_i=0.24$) assuming $R_i \sim N(L_i, \sigma_i^2)$. **B**, The prior distribution of bar lengths. **C**, The posterior distribution of L_i given a particular measurement (here $R_i=0.85$), calculated using Bayes' rule. **D**, The probability that $L_1 > L_2$ for different values of R_1 and R_2 , computed using the posteriors. The black line corresponds to the values of R_1 and R_2 such that $\Pr(L_1 > L_2 | R_1, R_2) = 0.5$ (here, $\sigma_1 = 0.24$ and $\sigma_2 = 0.13$). **E**, Response curve of the model on the impossible trials in which $L_1 = L_2$. doi:10.1371/journal.pone.0019551.g005

corresponds to values of (R_1, R_2) such that $\Pr[L_1 > L_2 | R_1, R_2] = 0.5$. Note that the slope of this curve is smaller than 1. This results from the assumption that $\sigma_1 > \sigma_2$, reflecting the fact that L_1 has to be stored in memory, a process that may contribute additional noise to the representation of L_1 . An ideal Bayesian observer, who has access to R_1 and R_2 , would report ' $L_1 > L_2$ ' in trials in which $\Pr[L_1 > L_2 | R_1, R_2] > 0.5$ and ' $L_1 < L_2$ ' in trials in which $\Pr[L_1 > L_2 | R_1, R_2] < 0.5$. Therefore, the probability that a model would report ' $L_1 > L_2$ ' in a trial in which L_1 and L_2 are presented is given by

$$\Pr['L_1 > L_2' | L_1, L_2] = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \Pr[R'_1 | L_1] \cdot \Pr[R'_2 | L_2] \cdot \Psi(R'_1, R'_2) dR'_1 dR'_2 \quad (3)$$

where $\Psi(R'_1, R'_2) = \begin{cases} 1 & \text{if } \Pr[L_1 > L_2 | R_1, R_2] > 0.5 \\ 0 & \text{otherwise} \end{cases}$.

In order to construct the response curve we compute $\Pr['L_1 > L_2' | L_1, L_2]$ (Figure 5E). For further insights into the Bayesian computation, we consider the simple example in which the level of uncertainty in the representation of L_1 is infinite, whereas there is no uncertainty in the representation of L_2 . In other words, $\sigma_1^2 \rightarrow \infty$ and $\sigma_2^2 = 0$. In this case, Eq. (1) becomes $\Pr[L_1 | R_1] = \Pr[L_1]$ and $\Pr[L_2 | R_2] = \delta(L_2 - R_2)$, Eq. (2) becomes $\Pr[L_1 > L_2 | R_1, R_2] = \int_{R_2}^{\infty} \Pr[L'_1] dL'_1$ and therefore the subject would report ' $L_1 > L_2$ ' if R_2 is larger than the median of L_1 . In trials in which $L_1 = L_2$, Eq. (3) dictates that he would

report ' $L_1 > L_2$ ' in trials in which L_2 is larger than the median and ' $L_1 < L_2$ ' otherwise.

Data analysis

Slope of line fitted to response curve. All slopes were computed after normalizing the range of lengths to 0 and 1 in the logarithmic space.

Bootstrap confidence intervals. We used a pairs bootstrap resampling procedure [27] in order to calculate confidence intervals for the slope of the regression lines. The bootstrap algorithm is as follows: repeated 5,000 times, we sampled (with replacement) from each subject's impossible trials in order to obtain a bootstrap dataset and fitted a regression line to the averaged response curve of each bootstrap dataset. This procedure resulted in 5,000 bootstrap slopes that could be used for calculating a CI for the slope of the regression line fitted to the experimentally obtained data points. The CIs reported in the text are 95% basic bootstrap intervals [27].

In order to compare the response curve slopes between subjects who saw 50–200 pixel lines and those who saw 150–600 pixel lines we sampled from each group independently using the algorithm above, and then constructed a 95% confidence interval on the difference between the bootstrap slopes of the two groups.

In order to compare trials with and without the interference task we calculated the difference in the bootstrap slope of each subjects' standard and interfered trials, and found the 95% confidence interval of this difference. The same method was also used to compare the slope of the response curve in the first 20 impossible trials of the experiment to the slope of the response curve in the last 20 impossible trials of the experiment.

Bayesian model fit. In order to compare behavioral performance to that predicted by the model, we used the model presented above to generate a set of response curves of ideal observers characterized by different values of σ_1 and σ_2 . These curves were compared to the experimentally measured response curves as described below:

Note that subjects exhibited a small bias in favor of reporting ' $L_2 > L_1$ ' in the 50–200 and 150–600 standard experiments. Subjects reported that ' $L_1 > L_2$ ' in the impossible trials in 41% and 46% respectively. This tendency has been reported previously [28,29]. In principle, such a bias can be explained in our Bayesian framework by claiming that the prior distribution that the subjects use in their Bayesian computation is biased in favor of small magnitudes, as was observed for speed perception [18]. In this framework, it is predicted that in the modified experiment (Figure 4A), the global bias should be larger in the trials interfered by the color task than in the standard trials. In fact we found that the global bias was larger in the modified trials (42% vs. 38%). However, this effect was not statistically significant ($p = 0.58$, two tailed t -test). More importantly, this explanation is circular because a bias in the opposite direction could equally well have been explained by arguing that the prior distribution is biased in favor of large magnitudes. Therefore we did not attempt to account for the global bias and subtracted it before fitting, assuming that it is

generated by a different mechanism. Thus, for the purpose of finding the parameters we added a constant to each of the response curves to normalize them such that $\text{mean}(\text{Pr}[\hat{L}_1 > L_2]) = 0.5$. For purposes of comparison, the range of the logarithm of bar lengths was normalized to lie between 0 and 1 and we used a least square fit to find the parameters that best fit the population-average experimental data. We found that the best fit model parameters for the groups who saw 50–200 and 150–600 pixel-long bars were given by $\sigma_1 = 0.13$, $\sigma_2 = 0.1$; The best fits for trials not interfered by the distracting task and those that had the distracting task were $\sigma_1 = 0.11$, $\sigma_2 = 0.09$, and $\sigma_1 = 0.24$, $\sigma_2 = 0.13$, respectively.

Acknowledgments

We thank Sebastian Seung for his encouragement and support and Merav Ahissar, Konrad Körding, Ofri Raviv, and Hanan Shteingart for fruitful discussions.

Author Contributions

Conceived and designed the experiments: YL PA. Performed the experiments: YL PA. Analyzed the data: YL PA. Contributed reagents/materials/analysis tools: YL PA. Wrote the paper: YL PA.

References

- Hollingworth HL (1910) The Central Tendency of Judgment. *The Journal of Philosophy, Psychology and Scientific Methods* 7: 461–469.
- Berliner JE, Durlach NI, Braida LD (1977) Intensity perception. VII. Further data on roving-level discrimination and the resolution and bias edge effects. *The Journal of the Acoustical Society of America* 61: 1577–1585.
- Hellstrom A (1985) The time-order error and its relatives: Mirrors of cognitive processes in comparing. *Psychological Bulletin* 97: 35–61.
- Marks LE (1993) Contextual processing of multidimensional and unidimensional auditory stimuli. *Journal of experimental psychology: Human perception and performance* 19: 227–249.
- Wilken P, Ma WJ (2004) A detection theory account of change detection. *J Vis* 4: 1120–1135.
- Raviv O, Ahissar M, Loewenstein Y (2009) Abstracts of the Israel Society for Neuroscience 18th Annual Meeting. *Journal of Molecular Neuroscience* 39 Suppl 1: S3–132.
- Schwartz E, Romo R, Loewenstein Y (2008) The computational principles and neural mechanisms underlying contraction bias. Program No 19214/UU14 2008 Neuroscience Meeting Planner Washington, DC: Society for Neuroscience, Online.
- Poulton EC (1989) Bias in quantifying judgments. Hove East Sussex, U.K.; Hillsdale USA: L. Erlbaum. 304 p.
- Helson H (1964) Adaptation-level theory; an experimental and systematic approach to behavior. New York: Harper & Row.
- Parducci A (1965) Category judgment: a range-frequency model. *Psychological Review* 72: 407–418.
- Alais D, Burr D (2004) The ventriloquist effect results from near-optimal bimodal integration. *Current Biology* 14: 257–262.
- Ernst MO, Banks MS (2002) Integration of simultaneous visual and haptic information. In: Bulthoff HH, Gegenfurtner KA, Mallot HA, Ulrich R, eds. 47 p.
- Ernst MO, Bulthoff HH (2004) Merging the senses into a robust percept. *Trends in Cognitive Sciences* 8: 162–169.
- Trommershauser J, Gepshtein S, Maloney LT, Landy MS, Banks MS (2005) Optimal compensation for changes in task-relevant movement variability. *The Journal of Neuroscience* 25: 7169–7178.
- Welchman AE, Lam JM, Bulthoff HH (2008) Bayesian motion estimation accounts for a surprising bias in 3D vision. *Proceedings of the National Academy of Sciences* 105: 12087–12092.
- Stocker AA, Simoncelli EP (2006) Noise characteristics and prior expectations in human visual speed perception. *Nature Neuroscience* 9: 578–585.
- Griffiths TL, Tenenbaum JB (2006) Optimal predictions in everyday cognition. *Psychol Sci* 17: 767–773.
- Weiss Y, Simoncelli EP, Adelson EH (2002) Motion illusions as optimal percepts. *Nature Neuroscience* 5: 598–604.
- van Ee R, Adams WJ, Mamassian P (2003) Bayesian modeling of cue interaction: bistability in stereoscopic slant perception. *J Opt Soc Am A Opt Image Sci Vis* 20: 1398–1406.
- Körding KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. *Nature* 427: 244–247.
- Jazayeri M, Shadlen MN (2010) Temporal context calibrates interval timing. *Nature Neuroscience* 13: 1020–1026.
- Chalk M, Seitz AR, Series P (2010) Rapidly learned stimulus expectations alter perception of motion. *J Vis* 10: 2.
- Jacobs RA (1999) Optimal integration of texture and motion cues to depth. *Vision Research* 39: 3621–3629.
- Huttenlocher J, Hedges LV, Vevea JL (2000) Why do categories affect stimulus judgment? *Journal of experimental psychology: General* 129: 220–241.
- Hosseini P, Loewenstein Y (2005) Visual short-term memory and Bayesian decision making. Program No. 16.9. 2005 Neuroscience Meeting Planner; Washington, DC. Society for Neuroscience.
- Miller EK, Freedman DJ, Wallis JD (2002) The prefrontal cortex: categories, concepts and cognition. *Philos Trans R Soc Lond B Biol Sci* 357: 1123–1136.
- Efron B, Tibshirani R (1993) An introduction to the bootstrap. New York: Chapman & Hall.
- Yeshurun Y, Carrasco M, Maloney LT (2008) Bias and sensitivity in two-interval forced choice procedures: Tests of the difference model. *Vision Research* 48: 1837–1851.
- Fechner GT (1860) *Elemente der Psychophysik*. Leipzig.