



Opinion Dynamics with Confirmation Bias

Armen E. Allahverdyan^{1*}, Aram Galstyan²

¹ Department of Theoretical Physics, Yerevan Physics Institute, Yerevan, Armenia, ² USC Information Sciences Institute, Marina del Rey, California, United States of America

Abstract

Background: Confirmation bias is the tendency to acquire or evaluate new information in a way that is consistent with one's preexisting beliefs. It is omnipresent in psychology, economics, and even scientific practices. Prior theoretical research of this phenomenon has mainly focused on its economic implications possibly missing its potential connections with broader notions of cognitive science.

Methodology/Principal Findings: We formulate a (non-Bayesian) model for revising subjective probabilistic opinion of a confirmationally-biased agent in the light of a persuasive opinion. The revision rule ensures that the agent does not react to persuasion that is either far from his current opinion or coincides with it. We demonstrate that the model accounts for the basic phenomenology of the social judgment theory, and allows to study various phenomena such as cognitive dissonance and boomerang effect. The model also displays the order of presentation effect—when consecutively exposed to two opinions, the preference is given to the last opinion (recency) or the first opinion (primacy) —and relates recency to confirmation bias. Finally, we study the model in the case of repeated persuasion and analyze its convergence properties.

Conclusions: The standard Bayesian approach to probabilistic opinion revision is inadequate for describing the observed phenomenology of persuasion process. The simple non-Bayesian model proposed here does agree with this phenomenology and is capable of reproducing a spectrum of effects observed in psychology: primacy-recency phenomenon, boomerang effect and cognitive dissonance. We point out several limitations of the model that should motivate its future development.

Citation: Allahverdyan AE, Galstyan A (2014) Opinion Dynamics with Confirmation Bias. PLoS ONE 9(7): e99557. doi:10.1371/journal.pone.0099557

Editor: Pablo Branas-Garza, Middlesex University London, United Kingdom

Received: December 8, 2013; **Accepted:** May 15, 2014; **Published:** July 9, 2014

Copyright: © 2014 Allahverdyan, Galstyan. 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 DARPA grant No. W911NF-12-1-0034 and AFOSR MURI grant No. FA9550-10-1-0569. 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.

* Email: armen.allahverdyan@gmail.com

Introduction

Confirmation bias is the tendency to acquire or process new information in a way that confirms one's preconceptions and avoids contradiction with prior beliefs [52]. Various manifestations of this bias have been reported in cognitive psychology [5,67], social psychology [24,54], politics [46] and (media) economics [31,51,57,73]. Recent evidence suggests that scientific practices too are susceptible to various forms of confirmation bias [12,38,43,44,52], even though the imperative of avoiding precisely this bias is frequently presented as one of the pillars of the scientific method.

Here we are interested in the opinion revision of an agent \mathcal{P} who is persuaded (or advised) by another agent \mathcal{Q} [10,13,52]. (Below we use the terms *opinion* and *belief* interchangeably.) We follow the known framework for representing uncertain opinions of both agents via the subjective probability theory [13]. Within this framework, the opinion of an agent about propositions (events) is described by probabilities that quantify his degree of confidence in the truth of these propositions [13]. As we argue in the next section, the standard Bayesian approach to opinion revision is inadequate for describing persuasion. Instead, here we study confirmationally-biased persuasion within the opinion combination approach developed in statistics; see [21,30] for reviews.

We suggest a set of conditions that model cognitive aspects of confirmation bias. Essentially, those conditions formalize the

intuition that the agent \mathcal{P} does not change his opinion if the persuasion is either far away or identical with his existing opinion [15,60]. We then propose a simple opinion revision rule that satisfies those conditions and is consistent with the ordinary probability theory. The rule consists of two elementary operations: *averaging* the initial opinion with the persuading opinion via linear combination, and then *projecting* it onto the initial opinion. The actual existence of these two operations has an experimental support [8,9,18,72,].

We demonstrate that the proposed revision rule is consistent with the *social judgment theory* [10], and reproduces the so called *change-discrepancy* relationship [10,35,40,45,69]. Furthermore, the well-studied *weighted average* approach [9,27] for opinion revision is shown to be a particular case of our model.

Our analysis of the revision rule also reveals novel effects. In particular, it is shown that within the proposed approach, the recency effect is related to confirmation bias. Also, repeated persuasions are shown to hold certain monotonicity features, but do not obey the law of diminishing returns. We also demonstrate that the rule reproduces several basic features of the *cognitive dissonance* phenomenon and predicts new scenarios of its emergence. Finally, the so called *boomerang (backfire)* effect can emerge as an extreme form of confirmation bias. The effect is given a straightforward mathematical description in qualitative agreement with experiments.

The rest of this paper is organized as follows. In the next section we introduce the problem setup and provide a brief survey of relevant work, specifically focusing on inadequacy of the standard Bayesian approach to opinion revision under persuasion. In the third section we define our axioms and introduce the confirmation-biased opinion revision rule. The fourth section relates our setup to the social judgment theory. Next two sections describe how our model accounts for two basic phenomena of experimental social psychology: opinion change versus discrepancy and the order of presentation effect. The seventh section shows how our model formalizes features of cognitive dissonance, followed by analysis of opinion change under repeated persuasion. Then we study the *boomerang effect*—the agent changes his opinion not towards the persuasion, but against it—as a particular case of our approach. We summarize and conclude in the last section.

The Set-Up and Previous Research

Consider two agents \mathcal{P} and \mathcal{Q} . They are given an uncertain quantity (random variable) \mathcal{X} with values $k=1, \dots, N$, e.g. $k=(rain, norain)$, if this is a weather forecast. \mathcal{X} constitutes the state of the world for \mathcal{P} and \mathcal{Q} . The opinions of the agents are quantified via probabilities

$$p = \{p_k\}_{k=1}^N \text{ and } q = \{q_k\}_{k=1}^N, \quad \sum_{k=1}^N p_k = \sum_{k=1}^N q_k = 1 \quad (1)$$

for \mathcal{P} and \mathcal{Q} respectively.

Let us now assume that \mathcal{P} is persuaded (or advised) by \mathcal{Q} . (Persuasion and advising are not completely equivalent [71]. However, in the context of our discussion it will be useful to employ both terms simultaneously stressing their common aspects.) Throughout this paper we assume that the state of the world does not change, and that the agents are aware of this fact. Hence, \mathcal{P} is going to change his opinion only under influence of the opinion of \mathcal{Q} , and not due to any additional knowledge about \mathcal{X} (For more details on this point see [3,41] and the second section of File S1.)

The normative standard for opinion revision is related to the Bayesian approach. Below we discuss the main elements of the Bayesian approach, and outline certain limitations that motivates the non-Bayesian revision rule suggested in this work.

Within the Bayesian approach, the agent \mathcal{P} treats his own probabilistic opinion $p = \{p_k\}_{k=1}^N$ as a prior, and the probabilistic opinion $q = \{q_k\}_{k=1}^N$ of \mathcal{Q} as an evidence [28,30,47]. Next, it is assumed that \mathcal{P} is endowed with conditional probability densities $\Pi(q|k)$, which statistically relate q to the world state k . Upon receiving the evidence from \mathcal{Q} , agent \mathcal{P} modifies his opinion from p_k to $p(k|q)$ via the Bayes rule:

$$p(k|q) = \Pi(q|k)p_k / \sum_{l=1}^N \Pi(q|l)p_l. \quad (2)$$

One issue with the Bayesian approach is that the assumption on the existence and availability of $\Pi(q|k)$ may be too strong [13,25,30]. Another issue is that existing empirical evidence suggests that people do not behave according to the Bayesian approach [13,61], e.g. they demonstrate the order of presentation effect, which is generally absent within the Bayesian framework.

In the context of persuasion, the Bayesian approach (2) has two additional (and more serious) drawbacks. To explain the first drawback, let us make a generic assumption that there is a unique

index \hat{k} for which $\Pi(q|k)$ is maximized as a function of k (for a given q): $\Pi(q|\hat{k}) > \Pi(q|k)$ for $\hat{k} \neq k$.

Now consider repeated application of (2), which corresponds to the usual practice of repeated persuasion under the same opinion q of \mathcal{Q} . The opinion of the agent then tends to be completely polarized, i.e. $p_{\hat{k}} \rightarrow 1$ and $p_k \rightarrow 0$ for $k \neq \hat{k}$. In the context of persuasion or advising, we would rather expect that under repeated persuasion the opinion of \mathcal{P} will converge to that of \mathcal{Q} .

The second issue is that, according to (2), \mathcal{P} will change his opinion even if he has the same opinion as \mathcal{Q} : $p = q$. This feature may not be realistic: we do not expect \mathcal{P} to change his opinion, if he is persuaded towards the same opinion he has already. This drawback of (2) was noted in [28]. (Ref. [28] offers a modification of the Bayesian approach that complies with this point, as shown in [28] on one particular example. However, that modification betrays the spirit of the normative Bayesianism, because it makes conditional probabilities depending on the prior probability.)

It is worthwhile to note that researchers have studied several aspects of confirmation bias by looking at certain deviations from the Bayes rule, e.g. when the conditional probability are available, but the agent does not apply the proper Bayes rule deviating from it in certain aspects [31,51,57,73]. One example of this is when the (functional) form of the conditional probability is changed depending on the evidence received or on the prior probabilities. Another example is when the agent does not employ the full available evidence and selects only the evidence that can potentially confirm his prior expectations [39,48,67]. More generally, one has to differentiate between two aspects of the confirmation bias that can be displayed either with respect to information acquiring, or information assimilation (or both) [52]. Our study will concentrate on information assimilation aspect; first, because this aspect is not studied sufficiently well, and second, because because it seems to be more directly linked to cognitive limitations [52]. We also stress that we focus on the belief revision, and not on actions an agent might perform based on those beliefs.

Opinion Revision Rule

We propose the following conditions that the opinion revision rule should satisfy.

1. The revised opinion \tilde{p}_k of \mathcal{P} is represented as

$$\tilde{p}_k = F[p_k, q_k] / \sum_{l=1}^N F[p_l, q_l], \quad (3)$$

where $F[x_1, x_2]$ is defined over $x_1 \in [0, \infty)$ and $x_2 \in [0, \infty)$. We enlarged the natural range $x_1 \in [0, 1]$ and $x_2 \in [0, 1]$, since below we plan to consider probabilities that are not necessarily normalized to 1. There are at least two reasons for doing so: First, experimental studies of opinion elicitation and revision use more general normalizations [8,9]. For example, if the probability is elicited in percents, the overall normalization is 100. Second, and more importantly, the axioms defining subjective (or logical) probabilities leave the overall normalization as a free parameter [22].

We require that $F[x_1, x_2]$ is continuous for $x_1 \in [0, \infty)$ and $x_2 \in [0, \infty)$ and infinitely differentiable for $x_1 \in (0, \infty)$ and $x_2 \in (0, \infty)$. Such (or similar) conditions are needed for features that are established for certain limiting values of the arguments of F (cf. (5, 6)) to hold approximately whenever the arguments are close to those limiting values. F can also depend on model parameters, as seen below.

Eq. (3) means that \mathcal{P} first evaluates the (non-normalized) weight $F[p_k, q_k]$ for the event k based *solely* on the values of p_k and q_k , and then applies overall normalization. A related feature of (3) is that it

is local: assume that $N \geq 3$ and only the probability q_1 is communicated by \mathcal{Q} to \mathcal{P} . This suffices for \mathcal{P} to revise his probability from p_1 to \tilde{p}_1 , and then adjust other probabilities via renormalization:

$$\begin{aligned} \tilde{p}_1 &= F[p_1, q_1] / (F[p_1, q_1] + 1 - p_1), \\ \tilde{p}_k &= p_k / (F[p_1, q_1] + 1 - p_1) \quad k \geq 2. \end{aligned} \tag{4}$$

Eq. (3) can be considered as a succession of such local processes.
2. If $p_k = 0$ for some k , then $\tilde{p}_k = 0$:

$$F[0, y] = 0. \tag{5}$$

The rationale of this condition is that if \mathcal{P} sets the probability of a certain event strictly to zero, then he sees logical (or factual) reasons for prohibiting the occurrence of this event. Hence \mathcal{P} is not going to change this zero probability under persuasion.

3. If $p_k q_k = 0$ for all k , then $\tilde{p}_k = p_k$: \mathcal{P} cannot be persuaded by \mathcal{Q} if their opinions have no overlap.

4. If \mathcal{Q} 's and \mathcal{P} 's opinions are identical, then the latter will not change his opinion: $\{p_k\}_{k=1}^N = \{q_k\}_{k=1}^N$ (for all k) leads to $\{\tilde{p}_k\}_{k=1}^N = \{p_k\}_{k=1}^N$. This can be written as

$$F[x, x] = x. \tag{6}$$

Conditions **3** and **4** are motivated by experimental results in social psychology, which state that people are not persuaded by opinions that are either very far, or very close to their initial opinion [10,17,69].

(Recall that we do not allow the uncertain quantity \mathcal{X} to change during the persuasion or advising. If such a change is allowed, **4** may not be natural as the following example shows. Assume that \mathcal{P} holds a probabilistic opinion (0.1,0.9) on a binary \mathcal{X} . Let \mathcal{P} learn that \mathcal{X} changed, but he does not know in which specific way it did. Now \mathcal{P} meets \mathcal{Q} who has the same opinion (0.1,0.9). Provided that \mathcal{Q} does not echo the opinion of \mathcal{P} , the agent \mathcal{P} should perhaps change his opinion by decreasing the first probability (0.1) towards a smaller value, because it is likely that \mathcal{X} changed in that direction.)

5. F is a homogeneous function of order one:

$$F[\gamma x, \gamma y] = \gamma F[x, y] \text{ for } \gamma \geq 0. \tag{7}$$

The rationale for this condition comes from the fact that (depending on the experimental situation) the subjective probability may be expressed not in normalization one (i.e. not with $\sum_{k=1}^N p_k = \sum_{k=1}^N q_k = 1$), but with a different overall normalization (e.g. $\sum_{k=1}^N p_k = \sum_{k=1}^N q_k = \gamma$) [8,9,22]; cf. **1**. In this light, (7) simply states that any choice of the overall normalization is consistent with the sought rule provided that it is the same for \mathcal{P} and \mathcal{Q} . Any rescaling of the overall normalization by the factor γ will rescale the non-normalized probability by the same factor γ ; cf. (7).

6. Now we assume that the opinion assimilation by \mathcal{P} consists of two sub-processes. Both are related to heuristics of human judgement.

6.1 \mathcal{P} combines his opinion *linearly* with the opinion of \mathcal{Q} [8,9,18,29,30]:

$$\hat{p}_k = \epsilon p_k + (1 - \epsilon) q_k, \quad 0 \leq \epsilon \leq 1, \tag{8}$$

where ϵ is a weight. Several mathematical interpretations of the weight ϵ were given in statistics, where (8) emerged as one of the basic rules of probabilistic opinion combination [16,29]; see section I of File S1. One interpretation suggested by this approach is that ϵ and $1 - \epsilon$ are the probabilities—from the subjective viewpoint of \mathcal{P} —for, respectively, p and q to be the true description of states of the world [29]: it is not known to \mathcal{P} which one of these probabilities (p or q) conveys a more accurate reflection of the world state. Then $\{\hat{p}_k\}_{k=1}^N$ is just the marginal probability for the states of the world. There is also an alternative (normative) way of deriving (8) from maximization of an average utility that under certain natural assumptions can be shown to be the (negated) average information loss [16]; see section I of File S1 for more details.

Several qualitative factors contribute to the subjective assessment of ϵ . For instance, one interpretation is to relate ϵ to credibility of \mathcal{Q} (as perceived by \mathcal{P}): more credible \mathcal{Q} leads to a larger $1 - \epsilon$ [18]. Several other factors might affect ϵ : egocentric attitude of \mathcal{P} that tends to discount opinions, simply because they do not belong to him; or the fact that \mathcal{P} has access to internal reasons for choosing his opinion, while he is not aware of the internal reasons of \mathcal{Q} *etc* [18]. Taking into account various factors that contribute to the interpretation of ϵ , we will treat it as a free model parameter.

6.2 Note that (8) does not satisfy conditions **2** and **3** above. We turn to the last ingredient of the sought rule, which, in particular, should achieve consistency with conditions **2** and **3**.

Toward this goal, we assume that \mathcal{P} projects the linearly combined opinion \hat{p} (see (8)) onto his original opinion p . Owing to (3), we write this transformation as

$$\tilde{p}_k = \phi[p_k, \hat{p}_k] / \sum_{l=1}^N \phi[p_l, \hat{p}_l], \tag{9}$$

where the function ϕ is to be determined.

The above projection operation relates to *trimming* [18,72], a human cognitive heuristics, where \mathcal{P} tends to neglect those aspects of \mathcal{Q} 's opinion that deviate from a certain reference. In the simplest case this reference will be the existing opinion of \mathcal{P} .

To make the projection process (more) objective, we shall assume that it commutes with the probabilistic revision: whenever

$$p'_k = \frac{\gamma_k p_k}{\sum_{l=1}^N \gamma_l p_l}, \quad \hat{p}'_k = \frac{\gamma_k \hat{p}_k}{\sum_{l=1}^N \gamma_l \hat{p}_l}, \quad 1 \leq k \leq N, \tag{10}$$

where $\gamma_k = \Pr(\dots|k) > 0$ are certain conditional probabilities, \tilde{p} is revised via the same rule (10):

$$\tilde{p}'_k \equiv \frac{\phi[p'_k, \hat{p}'_k]}{\sum_{l=1}^N \phi[p'_l, \hat{p}'_l]} = \frac{\gamma_k \phi[p_k, \hat{p}_k]}{\sum_{l=1}^N \gamma_l \phi[p_l, \hat{p}_l]}. \tag{11}$$

This feature means that the projection is consistent with probability theory: it does not matter whether (3) is applied before or after (10).

It is known that (9) together with (10, 11) selects a unique function [30]:

$$\phi[p_k, \hat{p}_k] = p_k^\mu \hat{p}_k^{1-\mu}, \quad 0 < \mu \leq 1 \tag{12}$$

where μ quantifies the projection strength: for $\mu = 1$ the projection is so strong that \mathcal{P} does not change his opinion at all (conservatism), while for $\mu \rightarrow 0$, \mathcal{P} fully accepts \hat{p} (provided that $p_k > 0$ for all k). (The above commutativity is formally valid also for $\mu \leq 0$ or $\mu > 1$, but both these cases are in conflict with (5).) In particular, $\epsilon \rightarrow 0$ and $\mu \rightarrow 0$ is a limiting case of a fully credulous agent that blindly follows persuasion provided that all his probabilities are non-zero. (For a sufficiently small μ , a small p_k is less effective in decreasing the final probability \tilde{p}_k ; see (12). This is because $p_k^\mu = e^{\mu \ln p_k}$ tends to zero for a fixed μ and $p_k \rightarrow 0$, while it tends to one for a fixed p_k and $\mu \rightarrow 0$. This interplay between $p_k \rightarrow 0$ and $\mu \rightarrow 0$ is not unnatural, since the initial opinion of a credulous agent is expected to be less relevant. The case of credulous agent is of an intrinsic interest and it does warrant further studies. However, since our main focus is confirmation bias, below we set $\mu = 1/2$ and analyze the opinion dynamics for varying ϵ .)

The final opinion revision rule reads from (12, 8, 9):

$$\tilde{p}_k = \frac{\sqrt{p_k[\epsilon p_k + (1-\epsilon)q_k]}}{\sum_{l=1}^N \sqrt{p_l[\epsilon p_l + (1-\epsilon)q_l]}}, \quad 0 \leq \epsilon \leq 1. \tag{13}$$

It is seen to satisfy conditions **1–5**.

(Note that the analogue of (11), $p'_k \propto \gamma_k p_k$, $q'_k \propto \gamma_k q_k$ does not leave invariant the linear function (8). First averaging, $\epsilon p_k + (1-\epsilon)q_k$ and then applying $p'_k \propto \gamma_k p_k$, $q'_k \propto \gamma_k q_k$ is equivalent to first applying the latter rules and then averaging with a different weight ϵ' . This is natural: once ϵ can be (in principle) interpreted as a probability it should also change under probabilistic revision process.)

The two processes were applied above in the specific order: first averaging (8), and then projection (9). We do not have any strong objective justifications for this order, although certain experiments on advising indicate on the order that led to (13) [72]. Thus, it is not excluded that the two sub-processes can be applied in the reverse order: first projection and then averaging. Then instead of (13) we get (3) with:

$$F[p_k, q_k; \epsilon, \mu] = \epsilon p_k + (1-\epsilon)p_k^\mu q_k^{1-\mu}, \quad 0 < \mu < 1. \tag{14}$$

Our analysis indicates that both revision rules (13) and (14) (taken with $\mu = 1/2$) produce qualitatively similar results. Hence, we focus on (13) for the remainder of this paper.

Returning to (1), we note that $k = x$ can be a continuous variable, if (for example) the forecast concerns the chance of having rain or the amount of rain. Then the respective probability densities are:

$$p(x) \text{ and } q(x), \quad \int dx p(x) = \int dx q(x) = 1. \tag{15}$$

Since the revision rule (13) is continuous and differentiable (in the sense defined after (3)), it supports a smooth transition between discrete probabilities and continuous and differentiable probability densities. In particular, (13) can be written directly for densities: for $p_k \simeq p(x_k)dx$ we obtain from (13)

$$\tilde{p}(x) = \frac{\sqrt{p(x)[\epsilon p(x) + (1-\epsilon)q(x)]}}{\int dx' \sqrt{p(x')[\epsilon p(x') + (1-\epsilon)q(x')]}}. \tag{16}$$

Social Judgment Theory and Gaussian Opinions

Opinion latitudes

Here we discuss our model in the context of the social judgment theory [59,10], and consider several basic scenarios of opinion change under the rule (16).

According to the social judgment theory, an agent who is exposed to persuasion perceives and evaluates the presented information by comparing it with his existing attitudes (opinions). The theory further postulates that an attitude is composed of three zones, or latitudes: *acceptance*, *non-commitment* and *rejection* [10,59]. The opinion that is most acceptable to \mathcal{P} , or the *anchor*, is located at the center of the latitude of acceptance. The theory states that persuasion does not change the opinion much, if the persuasive message is either very close to the anchor or falls within the latitude of rejection [10,59]. The social judgment theory is popular, but its quantitative modeling has been rather scarce. In particular, to the best of our knowledge, there has been no attempt to develop a consistent probabilistic framework for the theory. (The literature on the social judgment theory offers some formal mathematical expressions that could be fitted to experimental data [45]. There is also a more quantitative theory [34] whose content is briefly reminded in section III of File S1.)

Let us assume that $k = x$ is a continuous variable (cf. (15)) and that $p(x)$ and $q(x)$ are Gaussian with mean m_λ and dispersion v_λ ($\lambda = \mathcal{P}, \mathcal{Q}$):

$$p(x) = \frac{e^{-\frac{(x-m_P)^2}{2v_P}}}{\sqrt{2\pi v_P}}, \quad q(x) = \frac{e^{-\frac{(x-m_Q)^2}{2v_Q}}}{\sqrt{2\pi v_Q}}. \tag{17}$$

Effectively, Gaussian probabilistic opinions are produced in experiments, when the subjects are asked to generate an opinion with $\approx 95\%$ confidence in a certain interval [18]. Now we can identify the anchor with the most probable opinion m_λ , while v_λ^{-1} quantifies the opinion uncertainty.

The latitude of acceptance amounts to opinions not far from the anchor, while the latitude of rejection contains close-to-zero probability events, since \mathcal{P} does not change his opinion on them; recall point **2** from the previous section. One can also identify the three latitudes with appropriately chosen zones in the distribution. For instance, it is plausible to define the latitudes of acceptance and rejection by, respectively, the following formulas of the 3σ rule known in statistics

$$x \in [m_P - 2\sqrt{v_P}, m_\lambda + 2\sqrt{v_P}], \tag{18}$$

$$x \in (-\infty, m_P - 3\sqrt{v_P}] \cup [m_P + 3\sqrt{v_P}, \infty), \tag{19}$$

where the latitude of non-commitment contains whatever is left out from (18, 19). Recall that the latitudes of acceptance, non-commitment and rejection carry (respectively) 95.4, 4.3 and 0.3% of probability.

While the definitions (18, 19) are to some extent arbitrary, they work well with the rule (16), e.g. if the opinions of \mathcal{P} and \mathcal{Q} overlap only within their rejection latitudes, then neither of them can effectively change the opinion of another. Also, \mathcal{P} is persuaded most strongly, if the anchor of the persuasion falls into the non-commitment latitude of \mathcal{P} . This is seen below when studying change-discrepancy relations.

Weighted average of anchors

Next, we demonstrate that the main quantitative theory of persuasion and opinion change—the weighted average approach [9,27]—is a particular case of our model. We assume that the opinions $p(x)$ and $q(x)$ are given as

$$p(x) = f(x - m_P), \quad q(x) = g(x - m_Q), \quad (20)$$

$$\begin{aligned} f'(0) \equiv df/dx|_{x=0} &= g'(0) = 0, \\ f''(0) < 0, \quad g''(0) < 0, \end{aligned} \quad (21)$$

where both $f(x)$ and $g(x)$ have a unique maximum at $x = 0$. Hence $p(x)$ (resp. $q(x)$) has a single anchor (maximally probable opinion) m_P (resp. m_Q); see (17) for concrete examples.

If $|m_P - m_Q|$ is sufficiently small, $\tilde{p}(x)$ given by (20, 16) has a single anchor which is shifted towards that of $q(x)$; see Fig. 1(a). We now look for the maximum $m_{\tilde{P}}$ of $\tilde{p}(x)$ by using (20) in (16). We neglect factors of order $\mathcal{O}[(m_P - m_Q)^2/v_P]$ and $\mathcal{O}[(m_P - m_Q)^2/v_Q]$ and deduce:

$$m_{\tilde{P}} = (1 - \alpha_Q) m_P + \alpha_Q m_Q, \quad (22)$$

$$\alpha_Q \equiv \frac{(1 - \epsilon) \frac{|g''(0)|}{g(0)}}{(1 - \epsilon) \frac{|g''(0)|}{g(0)} + \frac{|f''(0)|}{f(0)} \left[2\epsilon \frac{f(0)}{g(0)} + 1 - \epsilon \right]}. \quad (23)$$

Eq. (22) is the main postulate of the weighted average approach; see [9,27] for reviews. Here α_Q and $1 - \alpha_Q$ are the weights of \mathcal{Q} and of \mathcal{P} , respectively. For the Gaussian case (17), we have

$$\alpha_Q = \left[1 + y \left(1 + \frac{2\epsilon\sqrt{y}}{1 - \epsilon} \right) \right]^{-1}, \quad y \equiv \frac{v_Q}{v_P}. \quad (24)$$

Furthermore, we have

$$\frac{\partial \alpha_Q}{\partial y} = -\alpha_Q^2 \left[1 + \frac{3\epsilon\sqrt{y}}{1 - \epsilon} \right], \quad \frac{\partial \alpha_Q}{\partial \epsilon} = -\frac{2y^{3/2}\alpha_Q^2}{(1 - \epsilon)^2}, \quad (25)$$

Thus, α_Q 's dependence on the involved parameters is intuitively correct: it increases with the confidence $1/v_Q$ of \mathcal{Q} , and decreases with the confidence $1/v_P$ of \mathcal{P} . Note also that α_Q decreases with ϵ .

Now let $p(x)$ and $q(x)$ (and hence $\tilde{p}(x)$) have the same maximum $m_P = m_Q$, but $v_P \approx v_Q$; see (17). Expanding (16, 17) over $v_P - v_Q$ and keeping the first-order term only we get

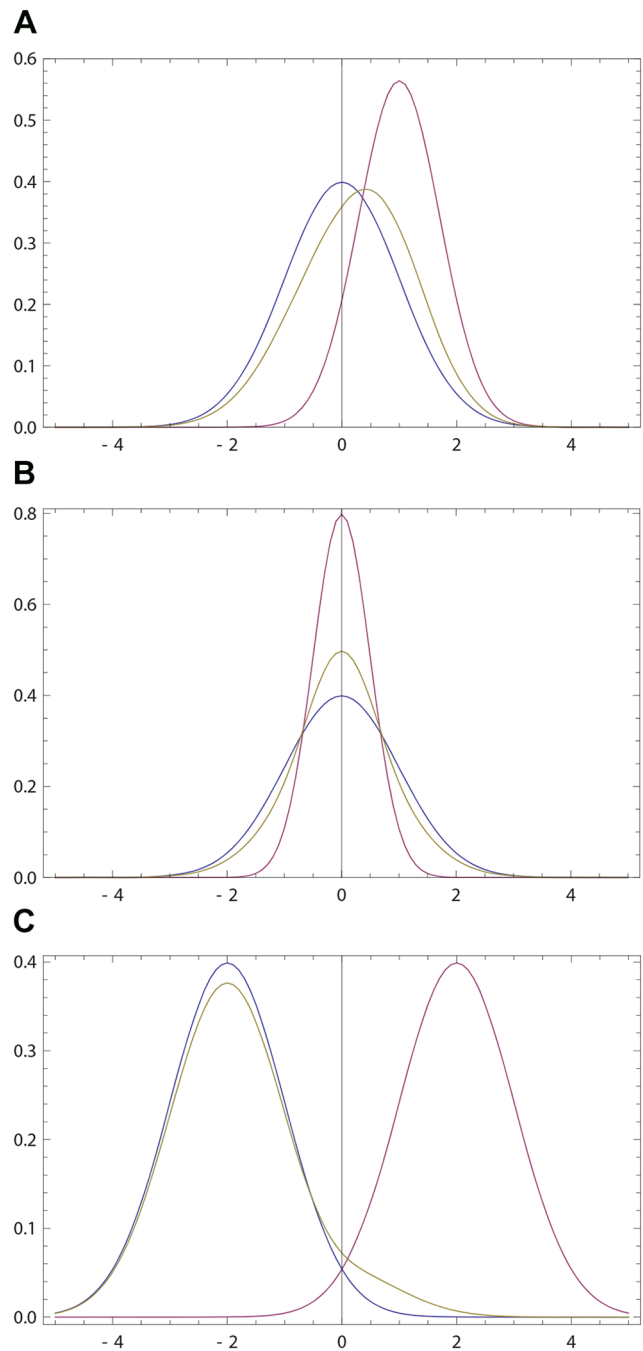


Figure 1. Opinions described via Gaussian densities (17). The initial opinion of \mathcal{P} is described by Gaussian probability density $p(x)$ (blue curve) centered at zero; see (17). The opinion of \mathcal{Q} amounts to Gaussian probability density $q(x)$ (purple curve) centered at a positive value. For all three figures continuous density $f(x)$ ($f = p, q, \tilde{p}$) were approximated by 100 points $\{f(x_k)\}_{k=1}^{100}$, $x_{k+1} - x_k = 0.1$. The resulting opinion $\tilde{p}(x)$ of \mathcal{P} is given by (16) with $\epsilon = 0.5$ (olive curve). (a) The opinion of \mathcal{P} moves towards that of \mathcal{Q} ; $m_P = 0$, $\sigma_P = 1$, $m_Q = 1$, $\sigma_Q = 0.5$. (b) The maximally probable opinion of \mathcal{P} is reinforced; $m_P = 0$, $\sigma_P = 1$, $m_Q = 0$, $\sigma_Q = 0.25$. (c) The change of the opinion of \mathcal{P} is relatively small provided that the Gaussian densities overlap only in the region of non-commitment; cf. (18), (19). Whenever the densities overlap only within the rejection range the difference between $p(x)$ and $\tilde{p}(x)$ is not visible by eyes. For example, if $p(x)$ and $q(x)$ are Gaussian with, respectively, $m_P = -3$, $m_Q = 3$, $v_P = v_Q = 1$, the Hellinger distance (see (30) for definition) $h[p, q] = 0.99$ is close to maximally far, while the opinion change is small: $h[p, \tilde{p}] = 3.48 \times 10^{-2}$. doi:10.1371/journal.pone.0099557.g001

$$v_{\tilde{\mathcal{P}}} = \frac{1-\epsilon}{2} v_{\mathcal{Q}} + \frac{1+\epsilon}{2} v_{\mathcal{P}}, \tag{26}$$

where $v_{\tilde{\mathcal{P}}}$ is the dispersion of (non-Gaussian) $\tilde{p}(x)$. Eq. (26) implies

$$(v_{\tilde{\mathcal{P}}} - v_{\mathcal{P}})(v_{\mathcal{Q}} - v_{\mathcal{P}}) = \frac{1-\epsilon}{2} (v_{\mathcal{Q}} - v_{\mathcal{P}})^2 \geq 0, \tag{27}$$

i.e. if $1/v_{\mathcal{Q}} > 1/v_{\mathcal{P}}$ (resp. $1/v_{\mathcal{Q}} < 1/v_{\mathcal{P}}$), the final opinion of \mathcal{P} becomes more (resp. less) narrow than his initial opinion. Fig. 1(b) shows that $(v_{\tilde{\mathcal{P}}} - v_{\mathcal{P}})(v_{\mathcal{Q}} - v_{\mathcal{P}}) \geq 0$ holds more generally.

Thus, the weighted average approach is a particular case of our model, where the agent \mathcal{P} is persuaded by a slightly different opinion. Note also that our model suggests a parameter structure of the weighted average approach.

Opinions and bump-densities

Gaussian densities (with three latitudes) do correspond to the phenomenology of social psychology. However, in certain scenarios one might need other forms of densities, e.g., when the probability is strictly zero outside of a finite support. Such opinions can be represented by bump-functions

$$X(x; b) = \mathcal{N}(b) \exp\left[\frac{b}{x^2 - 1}\right] \text{ for } |x| < 1 \tag{28}$$

$$= 0 \text{ for } x \leq -1 \text{ and } x \geq 1.$$

where $b > 0$ is a parameter, $\mathcal{N}(b)$ is the normalization and the support of the bump function was chosen to be $[-1, 1]$ for concreteness. The advantage of the bump function that is infinitely differentiable despite of having a finite support.

For sufficiently large b , $\chi(x; b)$ is close to a Gaussian, while for small b , $\chi(x; b)$ represents an opinion that is (nearly) homogeneous on the interval $[-1, 1]$; see Fig. 2. The opinion revision with bump densities follows to the general intuition of rule (16); see Fig. 2.

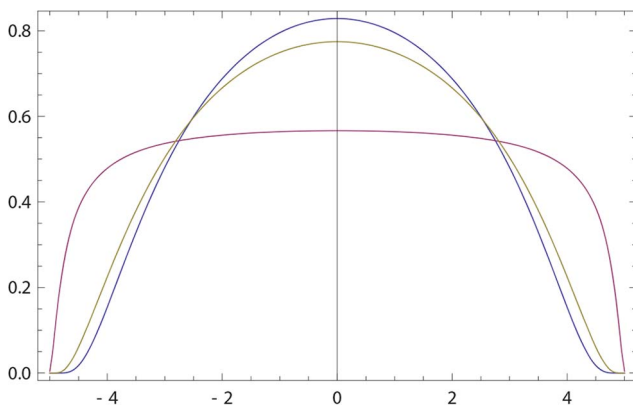


Figure 2. Opinions described via bump densities (28). Blue curve: the initial opinion of \mathcal{P} given by (28) with $b = 1$. Purple curve: the opinion of \mathcal{Q} described by (28) with $b = 0.001$. Olive curve: the resulting opinion of \mathcal{P} obtained via (16) with $\epsilon = 0.5$. doi:10.1371/journal.pone.0099557.g002

Opinion Change vs Discrepancy

One of extensively studied questions in social psychology is how the opinion change is related to the discrepancy between the initial opinion and the position conveyed by the persuasive message [10,35,40,45,69]. Initial studies suggested a linear relationship between discrepancy and the opinion change [35], which agreed with the prediction of the weighted average model. Indeed, (22) yields the following linear relationship between the change in the anchor and the initial opinion discrepancy of \mathcal{P} and \mathcal{Q} :

$$m_{\tilde{\mathcal{P}}} - m_{\mathcal{P}} = \alpha_{\mathcal{Q}} (m_{\mathcal{Q}} - m_{\mathcal{P}}). \tag{29}$$

However, consequent experiments revealed that the linear regime is restricted to small discrepancies only and that the actual behavior of the opinion change as a function of the discrepancy is non-monotonic: the opinion change reaches its maximal value at some discrepancy and decreases afterward [10,40,45,69].

To address this issue within our model, we need to define distance $h[p, q]$ between two probability densities $p(x)$ and $q(x)$. Several such distances are known and standardly employed [32]. Here we select the Hellinger distance (metric)

$$h[p, q] \equiv \frac{1}{\sqrt{2}} \left[\int dx [\sqrt{p(x)} - \sqrt{q(x)}]^2 \right]^{1/2}, \tag{30}$$

$$= \left[1 - \int dx \sqrt{p(x)q(x)} \right]^{1/2}. \tag{31}$$

Since $\sqrt{p(x)}$ is a unit vector in the ℓ_2 norm, Eq. (30) relates to the Euclidean (ℓ_2 -norm) distance. It is applicable to discrete probabilities by changing the integral in (30, 31) to sum. For Gaussian opinions (17) we obtain

$$h[p, q] = \left[1 - \left[\frac{(v_{\mathcal{Q}} v_{\mathcal{P}})^{1/2}}{\frac{v_{\mathcal{Q}} + v_{\mathcal{P}}}{2}} \right]^{1/2} e^{-\frac{(m_{\mathcal{Q}} - m_{\mathcal{P}})^2}{4(v_{\mathcal{Q}} + v_{\mathcal{P}})}} \right]^{1/2}. \tag{32}$$

A virtue of the Hellinger distance is that it is a measure of overlap between the two densities; see (31). We stress, however, that there are other well-known distances measures in statistics [32]. All results obtained below via the Hellinger distance will be checked with one additional metric, the total variation (ℓ_1 -norm distance):

$$\delta[p, q] = \frac{1}{2} \int dx |p(x) - q(x)|. \tag{33}$$

(To motivate the choice of (33), let us recall two important variational features of this distance [32]: (1) $\delta[p, q] = \max_{\Omega \in \mathcal{R}^1} |\int_{\Omega} dx (p(x) - q(x))|$. (2) Define two (generally dependent) random variables X, Y with joint probability density $g(x, y)$ such that $\int dx g(x, y) = q(y)$, $\int dy g(x, y) = p(x)$. Now it holds that $\delta[p, q] = \min[\Pr(X \neq Y)]$, where $\Pr(X \neq Y) = 1 - \int dx g(x, x)$, and the minimization is taken over all $g(x, y)$ with fixed marginals equal to $p(x)$ and $q(y)$, respectively.)

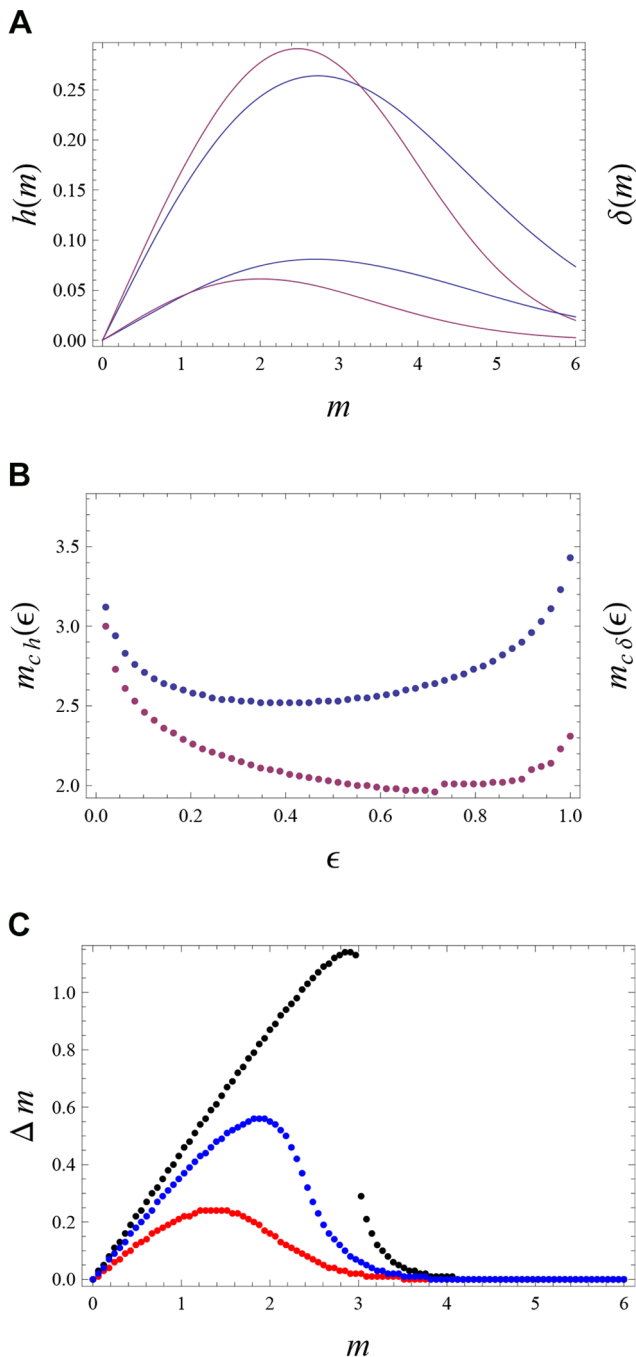


Figure 3. Opinion change versus discrepancy. (a) The opinion change is quantified via the Hellinger distance $h = h[p, \tilde{p}]$ between the old and new opinion of \mathcal{P} (blue curves); see (30) for the definition. For comparison we also include the total variance distance $\delta = \delta[p, \tilde{p}]$ (purple curves); see (33). These two distances are plotted versus the discrepancy $m = |m_P - m_Q|$. The initial opinion of the agent \mathcal{P} is Gaussian with $m_P = 0$ and $v_P = 1$; see (17). The opinion of \mathcal{Q} is Gaussian with $m_Q = m$ and $v_Q = 1$. Thus m quantifies the initial distance between the opinions of \mathcal{P} and \mathcal{Q} . The final opinion $\tilde{p}(x)$ is given by (13). Different curves correspond to different ϵ . Blue curves: $h(m) = h[p, \tilde{p}]$ for $\epsilon = 0.1$ (upper curve) and $\epsilon = 0.5$ (lower curve). Purple curves: $\delta(m) = \delta[p, \tilde{p}]$ for $\epsilon = 0.1$ (upper curve) and $\epsilon = 0.5$ (lower curve). The maximum of $h(m)$ ($\delta(m)$) is reached at m_{ch} (m_{δ}). (b) m_{ch} (m_{δ}) is the point where $h(m)$ ($\delta(m)$) achieves its maximum as a function of m . Blue points: $m_{ch}(\epsilon)$ versus ϵ for same parameters as in (a). $m_{ch}(\epsilon)$ grows both for $\epsilon \rightarrow 1$ and $\epsilon \rightarrow 0$, e.g. $m_{ch}(0.01) = 3.29972$, $m_{ch}(0.0001) = 4.53052$, $m_{ch}(0.9) = 2.94933$, $m_{ch}(0.999) = 4.12861$. Purple points: $m_{\delta}(\epsilon)$ versus ϵ

for same parameters as in (a). (c) The difference of the anchors (maximally probable values) $\Delta m = m_{\tilde{P}} - m_P$ versus $m_Q = m$ for the initial opinions of \mathcal{P} and \mathcal{Q} given by (17) under $m_P = 0$, $v_P = 1$, $m_Q = m$ and $v_Q = 1$. The final opinion $\tilde{p}(x)$ of \mathcal{P} (and its maximally probable value $m_{\tilde{P}}$) if found from (13) under $\epsilon = 0.1$ (black points), $\epsilon = 0.25$ (blue points) and $\epsilon = 0.5$ (red points).
doi:10.1371/journal.pone.0099557.g003

The opinion change is characterized by the Hellinger distance $h[p, \tilde{p}]$ between the initial and final opinion of \mathcal{P} , while the discrepancy is quantified by the Hellinger distance $h[p, q]$ between the initial opinion of \mathcal{P} and the persuading opinion. For concreteness we assume that the opinion strengths $1/v_P$ and $1/v_Q$ are fixed. Then $h[p, q]$ reduces to the distance $m = |m_P - m_Q|$ between the anchors (peaks of $p(x)$ and $q(x)$); see (32).

Fig. 3(a) shows that the change $h[p, \tilde{p}]$ is maximal at $m = m_{ch}$; it decreases for $m > m_{ch}$, since the densities of \mathcal{P} and \mathcal{Q} have a smaller overlap. The same behavior is shown by the total variation $\delta[p, \tilde{p}]$ that maximizes at $m = m_{\delta}$; see Fig. 3(a).

The dependence of m_{ch} (and of m_{δ}) on ϵ is also non-monotonic; Fig. 3(b). This is a new prediction of the model. Also, m_{ch} and m_{δ} are located within the latitude of non-commitment of \mathcal{P} (this statement does not apply to m_{ch} , when ϵ is close to 1 or 0); cf. (18, 19). This point agrees with experiments [10,69].

Note that experiments in social psychology are typically carried out by asking the subjects to express one preferred opinion under given experimental conditions [10,35,40,45,69]. It is this single opinion that is supposed to change under persuasion. It seems reasonable to relate this single opinion to the maximally probable one (anchor) in the probabilistic set-up. Thus, in addition to calculating distances, we show in Fig. 3(c) how the final anchor $m_{\tilde{P}}$ of \mathcal{P} deviates from his initial anchor m_P .

Fig. 3(c) shows that for $\epsilon > 0.25$, the behavior of $\Delta m = |m_{\tilde{P}} - m_P|$ as a function of $m = |m_P - m_Q|$ has an inverted-U shape, as expected. It is seen that Δm saturates to zero much faster compared to the distance $h[p, \tilde{p}]$. In other words, the full probability \tilde{p} keeps changing even when the anchor does not show any change; cf. Fig. 3(c) with Fig. 3(a).

A curious phenomenon occurs for a sufficiently small ϵ ; see Fig. 3(c) with $\epsilon = 0.1$. Here Δm drops suddenly to a small value when m passes certain critical point; Fig. 3(c). The mechanism behind this sudden change is as follows: when the main peak of $p(x)$ shifts towards m_Q , a second, sub-dominant peak of $\tilde{p}(x)$ appears at a value smaller than m_P . This second peak grows with m and at some critical value it overcomes the first peak, leading to a bistability region and an abrupt change of Δm . The latter arises due to a subtle interplay between the high credibility of \mathcal{Q} (as expressed by a relatively small value of ϵ) and sufficiently large discrepancy between \mathcal{P} and \mathcal{Q} (as expressed by a relatively large value of m). Recall, however, that the distance $h[p, \tilde{p}]$ calculated via the full probability does not show any abrupt change.

The abrupt change of Δm is widely discussed (and experimentally confirmed) in the attitude change literature; see [49] for a recent review. There the control variables for the attitude change—information and involvement [49]—differ from ϵ and m . However, one notes that the weight ϵ can be related to the involvement: more \mathcal{P} is involved into his existing attitude, larger is ϵ , while the discrepancy m connects to the (new) information contained in the persuasion ($m = 0$ naturally means zero information).

Let us finally consider a scenario where the change-discrepancy relationship is monotonic. It is realized for $m_P = m_Q$ (coinciding anchors), where the distance (32) between $p(x)$ and $q(x)$ is controlled by v_Q (for a fixed v_P). In this case, the change $h[p, \tilde{p}]$ is a

monotonic function of discrepancy $h[p,q]$: a larger discrepancy produces larger change. This example is interesting, but we are not aware of experiments that have studied the change-discrepancy relation in the case of two identical anchors.

Order of Presentation

Recency versus primacy

When an agent is consecutively presented with two persuasive opinions, his final opinion is sensitive to the order of presentation [10,13,25,34,35,50,52]. While the existence of this effect is largely established, its direction is a more convoluted matter. (Note that the order of presentation effect is not predicted by the Bayesian approach; see (2).) Some studies suggest that the first opinion matters more (primacy effect), whereas other studies advocate that the last interaction is more important (recency effect). While it is not completely clear which experimentally (un)controlled factors are responsible for primacy and recency, there is a widespread tendency of relating the primacy effect to confirmation bias [13,52]. This relation involves a qualitative argument that we scrutinize below.

We now define the order of presentation effect in our situation. The agent \mathcal{P} interacts first with \mathcal{Q} (with probability density $q(x)$), then with \mathcal{Q}' with probability density $q'(x)$. To ensure that we compare *only* the order of \mathcal{Q} and \mathcal{Q}' and not different magnitudes of influences coming from them, we take both interactions to have the same parameter $0 < \epsilon < 1$. Moreover, we make \mathcal{Q} and \mathcal{Q}' symmetric with respect to each other and with respect to \mathcal{P} , e.g. if $p(x)$, $q(x)$ and $q'(x)$ are given by (17) we assume

$$v_{\mathcal{Q}'} = v_{\mathcal{Q}}, \quad m_{\mathcal{Q}'} - m_{\mathcal{P}} = m_{\mathcal{P}} - m_{\mathcal{Q}}. \quad (34)$$

We would like to know whether the final opinion $p(x|q,q')$ of \mathcal{P} is closer to $q(x)$ (primacy) or to $q'(x)$ (recency).

In the present model (and for $0 \leq \epsilon < 1$), the final opinion $p(x|q,q')$ is *always* closer to the last opinion $q'(x)$, both in terms of maximally probable value and distance. In other words, the model unequivocally predicts the recency effect. In terms of the Hellinger distance (30)

$$h[p(x|q,q'),q'] < h[p(x|q,q'),q]. \quad (35)$$

See Fig. 4 for an example (In our model primacy effect exists in the boomerang regime $\epsilon > 1$; see below.)

To illustrate (35) analytically on a specific example, consider the following (binary) probabilistic opinion of \mathcal{P} , \mathcal{Q} and \mathcal{Q}'

$$p = (1/2, 1/2), \quad q = (0, 1), \quad q' = (1, 0). \quad (36)$$

\mathcal{P} is completely ignorant about the value of the binary variable, while \mathcal{Q} and \mathcal{Q}' are fully convinced in their opposite beliefs. If \mathcal{P} interacts first with \mathcal{Q} and then with \mathcal{Q}' (both interactions are given by (13) with $\epsilon = \frac{1}{2}$), the opinion of \mathcal{P} becomes $(0.52727, 0.47273)$. This is closer to the last opinion (that of \mathcal{Q}').

The predicted recency effect in our model seems rather counterintuitive. Indeed, since the first interaction shifts the opinion of \mathcal{P} towards that of \mathcal{Q} , one would think that the second

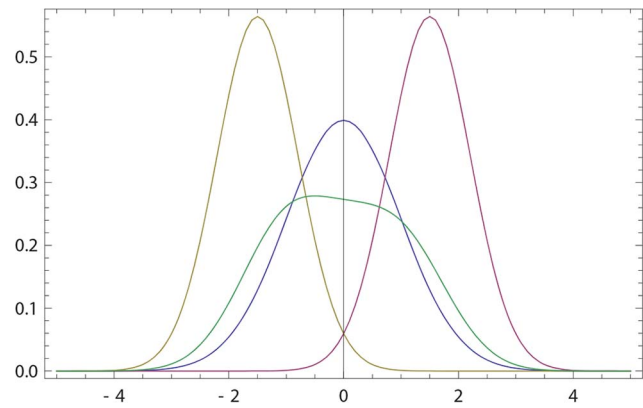


Figure 4. Order of presentation effect. Blue curve: The initial opinion of \mathcal{P} is described by Gaussian probability density $p(x)$ with $m_{\mathcal{P}} = 0$ and $v_{\mathcal{P}} = 1$; see (17). Purple (resp. olive) curve: the initial opinion of \mathcal{Q} (resp. \mathcal{Q}') are given by (17) with $m_{\mathcal{Q}} = 1.5$ (resp. $m_{\mathcal{Q}'} = -1.5$) and $v_{\mathcal{Q}} = 0.5$ (resp. $v_{\mathcal{Q}'} = 0.5$). Green curve: the resulting opinion of \mathcal{P} after interacting first with \mathcal{Q} and then with \mathcal{Q}' . Both interactions use $\epsilon = 0.5$. The final opinion of \mathcal{P} is inclined to the most recent opinion (that of \mathcal{Q}') both with respect to its maximally probable value and distance. The final opinion of \mathcal{P} has a larger width than the initial one. doi:10.1371/journal.pone.0099557.g004

interaction with \mathcal{Q}' should influence \mathcal{P} 's opinion less, due to a smaller overlap between the opinions of \mathcal{Q}' and \mathcal{P} before the second interaction. In fact, this is the standard argument that relates primacy effect to the confirmation bias [13,52]: the first interaction shapes the opinion of \mathcal{P} and makes him confirmation-biased against the second opinion. This argument does not apply to the present model due to the following reason: even though the first interaction shifts \mathcal{P} 's anchor towards \mathcal{Q} 's opinion, it also deforms the shape of the opinion; see Fig. 1(a). And the deformation produced by our revision rule happens to favor the second interaction more.

To get a deeper understanding of the recency effect, let us expand (13) for small $\eta \equiv 1 - \epsilon$:

$$\begin{aligned} \tilde{p}_k &= p_k + \frac{\eta}{2}(q_k - p_k) \\ &+ \frac{\eta^2}{8}(p_k - 1) \sum_l \frac{(q_l - p_l)^2}{p_l} + \mathcal{O}[\eta^3]. \end{aligned} \quad (37)$$

If now \mathcal{P} interacts with an agent \mathcal{Q}' having opinion q' , the resulting opinion $p(q,q')$ reads from (37):

$$\begin{aligned} p_k(q,q') &= p_k \\ &+ \frac{\eta}{2}(q_k - p_k) + \frac{\eta^2}{8}(p_k - 1) \sum_l \frac{(q_l - p_l)^2}{p_l} \\ &+ \frac{\eta}{2}(q'_k - p_k) + \frac{\eta^2}{8}(p_k - 1) \sum_l \frac{(q'_l - p_l)^2}{p_l} \end{aligned}$$

$$+ \frac{\eta^2}{4}(p_k - q_k) + \mathcal{O}[\eta^3]. \quad (38)$$

Hence in this limit $p_k(q, q') - p_k(q', q)$ depends only on $q'_k - q_k$ (and not e.g. on $q_{l \neq k}$):

$$p_k(q, q') - p_k(q', q) = \eta^2 [q'_k - q_k] / 4 + \mathcal{O}[\eta^3]. \quad (39)$$

It is seen that the more probable persuasive opinion (e.g. the opinion of \mathcal{Q}' if $q'_k > q_k$) changes the opinion of \mathcal{P} if it comes later. This implies the recency effect. Indeed, due to symmetry conditions for checking the order of presentation effect we can also look at $h[p(q, q'), q] - h[p(q', q), q]$. Using (39) we get for this quantity: $\frac{\eta^2}{16h[p(q', q), q]} \sum_k [q_k - q'_k] \sqrt{q_k/p_k} > 0$, again due to symmetry conditions.

Note that this argument on recency directly extends to more general situations, where the agent is exposed to different opinions multiple times. For instance, consider an exposure sequence $q q' q'$ and its reverse $q' q' q$. It can be shown that the model predicts a recency effect in this scenario as well. For this case, we get instead of (39): $p_k(q, q') - p_k(q', q) = \eta^2 [q_k - q'_k] + \mathcal{O}[\eta^3]$.

Note that the primacy-recency effect is only one (though important!) instance of contextual and non-commutative phenomena in psychology; see [11,66] and references therein. Hence in section IV of File S1 we study a related (though somewhat less interesting) order of presentation effect, while below we discuss our findings in the context of experimental results.

Experimental studies of order of presentation effect

We now discuss our findings in this section in the context of experimental results on primacy and recency. The latter can be roughly divided into several groups: persuasion tasks [10,50], symbol recalling [70], inference tasks [34], and impression formation [7,9]. In all those situations one generally observes both primacy and recency, though in different proportions and under different conditions [34]. Generally, the recency effect is observed whenever the *retention* time (the time between the last stimulus and the data taking) is short. If this time is sufficiently long, however, the recency effect changes to the primacy effect [10,50,62,70]. The general interpretation of these results is that there are two different processes involved, which operate on different time-scales. These processes can be conventionally related to short-term and long-term memory [70], with the primacy effect related to the long-term memory. In our model the longer time process is absent. Hence, it is natural that we see only the recency effect. The prevalence of recency effects is also seen in inference tasks, where the analogue of the short retention time is the incremental (step-by-step) opinion revision strategy [34].

At this point, let us remind the importance of symmetry conditions [such as (34)] for observing a *genuine* order of presentation effect. Indeed, several experimental studies—in particular those on impression formation—suggest that the order of presentation exists *due* to different conditions in the first versus the second interaction [7,10,34,68]. (In our context, this means different parameters ϵ and ϵ' for each interaction). For instance, Refs. [7,10] argue that the primacy effect is frequently caused by attention decrement (the first action/interaction gets more

attention); see also [68] in this context. This effect is trivially described by our model, if we assume ϵ to be sufficiently smaller than ϵ' . In related experiments, it was shown that if the attention devoted to two interactions is balanced, the recency effect results [33], which is consistent with the prediction of our model.

At the same time, in another interesting study based on subjective probability revision, where the authors had taken special measures for minimizing the attention decrement, the results indicated a primacy effect [55].

We close this section by underlining the advantages and drawbacks of the present model concerning the primacy-recency effect: the main advantage is that it demonstrates the recency effect and shows that the well-known argument on relating confirmation bias to primacy does not hold generally. The main drawback is that the model does not involve processes that are supposedly responsible for the experimentally observed interplay between recency and primacy. In the concluding section we discuss possible extensions of the model that can account for this interplay.

Cognitive Dissonance

Consider an agent whose opinion probability density has two peaks on widely separated events. Such a density—with the most probable opinion being different from the average—is indicative of cognitive dissonance, where the agent believes in mutually conflicting things [10,26].

The main qualitative scenario for the emergence of cognitive dissonance is when an agent—who initially holds a probabilistic opinion with a single peak—is exposed to a conflicting information coming from a sufficiently credible source [10,26]. We now describe this scenario quantitatively.

Consider again the opinion revision model (16, 17), and assume that $|m_{\mathcal{P}} - m_{\mathcal{Q}}|$ is neither very large nor very small (in both these cases no serious opinion change is expected), $v_{\mathcal{Q}}/v_{\mathcal{P}} < 1$ (self-assured persuasive opinion) and $0 < \epsilon < 1$. In this case, we get two peaks (anchors) for the final density $\tilde{p}(x)$. The first peak is very close to the initial anchor of $p(x)$, while the second closer to the anchor of $q(x)$; see Fig. 5(a). Thus, persuasion from \mathcal{Q} whose opinion is sufficiently narrow and is centered sufficiently close (but not too close) to \mathcal{P} 's initial anchor leads to cognitive dissonance: \mathcal{P} holds simultaneously two different anchors, the old one and the one induced by \mathcal{Q} .

There are 3 options for reducing cognitive dissonance:

(i) Increase ϵ making it closer to 1, i.e. making \mathcal{Q} less credible; see Fig. 5(b).

(ii) Decrease the width of the initial opinion of \mathcal{P} .

(iii) Decrease ϵ making \mathcal{Q} more credible. In this last case, the second peak of $\tilde{p}(x)$ (the one close to the anchor of \mathcal{Q}) will be dominant; see Fig. 5(c).

To understand the mechanism of the cognitive dissonance as described by this model, let us start from (1) and assume for simplicity that the opinion of \mathcal{Q} is certain: $q_k = 0$ for $k \neq l$ and $q_l = 1$. We get from (13):

$$\tilde{p}_k = \frac{p_k}{1 - p_l + p_l \sqrt{\epsilon + (1 - \epsilon)p_l^{-1}}} \text{ for } k \neq l, \quad (40)$$

$$\tilde{p}_l = \frac{p_l \sqrt{\epsilon + (1 - \epsilon)p_l^{-1}}}{1 - p_l + p_l \sqrt{\epsilon + (1 - \epsilon)p_l^{-1}}}. \quad (41)$$

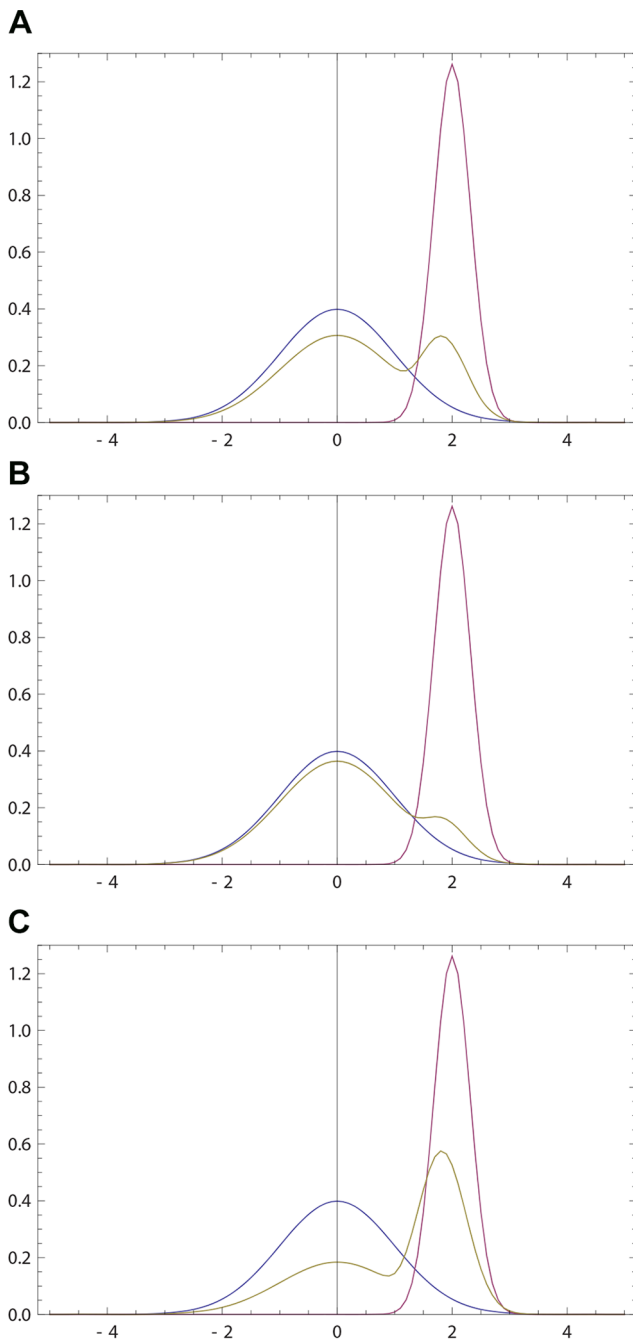


Figure 5. Cognitive dissonance. (a) Blue (resp. purple) curve: the initial opinion of agent \mathcal{P} (resp. \mathcal{Q}) described by probability density $p(x)$ (resp. $q(x)$). Olive curve: the final opinion $\bar{p}(x)$ of \mathcal{P} as given by (16) with $\epsilon=0.35$. Here $p(x)$ and $q(x)$ are defined by (17) with $m_P=0$, $v_P=1$, $m_Q=2$, $v_Q=0.1$. The final opinion develops two peaks of comparable height (cognitive dissonance). (b) Avoiding the cognitive dissonance due to a larger $\epsilon=0.75$: the second peak is much smaller (other parameters are those of (a)). (c) Avoiding the cognitive dissonance due to a smaller $\epsilon=0.05$: the first peak is much smaller (other parameters are those of (a)).
doi:10.1371/journal.pone.0099557.g005

Now $\bar{p}_l/p_l > 1 > \bar{p}_k/p_k$, where $k \neq l$; hence even if l was on the tail of $\{p_k\}_{k=1}^N$, it is possible to make it a local (or even the global) maximum of $\{\bar{p}_k\}_{k=1}^N$ provided that ϵ is not close to 1.

The existence of at least two widely different probable opinions is only one aspect of cognitive dissonance [10,26]. Another aspect (sometimes called Freud-Festinger’s law) is that people tend to avoid cognitive dissonance: if in their action they choose one of the two options (i.e. one of two peaks of the subjective probability), they re-write the history of their opinion revision so that the chosen option becomes the most probable one [10,26]. This aspect of cognitive dissonance found applications in economics and decision making [2,73]. The above points (i)–(iii) provide concrete scenarios for a such re-writing.

Repeated Persuasion

Here we analyze the opinion dynamics under repeated persuasion attempts. Our motivation for studying this problem is that repeated exposure to the same opinion is generally believed to be more persuasive than a single exposure.

Under certain conditions ($p_k q_k \neq 0$, for all k and $1 > \epsilon > 0$) we show that the target opinion converges to the persuading opinion after sufficient number of repetition. Below we also examine how exactly this convergence takes place.

Assume that \mathcal{P} revises his opinion repeatedly with the same opinion of \mathcal{Q} . Eq. (13) implies ($1 \leq k \leq N$)

$$p_k^{[n+1]} = \frac{\sqrt{p_k^{[n]} [\epsilon p_k^{[n]} + (1-\epsilon)q_k]}}{\sum_{l=1}^N \sqrt{p_l^{[n]} [\epsilon p_l^{[n]} + (1-\epsilon)q_l]}} \tag{42}$$

where $1 > \epsilon > 0$, and $n=1,2,\dots$ is the discrete time. For simplicity, we assume

$$p_k^{[1]} \equiv p_k > 0, q_k > 0 \text{ for } 1 \leq k \leq N. \tag{43}$$

Eq. (42) admits only one fixed point $q = \{q_k\}_{k=1}^N$. Section VI of File S1 shows that for any convex, $\frac{d^2 f(y)}{dy^2} \geq 0$, function $f(y)$ one has

$$\Phi[p^{[n+1]}; q] \leq \Phi[p^{[n]}; q], \tag{44}$$

$$\Phi[p; q] \equiv \sum_{k=1}^N q_k f(p_k/q_k). \tag{45}$$

Hence $\Phi[p; q]$ is a Lyapunov function of (42). Since $\Phi[p; q]$ is a convex function of p , $\Phi[p; q] \geq f(1) = \Phi[q; q]$ and $f(1)$ is the unique global minimum of $\Phi[p; q]$. Section VI of File S1 shows that the equality sign in (45) holds only for $p^{[n+1]} = p^{[n]}$. Thus $\Phi[p^{[n]}; q]$ monotonically decays to $f(1) = \Phi[q; q]$ showing that the fixed point q is globally stable. More generally, the convergence reads: $p_k^{[n]} \rightarrow \zeta [p_k^{[1]}] q_k / \sum_{l=1}^N \zeta [p_l^{[1]}] q_l$, where $\zeta(x > 0) = 1$ and $\zeta(0) = 0$.

To illustrate (44, 45), one can take $f(y) = -\sqrt{y}$. Then (44) amounts to decaying Hellinger distance (30). Many other reasonable measures of distance are obtained under various choices of f . For instance, $f(y) = |y - 1|$ amounts to decaying total variation distance (33), while $f(y) = -\ln y$ leads to the decaying relative entropy (Kullback-Leibler entropy).

As expected, $0 < \epsilon < 1$ influences the convergence time. We checked that this time is an increasing function of ϵ , as expected.

In section VI of File S1 we also show that the convergence to the fixed point respects the Le Chatelier principle known in thermodynamics [4]: the probabilities of those events that are overestimated from the viewpoint of \mathcal{Q} (i.e. $p_k^{[1]} > q_k$) tend to decay in the discrete time. Likewise, probabilities of the underestimated events (i.e. $p_k^{[1]} < q_k$) increase in time.

Let us consider the Hellinger distance $h_n = h[p^{[n+1]}, p^{[n]}]$ between two consecutive opinions of \mathcal{P} evolving as in (42). It is now possible that

$$\max_{1 \leq n < \infty} [h_n] = h_m \neq h_1, \tag{46}$$

i.e. the largest change of the opinion of \mathcal{P} comes not from the first, but from one of intermediate persuasions. A simple example of this situation is realized for $N = 3$, an initial probability vector $p = (0.98, 0.01, 0.01)$ and $q = (0.01, 0.01, 0.98)$ in (43). We then apply (42) under $\epsilon = 0.5$. The consecutive Hellinger distances read $h_1 = 0.1456 < h_2 = 0.1567 > h_3 = 0.1295 > h_4 \dots$. Hence the second persuasion changes the opinion more than others. For this to hold, the initial opinion p of \mathcal{P} has to be far from the opinion q of \mathcal{Q} . Otherwise, we get a more expected behavior $h_1 > h_2 > h_3 > h_4 \dots$ meaning that the first persuasion leads to the largest change.

(The message of (46) is confirmed by using the discrete version $\delta[p, q] = \frac{1}{2} \sum_k |p_k - q_k|$ of the distance (33). Define $\delta_n = \delta[p^{[n+1]}, p^{[n]}]$. Then with $p = (0.98, 0.01, 0.01)$ and $q = (0.01, 0.01, 0.98)$ we get $\delta_1 = 0.0834$, $\delta_2 = 0.1636$, $\delta_3 = 0.1717$, $\delta_4 = 0.1444$.)

We conclude by stressing that while repeated persuasions drive the opinion to its fixed point monotonically in the number of repetitions, it is generally not true that the first persuasion causes the largest opinion change, i.e. the law of diminishing returns does not hold. To obtain the largest opinion change, one should carefully choose the number of repetitions.

Finally, note that the framework of (42) can be applied to studying mutual persuasion (consensus reaching). This is described in Section VII of File S1; see also [23] in this context.

Boomerang (Backfire) Effect

Definition of the effect

The *boomerang* or *backfire* effect refers to the empirical observation that sometimes persuasion yields the opposite effect: the persuaded agent \mathcal{P} moves his opinion away from the opinion of the persuading agent, \mathcal{Q} , i.e. he enforces his old opinion [53,58,64,69]. Early literature on social psychology proposed that the boomerang effect may be due to persuading opinions placed in the latitude of rejection [69], but this was not confirmed experimentally [40].

Experimental studies indicate that the boomerang effect is frequently related with opinion formation in an affective state, where there are emotional reasons for (not) changing the opinion. For example, a clear evidence of the boomerang effect is observed when the persuasion contains insulting language [1]. Another interesting example is when the subjects had already announced their opinion publicly, and were not only reluctant to change it (as for the usual conservatism), but even enforced it on the light of the contrary evidence [64] (in these experiments, the subjects who did not make their opinion public behaved without the boomerang effect). A similar situation is realized for voters who decided to support a certain candidate. After hearing that the candidate is criticized, the voters display a boomerang response to this criticism and thereby increase their support [53,58].

Opinion revision rule

We now suggest a simple modification of our model that accounts for the basic phenomenology of the boomerang effect.

Recall our discussion (around (8)) of various psychological and social factors that can contribute into the weight ϵ . In particular, increasing the credibility of \mathcal{Q} leads to a larger $1 - \epsilon$. Imagine now that \mathcal{Q} has such a low credibility that

$$\epsilon > 1. \tag{47}$$

Recall that $\epsilon = 1$ means a special point, where no change of opinion of \mathcal{P} is possible whatsoever; cf. (13).

After analytical continuation of (13) for $\epsilon > 1$, the opinion revision rule reads

$$\tilde{p}_k = \frac{\sqrt{p_k |\epsilon p_k + (1 - \epsilon) q_k|}}{\sum_{l=1}^N \sqrt{p_l |\epsilon p_l + (1 - \epsilon) q_l|}}, \tag{48}$$

with obvious generalization to probability densities. The absolute values in (48) are necessary to ensure the positivity of probabilities.

It is possible to *derive* (rather simply postulate) (48). Toward this end, let us return to the point **6.1** and (8). During the opinion combination step, \mathcal{P} forms $\epsilon p_k + (1 - \epsilon) q_k$ which in view of $\epsilon > 1$ can take negative values and hence is a signed measure. Signed measures have all formal features of probability besides positivity [6,14,19,65]; see section V of File S1 for details. There is no generally accepted probabilistic interpretation of signed measures, but in section V of File S1 we make a step towards such an interpretation. There we propose to look at a signed measure as a partial expectation value defined via joint probability of the world's states and certain hidden degrees of freedom (e.g. emotional states). After plausible assumptions, the marginal probability of the world's states is deduced to be

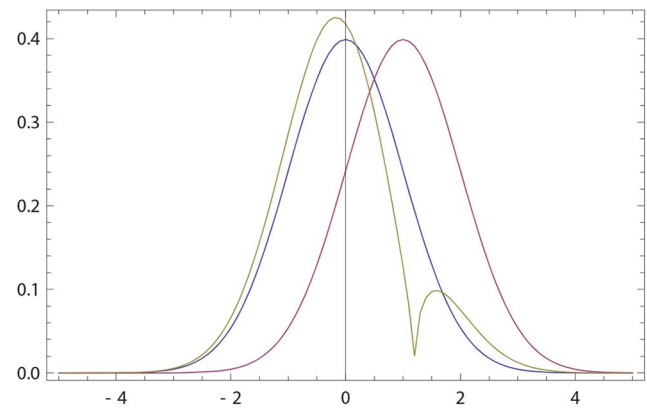


Figure 6. Opinion change in the boomerang regime. Blue (resp. purple) curve: the initial opinion of agent \mathcal{P} (resp. \mathcal{Q}) described by probability density $p(x)$ (resp. $q(x)$). Olive curve: the final opinion $\tilde{p}(x)$ of \mathcal{P} given by (16) with $\epsilon = 2$. Here $p(x)$ and $q(x)$ are given by (17) with $m_P = 0$ and $v_P = m_Q = v_Q = 1$. The anchor (maximally probable opinion) of \mathcal{P} not only moves away from the anchor of \mathcal{Q} ; but it is also enhanced: the (biggest) peak of $\tilde{p}(x)$ is larger than that of $p(x)$. The second (smaller) peak of $\tilde{p}(x)$ arises because the initial probability of \mathcal{P} located to the right from the anchor m_Q of \mathcal{Q} , moves away from m_Q ; $\tilde{p}(x)$ gets a local minimum close to m_Q . doi:10.1371/journal.pone.0099557.g006

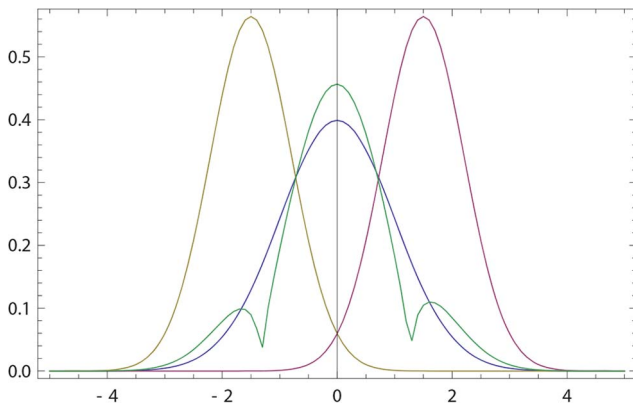


Figure 7. Order of presentation effect in the boomerang regime. The same as in Fig. 4 but for $\epsilon = 1.5$ (boomerang regime). Now the final opinion of \mathcal{P} is inclined to the first opinion (that of \mathcal{Q}) with respect to the distance. The initial maximally probable opinion of \mathcal{P} is still maximally probable. Moreover, its probability has increased and the width around it has decreased. The final opinion has 3 peaks. doi:10.1371/journal.pone.0099557.g007

$$\hat{p}_k = |\epsilon p_k + (1 - \epsilon)q_k| / \sum_{l=1}^N |\epsilon p_l + (1 - \epsilon)q_l|, \quad (49)$$

We obtain (48) after applying (9, 10) to (49).

Scenarios of opinion change

According to (47, 48) those opinions of \mathcal{P} which are within the overlap between p and q (i.e. $p_k q_k \neq 0$) get their probability decreased if $p_k/q_k \approx (\epsilon - 1)/\epsilon < 1$, i.e. if the initial p_k was already smaller than q_k . In this sense, \mathcal{P} moves his opinion away from that of \mathcal{Q} . Hence for continuous densities $p(x)$ and $q(x)$ there will be a point x_0 , where $\hat{p}(x_0)$ is close to 0. This point is seen in Figs. 6 and 7.

Fig. 6 illustrates the shape of $\tilde{p}(x)$ produced by (48) for initially Gaussian opinions (17) of \mathcal{P} and \mathcal{Q} . It is seen that \mathcal{P} 's anchor moves away from \mathcal{Q} 's anchor, while the width of $\tilde{p}(x)$ around the anchor is more narrow than that of $p(x)$; cf. with Fig. 4. To illustrate these points analytically, we return to (29, 24, 24) that for $v_{\mathcal{P}} \approx v_{\mathcal{Q}}$ and $m_{\mathcal{P}} \approx m_{\mathcal{Q}}$ predict $m_{\tilde{\mathcal{P}}} - m_{\mathcal{P}} = \frac{1 - \epsilon}{2}(m_{\mathcal{Q}} - m_{\mathcal{P}})$: for $\epsilon > 1$, \mathcal{P} 's anchor drifts away from \mathcal{Q} 's anchor.

Likewise, whenever the two anchors are equal, $m_{\mathcal{P}} = m_{\mathcal{Q}}$, inequality (27) is reversed in the boomerang regime (47).

Let us now consider the impact of the presentation order under this settings. We saw that for $0 < \epsilon < 1$ the model predicts recency effect. For $1 \lesssim \epsilon$ we expect the recency effect is still effective as implied by the argument (39). However, the situation changes drastically for ϵ sufficiently larger than 1, as indicated in Fig. 7. Now the primacy effect dominates, i.e. instead of (35) we get the opposite inequality. Fig. 7 also shows that interaction with two contradicting opinions (in the boomerang regime) enforces the initial anchor of \mathcal{P} .

To understand the primacy-recency effect analytically, consider the example (36), and recall that \mathcal{P} interacts first with \mathcal{Q} and then with \mathcal{Q} with the same parameter ϵ . The resulting opinion $p(q, q')$ of \mathcal{P} reads:

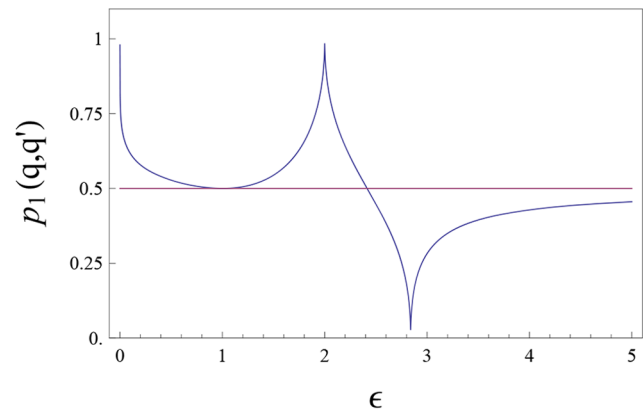


Figure 8. Illustration of the order of presentation effect in the boomerang regime. $p_1(q, q') = \frac{g(\epsilon)}{g(\epsilon) + 1}$ given by (50, 51) versus ϵ . doi:10.1371/journal.pone.0099557.g008

$$p(q, q') = \left(\frac{g(\epsilon)}{g(\epsilon) + 1}, \frac{1}{g(\epsilon) + 1} \right), \quad (50)$$

$$g(\epsilon) = \sqrt{\left| \frac{\sqrt{\epsilon + (1 - \epsilon)\sqrt{2 - \epsilon}}}{\sqrt{\epsilon(2 - \epsilon)}} \right|} \quad (51)$$

Fig. 8 shows how $p_1(q, q') = \frac{g(\epsilon)}{g(\epsilon) + 1}$ behaves as a function of ϵ . The recency effect holds for $\epsilon < 2 + \sqrt{2}$; for $\epsilon > 2 + \sqrt{2}$ we get primacy. Similar results are obtained for initially Gaussian opinions.

Thus, in the present model, the primacy effect (relevance of the first opinion) can be related to the boomerang effect.

We now examine the emergence of cognitive dissonance in the boomerang regime $\epsilon > 1$. Our results indicate that in this regime the agent is more susceptible to cognitive dissonance; cf. Fig. 6 with Figs. 1. The mechanism of the increased susceptibility is explained in Fig. 6: \mathcal{P} 's opinion splits easier, since the probability mass moves away (in different directions) from the anchor of \mathcal{Q} .

Let us now assume that \mathcal{P} repeatedly interacts with the same opinion of \mathcal{Q} [cf. (42)]:

$$p^{[n+1]}(x) = \frac{\sqrt{p^{[n]}(x) |\epsilon p^{[n]}(x) + (1 - \epsilon)q(x)|}}{\int dx' \sqrt{p^{[n]}(x') |\epsilon p^{[n]}(x') + (1 - \epsilon)q(x')|}}, \quad (52)$$

where $n = 1, 2, \dots$ is the discrete time. Starting from initially Gaussian opinion, \mathcal{P} develops two well-separated peaks, which is another manifestation of cognitive dissonance: the smaller peak moves towards the anchor of \mathcal{Q} and finally places itself within the acceptance latitude of \mathcal{Q} , where the larger peak becomes more narrow and drifts away from $q(x)$; see Fig. 9. After many iterations ($\approx 10^3$ for parameters of Fig. 9) the larger peak places itself within the rejection latitude of \mathcal{Q} , at which point $p^{[n]}(x)$ stops changing (stationary opinion). The above scenario suggests that in the boomerang regime there is a finite probability that the target agent

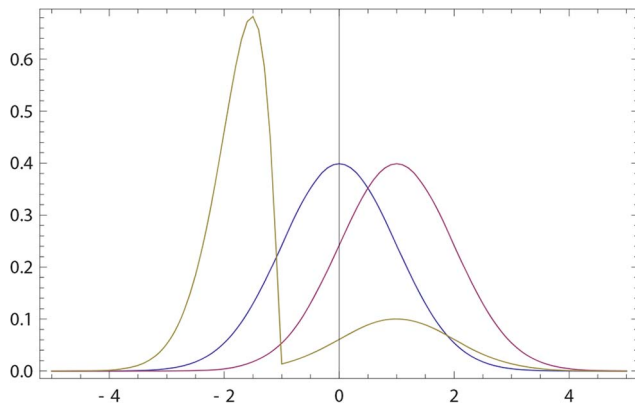


Figure 9. Repeated persuasion in the boomerang regime. Blue (resp. purple) curve: the initial opinion of agent \mathcal{P} (resp. \mathcal{Q}) described by probability density $p(x)$ (resp. $q(x)$) as given by (17) with $m_{\mathcal{P}}=0$, $v_{\mathcal{P}}=v_{\mathcal{Q}}=1$, $m_{\mathcal{Q}}=1$. Olive curve: the opinion of \mathcal{P} after 50 iterations (52) with $\epsilon=2$.

doi:10.1371/journal.pone.0099557.g009

will eventually be persuaded after repeated exposure to the same opinion.

Let us mention an experimental work that is relevant to our discussion above. Ref. [58] carried out experiments with subjects displaying boomerang effect, where each subject was exposed to sufficiently many different (but still similar) persuasive opinions. It was found that, sooner or later, the subjects exit the boomerang regime, i.e. they start to follow the persuasion [58]. Our set-up is somewhat different in that the subject (\mathcal{P}) is repeatedly exposed to the same persuading opinion. Modulo this difference, our conclusion is similar to the experimental finding: the agent starts following the persuasion with a certain probability.

Discussion

We presented a new model for opinion revision in the presence of confirmation bias. The model has three inputs: the subjective probabilistic opinions of the target agent \mathcal{P} and a persuading (advising) agent \mathcal{Q} , and the weight of \mathcal{Q} as perceived by \mathcal{P} .

The basic idea of the opinion revision rule is that no opinion change is expected if the persuasion is either too far or too close to the already existing opinion [15,36,60]. The opinion revision rule is not Bayesian, because the standard Bayesian approach does not apply to processes of persuasion and advising; see the second section for more details.

The model accounts for several key empirical observations reported in social psychology and quantitatively interpreted within the social judgment theory. In particular, the model allows to formalize the concept of opinion latitudes, explains the structure of the weighted average approach to opinion formation, and relates the initial discrepancy (between the opinions of \mathcal{P} and \mathcal{Q}) to the magnitude of the opinion change (shown by \mathcal{P}). In all these cases our model extends and clarifies previous empiric results, e.g. it elucidates the difference between monotonic and non-monotonic change-discrepancy relations, identifies conditions under which the opinion change is sudden, as well as provides a deeper perspective on the weighted average approach.

New effects predicted by the model are summarized as follows.

(i) For the order of presentation set-up (and outside of the boomerang regime) the model displays recency effect. We suggested that the standard argument that relates confirmation

bias to the primacy effect does not work in this model. In this context we recall a widespread viewpoint that *both* recency and primacy relate to (normative) irrationality; see e.g. [13]. However, the information which came later is generally more relevant for predicting future. Hence recency can be more rational than primacy.

In many experimental set-ups the recency changes to primacy upon increasing the retention time; see e.g. [70]. Our model demonstrates the primacy effect only in the boomerang regime (i.e. only in the special case). Hence, in future it needs to be extended by involving additional mechanisms, e.g. those related to “long-term memory” processes which could be responsible for the above experimental fact. Recall in this context there are several other theoretical approaches that address the primacy-recency difference [11,34,42,56,66].

(ii) The model can be used to describe the phenomenon of cognitive dissonance and to formalize the main scenario of its emergence.

(iii) Repeated persuasions display several features implying monotonous change of the target opinion towards the persuading opinion. However, the opinion changes do not obey the law of diminishing returns, or in other words, the first persuasion is not always leads to the largest change. These findings may contribute to better understanding the widespread use of repeated persuasions.

(iv) We proposed that the boomerang effect is related to the limit of this model, where the credibility of persuasion is (very) low. A straightforward implementation of this assumption led us to a revision rule that does describe several key observational features of the boomerang effect and predicts new ones; e.g. that in the boomerang regime the agent can be prone to primacy effect and to cognitive dissonance. There are, however, several open problems with the opinion revision rule in the boomerang regime. They should motivate future developments of this model. One problem concerns relations of the revision rule with signed measures that at a preliminary level were outlined in section V of File S1. Another problem is that the revision rule in the boomerang regime (and only there) is not completely smooth, since it includes the function $|x|$, whose second derivative is singular. We do hope to clarify these points in future.

In this paper we restricted ourselves by studying few (two or three) interacting agents with opinions described via subjective probabilities. However, these probabilities can also represent an ensemble of agents each one having a fixed (single) opinion, a useful viewpoint on subjective probabilities advocated in Ref. [37]. In future we plan to explore this point and also address the opinion dynamics for collectives of agents. This last aspect was recently extensively studied via methods of statistical physics; see [20,63] for reviews.

Supporting Information

File S1
(PDF)

Acknowledgments

We thank Seth Frey for useful remarks and suggestions.

Author Contributions

Conceived and designed the experiments: AA AG. Performed the experiments: AA. Analyzed the data: AA AG. Wrote the paper: AA.

References

- Abelson RP (1986). Beliefs are like possessions. *Journal for the Theory of Social Behaviour* 16: 223–250.
- Akerlof G, Dickens WT (1982). The economic consequences of cognitive dissonance. *Amer. Econ. Rev.* 72: 307–319.
- Alchourrón CE, Gärdenfors G, Makinson D (1985). On the logic of theory change. *J. Symb. Logic* 50: 510–530.
- Allahverdyan AE, Galstyan A (2011). Le Chatelier principle in replicator dynamics. *Physical Review E* 84: 041117
- Allakhverdov VM, Gershkovich VA (2010). Does consciousness exist? in what sense? *Integrative Psychological and Behavioral Science* 44: 340–347.
- Allen EH (1976). Negative Probabilities and the Uses of Signed Probability Theory. *Philosophy of Science* 43: 53–70.
- Anderson NN (1965). Primacy effect in personality impression formation using generalized order effect paradigm. *Journal of Personality and Social Psychology* 2: 1–9.
- Anderson NH (1971). Integration theory and attitude change. *Psychological Review* 78: 171–206.
- Anderson NN (1981). *Foundations of information integration theory*. Academic Press, New York.
- Aronson E (2007). *The Social Animal*. Palgrave Macmillan, 10th revised edition.
- Atmanspacher H, Roemer H (2012). Order effects in sequential measurements of non-commuting psychological observables. *Journal of Mathematical Psychology* 56: 274280.
- Austerweil JL, Griffiths TL (2011). Seeking confirmation is rational for deterministic hypotheses. *Cognitive Science* 35: 499–526
- Baron J (2008). *Thinking and deciding*. Cambridge University Press, Cambridge.
- Bartlett MS (1945). Negative probability. *Mathematical Proceedings of the Cambridge Philosophical Society* 41: 71–73.
- beim Graben P (2006). Pragmatic information in dynamic semantics. *Mind and Matter*, 4: 169–193.
- Bernardo JM (1979). Expected information as expected utility. *Annals of Statistics* 7: 686–690.
- Bochner S, Insko CA (1966). Communicator discrepancy, source credibility, and opinion change. *Journal of Personality and Social Psychology*, 4: 133–140.
- Bonaccio S, Dalal RS (2006). Advice taking and decision-making: An integrative literature review, and implications for the organizational sciences. *Organizational Behavior and Human Decision Processes* 101: 127151.
- Burgin M, Meissner G (2012). Negative probabilities in financial modeling. *Wilmott Magazine* 58: 6065.
- Castellano C, Fortunato S, Loreto V (2009). Statistical physics of social dynamics. *Reviews of Modern Physics* 81: 591–646.
- Clemen RT, Winkler RL (1999). Combining probability distributions from experts in risk analysis. *Risk Analysis* 19: 187–203.
- Cox RT (1946). Probability, Frequency and Reasonable Expectation. *American Journal of Physics* 14: 1–13.
- Curtis JP, Smith FT (2008). *Mathematical Models of Persuasion*. American Conference on Applied Mathematics (MATH '08), Harvard, Massachusetts: 60–65.
- Darley JM, Gross PH (1983). A hypothesis-confirming bias in labeling effects. *Journal of Personality and Social Psychology* 44: 20–33.
- Diaconis P, Zabel SL (1982). Updating subjective probability. *Journal of the American Statistical Association*, 77: 822–830.
- Festinger L (1957). *A Theory of Cognitive Dissonance*. Stanford University Press, Stanford, CA.
- Fink EL, Kaplowitz SA, Bauer CL (1983). Positional discrepancy, psychological discrepancy, and attitude change: Experimental tests of some mathematical models. *Communication Monographs* 50: 413–430.
- French S (1980). Updating of Belief in the Light of Someone Else's Opinion. *Journal of the Royal Statistical Society, Series A (General)*: 143 43–48.
- Genest C, McConway KJ (1990). Allocating the weights in the linear opinion pool. *Journal of Forecasting* 9: 53–73.
- Genest C, Zidek JV (1986). Combining probability distributions: A critique and an annotated bibliography. *Statistical Science* 1: 114–135.
- Gentzkow M, Shapiro JM (2006). Media bias and reputation. *Journal of Political Economy* 114: 280–316.
- Gibbs AL, Su FE (2002). On choosing and bounding probability metrics. *International statistical review* 70: 419–435.
- Hendrick C, Costantini AF (1970). Effects of varying trait inconsistency and response requirements on the primacy effect in impression formation. *Journal of Personality and Social Psychology* 15: 158–164.
- Hogarth RM, Einhorn HJ (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, 24: 1–55.
- Hovland CI (editor) (1957). *The order of presentation in persuasion*. Yale University Press, New Haven.
- Huhns MN, Singh MP (1997). Ontologies for agents. *IEEE Internet Computing*, 1: 81–83.
- Jaynes ET (1968). Prior Probabilities. *IEEE Trans. Syst. Science & Cyb*: 4 227–256.
- Jeng M (2005). A selected history of expectation bias in physics. *American Journal of Physics*, 74: 578–582.
- Jones M, Sugden R (2001). Positive Confirmation Bias In The Acquisition of Information. *Theory and Decision* 50: 59–99.
- Kaplowitz SA, Fink EL (1997). Message discrepancy and persuasion. *Progress in Communication Sciences XIII*. Edited by Barnett GA, Foster FJ. Ablex Publishing Corporation, Greenwich, Connecticut.
- Keller AM, Winslett M (1985). On the use of extended relational model to handle changing incomplete information. *IEEE Trans. Software Engineering SE-11*: 620–633.
- Khrennikov A (2006). Quantum-like brain: Interference of minds. *BioSystems* 84: 225–241.
- Klayman J, Ha YW (1987). Confirmation, disconfirmation, and information. *Psychological Review* 94: 211–228.
- Koehler JJ (1993). The influence of prior beliefs on scientific judgments of evidence quality. *Organizational Behavior and Human Decision Processes* 56: 28–55.
- Laroche M (1977). A model of attitude change in groups following a persuasive communication: An attempt at formalizing research findings. *Behavioral Science* 22: 246–257.
- Lazarsfeld PF, Berelson B, Gaudet H (1944). *The People's Choice. How the Voter Makes up his Mind in Presidential Campaign*. Columbia University Press, New York.
- Lindley DV, Tversky A, Brown RV (1979). On the reconciliation of probability assessments. *J. R. Statist. Soc. A* 142: 146–156.
- Lord C, Lepper MR, Ross L (1979). Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence. *Journal of Personality and Social Psychology* 37: 2098–2122.
- van der Maas HLJ, Kolstein R, van der Pligt J (2003). Sudden Transitions in Attitudes. *Sociological Methods & Research* 32: 125–152.
- Miller N, Campbell DT (1959). Recency and primacy in persuasion as a function of the timing of speeches and measurements. *The Journal of Abnormal and Social Psychology* 59: 1.
- Mullainathan S, Shleifer A (2005). The Market of News. *The American Economic Review* 95: 1031–1053.
- Nickerson RS (1998). Confirmation bias: a ubiquitous phenomenon in many guises. *Review of General Psychology*, 2: 175–220.
- Nyhan B, Reifler J (2010). When corrections fail: The persistence of political misperceptions. *Political Behavior* 32: 303–330.
- Oskamp (1965). Overconfidence in case-study judgments. *Journal of Consulting Psychology*, 29: 261–265.
- Peterson CR, DuCharme WM (1967). A primacy effect in subjective probability revision. *Journal of Experimental Psychology* 73: 6165.
- Pothos EM, Busemeyer JR (2009). A quantum probability explanation for violations of “rational” decision theory. *Proceedings of the Royal Society B* 276: 2171–2178.
- Rabin M, Schrag JL (1999). First Impressions Matter: A Model of Confirmatory Bias. *The Quarterly Journal of Economics* 114: 37.
- Redlawsk DP, Civettini AJW, Emmerson KM (2010). The Affective Tipping Point: Do Motivated Reasoners Ever “Get It”. *Political Psychology* 31: 563–593.
- Social Psychology: Handbook of Basic Principles* (2007). Edited by Kruglanski AW, Higgins EW. The Guilford Press, New York.
- Schreider YA (1970). *On the semantic characteristics of information*. edited by Saracevic T. *Introduction to Information Science*. Bowker, New York: 24–32.
- Tversky A, Kahneman D (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131.
- Stewart RH (1965). Effect of continuous responding on the order effect in personality impression formation. *Journal of Personality and Social Psychology* 1: 161–165.
- Stauffer D (2013). A biased review of sociophysics. *Journal of Statistical Physics* 151: 9–20.
- Sutherland S (1992). *Irrationality: The Enemy Within*. London, Constable.
- Szekely GJ (2005). Half of a coin: negative probabilities. *Wilmott Magazine* 50: 66–68.
- Trueblood JS, Busemeyer JR (2011). A quantum probability account of order effects in inference. *Cognitive science* 35: 1518–1552.
- Wason PC (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology*, 12: 129–140.
- Webster DM, Richter L, Kruglanski AW (1996). On leaping to conclusions when feeling tired: Mental fatigue effects on impression primacy. *Journal of Experimental Social Psychology*, 32: 181–195.
- Whittaker JO (1963). Opinion change as a function of communication-attitude discrepancy. *Psychological Reports* 13: 763–772.
- Wright AA, Santiago HC, Sands SF, Kendrick DF, Cook RG (1985). Memory processing of serial lists by pigeons, monkeys, and people. *Science* 229: 287–289.
- Yaniv I (2004). Receiving other people's advice: Influence and benefit. *Organizational Behavior and Human Decision Processes* 93: 113.
- Yaniv I (1997). Weighting and Trimming: Heuristics for Aggregating Judgments under Uncertainty. *Organizational Behavior and Human Decision Processes* 69: 237249.
- Yariv L (2002). I'll see it when I believe it - A Simple Model of Cognitive Consistency. *Cowles Foundation Discussion Paper # 1352*.