First regression
To provide a first evidence for possible treatment differences, we will run the following regression for each round separately:

$$|P_i - P_i^{Bayes}| = N_i + \varepsilon_i$$

where P_i and P_i^{Bayes} are the subject's stated and computed bayesian beliefs respectively. N_i is a dummy equal to one if the subject is assigned to the Narrative treatment.

Regression to test Bayesian behaviour

Following Charness and Dave (2017) and Kieren and Weber (2025).

Say culprits are either Marco (M state) or Andrea (A state).

Consider the i_{th} person reporting the probability of the M state at each point in a sequence of 10 rounds. If she believes that both states are initially equally likely ex ante should report a 0.5 probability for the M state as an initial prior. Further, a Bayesian would calculate the probability of the M state after round t as (assuming she disregards non-informative signals)

$$P_t^{\text{Bayes}} = P(\text{M} \mid z_t)^{\text{Bayes}} = \frac{\theta^{z_t}}{\theta^{z_t} + (1-\gamma-\theta)^{z_t}}, \quad z_t = m_t - a_t$$

where $m_t(a_t)$ denotes the number of infos towards Marco (Andrea) that have been drawn as of round t, the proportion of Marco's infos in the M state is $\theta = 0.45$, and the proportion of non-informative signals is $\gamma = 0.25$. Note that a Bayesian is indifferent regarding the order of the draws, since only the difference (z_t) is relevant. Thus, the natural log of the odds ratio for a Bayesian is given by

$$\pi_t = \ln\left(\frac{P(\text{M} \mid z_t)^{\text{Bayes}}}{P(\text{A} \mid z_t)^{\text{Bayes}}}\right) = \ln\left(\frac{\theta}{1-\gamma-\theta}\right) \times z_t = 0.4055 \times z_t \underset{>}{\overset{<}{}} 0,$$

and it is linear in info received; that is, the Bayesian log-odds ratio is updated by $\pm 0.4055 \times z_t$ after each new draw.
We next note that first-differencing both sides of equation yields:

$$\Delta\pi_t = \pi_t - \pi_{t-1} = \ln\left(\frac{\theta}{1-\gamma-\theta}\right) \times \Delta z_t = 0.4055 \times \Delta z_t,$$

where $\Delta\pi_t \in \{-0.4055, 0, 0.4055\}$ and $\Delta z_t \in \{-1, 0, 1\}$. In contrast to Charness and Dave (2017) and as in Kieren and Weber (2025) we have 0 as a possible value as we have the uninformative signals.
Finally, to isolate a 'confirming' information condition, we construct the following dummy variables:

$$C_t^M = \begin{cases} 1 & \text{if } \pi_{t-1} > 0 \text{ and } s_t = m \\ 0 & \text{otherwise} \end{cases}$$

$$C_t^A = \begin{cases} 1 & \text{if } \pi_{t-1} < 0 \text{ and } s_t = a \\ 0 & \text{otherwise} \end{cases}$$

where s_t is the info drawn in the current round. These variables measure whether a Bayesian receiver would view a received signal as confirming a belief. For example, if a Bayesian believed that the M state of the world was more likely to be in play (i.e., $\pi_{t-1} > 0$) for a given draw t and she receives a Marco signal ($s_t = m$), then we say the belief of a M state of the world is confirmed ($C_t^M = 1$). The same holds if that Bayesian believed that the A state of the world was more likely to be in play ($\pi_{t-1} < 0$) and she received an Andrea signal ($s_t = a$), in which case we would have that $C_t^A = 1$.

Moreover, we should take into account cases where null information is drawn. Let

$$C_t^{n+} = \begin{cases} 1 & \text{if } \pi_{t-1} > 0 \text{ and } s_t = n \\ 0 & \text{otherwise} \end{cases}$$

$$C_t^{n-} = \begin{cases} 1 & \text{if } \pi_{t-1} < 0 \text{ and } s_t = n \\ 0 & \text{otherwise} \end{cases}$$

be dummy variables that measure whether a Bayesian would view a null signal as confirming a belief.

Given the change in log-odds ($\Delta\pi_t$), and the dummies above, consider the following regression for a Bayesian:

$$\pi_t = \rho\pi_{t-1} + \beta\Delta z_t + \delta_1 C_t^M + \delta_2 C_t^A + \delta_3 C_t^{n+} + \delta_4 C_t^{n-} + \varepsilon_t$$

because a Bayesian would not be subject to either the conservatism/overreaction heuristic nor would she place any additional weight on confirming or uninformative signals, it must be the case that the coefficients satisfy

$$\rho = 1, \quad \beta = \ln\left(\frac{\theta}{1-\lambda-\theta}\right), \quad \delta_i = 0, \quad i = 1, 2, 3, 4$$

Actual Estimation

If you want to estimate this for a subject that is not Bayesian, within a round, the natural logarithm of an individual subject's odds ratio (the analog of π_t), based on her stated probability at each round t, that is $P_{it} = P_{it}(\mathrm{M} \mid z_t)$, is:

$$\lambda_{it} = \ln(\Lambda_{it}) = \ln\left(\frac{P_{it}(\mathrm{M} \mid z_t)}{1 - P_{it}(\mathrm{M} \mid z_t)}\right) \tag{1}$$

and may differ from π_t. Moreover, you need to either truncate data to lie in the [0.01, 0.99] interval, or hardcode beliefs equal to 0 or 1 to 0.01 and 0.99 respectively, so that λ_{it} is always defined.

As before, construct the dummy variables

$$C_{it}^{M} = \begin{cases} 1 & \text{if } \lambda_{it-1} > 0 \text{ and } s_t = m \\ 0 & \text{otherwise} \end{cases}$$

$$C_{it}^{A} = \begin{cases} 1 & \text{if } \lambda_{it-1} < 0 \text{ and } s_t = a \\ 0 & \text{otherwise} \end{cases}$$

$$C_{it}^{n+} = \begin{cases} 1 & \text{if } \lambda_{it-1} > 0 \text{ and } s_t = n \\ 0 & \text{otherwise} \end{cases}$$

$$C_{it}^{n-} = \begin{cases} 1 & \text{if } \lambda_{it-1} < 0 \text{ and } s_t = n \\ 0 & \text{otherwise} \end{cases}$$

Clearly if people deviate from Bayesian behavior due to conservatism (over-reaction), one would expect to see log-odds for people that are consistently smaller (larger) than $\pm 0.4055 \times z_t$. If people do weight evidence that confirms a previously held belief, or weight null information, then it would be the case that the dummies would predict log-odds; thus the regression for a subject, analogous to the one for a Bayesian, would be:

$$\lambda_{it} = \rho\lambda_{it-1} + \beta\Delta z_t + \delta_1 C_{it}^{M} + \delta_2 C_{it}^{A} + \delta_3 C_{it}^{n+} + \delta_4 C_{it}^{n-} + \varepsilon_{it}$$

clustering errors at the subject level, with a test of Bayesian behavior being that the estimate for ρ equals 1, the estimate for β equals $\ln\left(\frac{\theta}{1-\gamma-\theta}\right)$, and that the coefficients on the dummies C_{it}^{M}, C_{it}^{A}, C_{it}^{n+} and C_{it}^{n-}, namely, δ_i, $i = 1, 2, 3, 4$ be zero since a Bayesian would not care whether a particular signal (draw) confirmed a belief or is non-informative. Conservatism would manifest with an estimate of β less than $\ln\left(\frac{\theta}{1-\gamma-\theta}\right)$ whereas use of the over-reaction heuristic would imply an estimate of β greater than $\ln\left(\frac{\theta}{1-\gamma-\theta}\right)$. If people are affected by confirmation bias in that they place extra weight on a signal that confirms beliefs, then it would be the case that the estimates of δ_i, $i = 1, 2$ would be non-zero. If people use the uninformative signal to confirm they beliefs, then the estimates of δ_i, $i = 3, 4$ would be non-zero as well. In summary, a test of Bayesian behavior would be:

$$H_0 : \hat{\rho} = 1 \quad \text{and} \quad \hat{\beta} = 0.4055 \quad \text{and} \quad \hat{\delta}_i = 0, \quad i = 1, 2, 3, 4.$$

We will do this test for both treatments separately.

References

Charness, G., & Dave, C. (2017). Confirmation bias with motivated beliefs. *Games and Economic Behavior*, *104*, 1-23. Retrieved from `https://www.sciencedirect.com/science/article/pii/S0899825617300416` doi: https://doi.org/10.1016/j.geb.2017.02.015

Kieren, P., & Weber, M. (2025). Expectation formation under uninformative signals. *Management Science*, *0*(0), null. Retrieved from `https://doi.org/10.1287/mnsc.2023.03367` doi: 10.1287/mnsc.2023.03367