Sequential choice effects predict prevalence-induced concept change

Robert C. Wilson^{1,2,*}

¹ Department of Psychology, University of Arizona

² Cognitive Science Program, University of Arizona

* email: bob@arizona.edu

Abstract

Prevalence-induced concept change is the process by which the definition of a concept can shift as the prevalence of that concept changes. In a series of beautiful experiments, Levari and colleagues (2018) demonstrated that prevalence-induced concept change occurs across a range of situations, from color perception to threat assessment to ethical judgment. Here we show that in all of these cases the extent of prevalence-induced concept change is accounted for by sequential effects in which previous stimuli and responses bias the upcoming categorization. Crucially, these sequential effects are always active, even when the prevalence is does not change. Thus, by modeling these effects, we were able to predict the magnitude of prevalence-induced concept change based on participants' behavior when the prevalence is fixed. This suggests that sequential effects are the mechanism by which prevalence-induced concept change occurs.

Keywords: Prevalence-induced concept change, sequential effects, choice kernel, judgment, decision making

Introduction

Prevalence-induced concept change is a recently described phenomenon in which the definition of concepts can shift in response to changes in the prevalence of exemplars on either side of a concept boundary (Levari et al. 2018). For example, when asked to determine whether an ambiguously colored dot is purple or blue, people are more likely to say it is blue when the frequency of objectively blue dots is low. Similar effects occur when assessing the threat level of a face or the ethics of a research proposal, suggesting that prevalence-induced concept change is a domain general effect that is present in all categorical judgments.

Despite strong evidence for the existence of the effect, the mechanism by which prevalence-induced concept change occurs is unclear and there is no quantitative model of the effect. In their paper, Levari and colleagues propose a qualitative account in which concepts are determined by comparison with recently encountered examples. Thus, when blue stimuli are rare, most of the recent stimuli were purple meaning that an ambiguously blue/purple stimulus appears more blue by comparison with these purple examples. A key component of this account is that this comparison process should be always be active, even when the frequency of each concept is fixed. This suggests that a quantitative account of the comparison process would allow us to predict the extent prevalence-induced concept change *based on behavior when the prevalence is fixed*.

In this paper we develop such a quantitative model of prevalence-induced concept change. Using this model we show that prevalence-induced concept change is driven by two competing effects: (1) a comparison with recent stimuli as suggested by Levari et al, and (2) a tendency to repeat the previous response. Thus, the first effect tends to expand rare categories, as the average stimulus

is more reflective of the most prevalent category. Conversely, the second effect tends to contract rare categories, as repeating a response makes the rare category *less* likely to be selected. Empirically, we find that the first effect is larger than the second causing rare concepts to expand overall. Moreover, by fitting our model to participants' behavior when the prevalence is fixed, we were able to quantitatively predict the size of prevalence-induced concept change when the prevalence shifts, providing strong evidence that these sequential effects drive prevalence-induced concept change.

Results

Our results are built on data from Levari et al. (2018) which can be downloaded from (Levari, 2018). In this dataset, the authors demonstrate prevalence-induced concept change in seven different studies with three different modalities. In Studies 1-5, participants were asked classify colored stimuli as either purple or blue; in Study 6, they were asked to classify faces as threatening or non-threatening; and in Study 7, they were asked to classify research proposals as ethical or unethical. In all studies, behavior was compared between a condition in which the prevalence of each category was fixed over the course of the experiment (e.g. 50% objectively purple, 50% objectively blue), and conditions in which the prevalence of each category was initially fixed but then changed to take on a new value (e.g. 75% purple, 25% blue).

Across all seven of these studies, the key finding is that the category boundary shifts in the changing prevalence condition vs the fixed prevalence condition. Thus, in Studies 1-4, participants are more likely to classify stimuli as blue when the frequency of blue stimuli decreases (Figure 1A); in Study 5, participants are less likely to classify stimuli as blue when the

frequency of blue stimuli increases (Figure 1B); in Study 6, participants are more likely to classify a face as threatening when the frequency of threatening faces decreases (Figure 1C); and in Study 7, participants are more likely to classify a study as unethical when the frequency of unethical studies decreases (Figure 1D). Taken together these results suggest that prevalence-induced concept change is a domain general phenomenon in which rarely encountered concepts tend to expand and frequently encountered concepts tend to contract.

Figure 1 – Main findings from Levari et al. comparing behavior before and after the change in prevalence. (A) Aggregate behavior in Studies 1-4 in which the prevalence of blue stimuli decreases after the first 200 trials. In all of these experiments, participants are more likely to classify ambiguous stimuli as blue as the prevalence of blue stimuli declines. (B) Behavior in Study 5 where the prevalence of blue stimuli *increases* after the first 200 trials causing ambiguous stimuli to be classified as purple in the latter part of the experiment. (C) Ambiguously threatening stimuli are more likely to be classified as threatening following a decrease in the prevalence of threatening stimuli. (D) Ambiguously ethical research proposals are more likely to be classified as unethical as the prevalence of unethical proposals declines.

Stimulus and response on previous trial affect current decision

Levari et al. suggest that the effects in Figure 1 are due to people using a comparison process in order to make their decision. In particular, they propose that people make their judgment (about blueness, threat level and unethicity) by comparing the current stimulus with recently encountered stimuli. If this is indeed how the judgment proceeds then we should see evidence for this effect in the form of sequential effects whereby the objective value of the last stimulus should alter the judgment on the current trial.

To test this prediction we computed the extent to which the current decision was influenced by the identity of the previous stimulus (e.g. objectively purple or blue). In addition, inspired by the literature on choice kernels (e.g. Lau & Glimcher 2005), we also asked whether the current decision was further influenced by the previous response (e.g. subjectively purple or blue). To do this we computed the probability of identifying the target as blue, threatening or unethical in the four different situations defined by the objective and subjective value of the previous stimulus (Figure 2). Because the experiments were not properly balanced for this analysis (e.g. there were more objectively blue stimuli following a blue stimulus in the changing condition) we further corrected this probability according to the objective color of the current stimulus (see Methods). This gave us a measure of the *excess* choice probability of saying the target was blue, threatening or unethical over and above the objective value of the current stimulus.

Figure 2 – Model free analysis showing the existence of sequential effects in all seven experiments. Excess probability (see main text) of classifying a stimulus as blue (top row), threatening (middle row) or unethical (bottom row) as a function of the last stimulus and last response for the Stable (left column) and Changing (right column) conditions. In all cases apart from Study 7 in the changing condition, participants are more likely to choose the opposite option of the last stimulus and the same option as their last response.

These results are shown in Figure 2 where we plot the excess probability for the stable and changing prevalence conditions. Across experiments 1-6, (and numerically in the stable condition in Study 7), we see two main effects of the past trial. First, is a main effect of the objective value of last stimulus (for Studies 1-5 combined, main effect of past stimuli: $F(1, 584) = 106.34$, $p = 5.09 \times 10^{-23}$; for Study 6: $F(1, 131) = 35.02$, $p = 2.69 \times 10^{-8}$; for Study 7: $F(1, 241) =$

0.62, $p = 0.4$) such that participants are biased to choose the opposite identity on the next trial (e.g. they are more likely to choose blue in the trial immediately following an objectively purple stimulus). Second, is a main effect of past response (or the subjective value of the stimulus) (for Studies 1-5: $F(1, 584) = 10.20$, $p = 1.48 \times 10^{-3}$; For Study 6: $F(1, 131) = 51.72$, $p = 4.40 \times 10^{-11}$; For Study 7: $F(1, 241) = 3.16$, $p = 0.07$), whereby participants tend to repeat their last response (e.g. they are more likely to respond blue if they pressed blue last time). The effect of past stimuli is larger in the changing condition (as indicated by an interaction between condition and past stimuli, for Studies 1-5: $F(1, 584) = 48.77$, $p = 7.88 \times 10^{-12}$; for Study 6: $F(1, 131) = 15.08$, $p = 1.62 \times 10^{-4}$; for Study 7: $F(1, 241) = 0.50$, $p = 0.5$), although crucially, they are also present in the stable condition suggesting that these order effects are always active. Moreover, the sequential effects are identical between the two conditions before the prevalence change (Supplementary Figure 1), including in the changing condition in Study 7.

Regression analysis shows sequential effects extend multiple trials into the past

To further quantify the sequential effects beyond the last trial we turned to logistic regression. In this analysis, we assume that the probability of making a particular choice (i.e. the probability of classifying the current stimulus as blue, threatening or unethical) is a function of the current stimulus, f_t , (which ranges from -1 for very purple, non-threatening, or ethical, to +1 for very blue, threatening or unethical), past stimuli, f_{t-i} , for $i = 1$ to N , and past responses, c_{t-i} (-1 for purple, non-threatening or ethical, and +1 for blue, threatening or unethical); i.e.

$$p_t = 1 - \frac{1}{1 + \exp(\beta^{bias} + \beta^f f_t + \sum_{i=1}^N \beta_i^{fPast} f_{t-i} + \sum_{i=1}^N \beta_i^{cPast} c_{t-i})} \quad (\text{Eq. 1})$$

where β^{bias} captures the overall bias for pressing 1 (i.e. classifying the stimulus as blue, threatening or unethical), β^f captures the effect of the current stimulus, β_i^{fPast} captures the effect of the stimulus from i trials ago, and β_i^{cPast} captures the effect of the responses from i trials ago. These regression coefficients were fit separately to each subject using Matlab's glmfit function. We used a value of $N = 5$ for Studies 1-6, and $N = 2$ for Study 7 due to the smaller number of trials in this study. The average, over subjects, of these regression coefficients for each experiment are shown in Figure 3.

The key results in Figure 3 are for the past stimuli β_i^{fPast} and past responses β_i^{cPast} . For past stimuli, for Studies 1-6 all five regression coefficients are negative (see Supplementary Table 1 for complete statistics) indicating that the previous five stimuli have a negative effect on the choice. For Study 7, there is a hint that the last trial has a negative effect on choice in the changing condition ($t(44) = -2.08$, $p = 0.04$), but not the Stable condition ($t(38) = 0.01$, $p = 0.99$). A negative effect of past stimuli is consistent with prevalence-induced context change whereby people become biased towards choosing the less prevalent stimulus. Conversely, for past responses, for Studies 1-6, all five regression coefficients are positive (see Supplementary Table 2 for complete statistics) indicating that participants are more likely to repeat choices up to five trials back. Again the results for Study 7 are weaker, with a hint that the last response has a positive effect in the changing condition ($t(44) = 1.86$, $p = 0.07$), but not in the stable condition ($t(38) = 0.89$, $p = 0.38$). This positive effect of past responses works against prevalence-induced context change as it encourages repetition of more frequent responses.

For both past stimuli and past responses, the magnitude of the effect decreases with trials into the past. In addition, we see that the effect of past stimuli, but not past responses, appears to be

stronger in the changing condition than the stable condition, at least in Studies 1-6 (for past stimuli, in Study 5: main effect of condition: $F(1, 860) = 36.98$, $p = 5.39 \times 10^{-9}$; in Study 6: $F(1, 180) = 5.83$, $p = 0.02$; in Study 7: $F(1, 82) = 2.04$, $p = 0.16$; For past responses the main effect of condition is not significant for any experiment, lowest p value = 0.17 in Studies 1-5). Finally, we note that the results for Study 7 are more ambiguous overall and there is much weaker evidence for an effect of previous stimuli or previous responses. This ambiguous result may reflect the much smaller number of trials in this Study (240 vs 800-1010 in Studies 1-6) and the fact that the effect size of prevalence-induced context change is much smaller in this Study (Figure 1).

Figure 3 – Mean regression weights from the logistic regression model. From left to right: the overall bias, β^{bias} , for blue (top row), threatening (middle row) or unethical (bottom row), the weight given to the current stimulus, β^f , the weight given to past

stimuli, β^{fPast} , and the weight given to past responses, β^{cPast} . Note that the weights given to past stimuli are all negative for Studies 1-6, consistent with prevalence-induced context change, while weights given to past responses are positive, which works against prevalence-induced context change.

Sequential effects when the prevalence is fixed predicts concept change when the prevalence shifts

If the sequential effects shown in Figures 2 and 3 really underlie prevalence-induced context change, then we should be able to use estimates of these effects when the prevalence is fixed to predict the amount of concept change when the prevalence shifts. That is, if we can quantify the sequential effects in each person's behavior *before* the prevalence changes, we should be able to predict the extent of concept change *after* the prevalence changes.

More specifically, we use the following three step approach to predict prevalence change in the data. First, we fit the parameters of a model of the sequential effects on the first 200 (or, in Study 7, 48) trials, before the prevalence changes in the changing condition. Next, we use these parameter values to simulate behavior of the model when the prevalence shifts. Finally, we compare the change in concept boundaries predicted by the simulated behavior to the actual change in concept boundaries observed in human behavior.

In principle, this prediction could be done using the regression model in the previous section. In practice, however, the full regression model has a large number of parameters (12 if we look five trials into the past), which leads to relatively poor fitting when the number of trials is small (which is especially problematic for Study 7). Instead, to mitigate this problem, we use a reduced form of the regression model with 6 free parameters. In this model, we assume that the weight given to past stimuli and past choices falls off exponentially with trials into the past. This allows

us to write the choice probability in terms of the exponentially weighted sum of past stimuli, F_t , and past responses, C_t , as

$$p_t = 1 - \frac{1}{1 + \exp(\beta_0 + \beta_f f_t + \beta_F F_t + \beta_C C_t)} \quad (\text{Eq. 2})$$

where F_t is the exponentially weighted sum of past stimuli, which updates according to

$$F_{t+1} = \lambda_F F_t + f_t \quad (\text{Eq. 3})$$

and C_t is the exponentially weighted sum of past responses, which updates as

$$C_{t+1} = \lambda_C C_t + c_t \quad (\text{Eq. 4})$$

In these equations, the parameters λ_F and λ_C control the rate of decay of the exponential weighting with larger values corresponding to slower decay.

Overall this reduced model has six free parameters ($\beta_0, \beta_f, \beta_F, \beta_C, \lambda_F, \lambda_C$), which can be estimated for each subject from behavior using a standard maximum likelihood approach (Daw 2011). Average values for these parameters are shown in Figure 4. As with the regression analysis we see that past stimuli tend to have a negative effect on the current choice, at least for Studies 1-6, (for Studies 1-5: in the stable condition, $t(78) = -14.78$, $p = 2.52 \times 10^{-24}$, in the changing condition, $t(137) = -18.34$, $p = 1.05 \times 10^{-38}$; for Study 6: stable, $t(26) = -6.95$, $p = 2.24 \times 10^{-7}$, changing, $t(21) = -6.32$, $p = 2.87 \times 10^{-6}$; Study 7: stable, $t(38) = 0.23$, $p = 0.82$, changing, $t(44) = -0.23$, $p = 0.82$), consistent with a role in prevalence-induced concept change, while past choices have a positive effect (for Studies 1-5: in the stable condition: $t(78) = 10.98$, $p = 1.72 \times 10^{-17}$, in the changing condition: $t(137) = 15.11$, $p = 5.56 \times 10^{-31}$; for Study 6: stable, : $t(26) = 9.53$, $p = 5.71 \times 10^{-10}$, changing, $t(21) = 7.66$, $p = 1.65 \times 10^{-7}$; Study 7: stable, $t(38) = 0.18$,

$p = 0.86$, changing $t(44) = 0.26$, $p = 0.80$), consistent with a role that opposes prevalence-induced concept change. Again there are small differences between the parameter values for the stable and changing conditions, but they do not reach significance and, as shown below, these differences have only a minor effect on the ability of the model to predict how the concepts shift when the prevalence changes.

Figure 4 – Mean parameter values for the reduced model in each of the seven experiments. From left to right we plot: the overall bias for blue, threatening or unethical, β_0 , the weight given to the current stimulus, β_f , the weight given to past stimuli, β_F , the decay parameter for past stimuli, λ_F , the weight given to past responses, β_C , and the decay parameter for past responses, β_C . Again we see that past stimuli and past responses have opposite effects on the choice, with past stimuli biasing people to choose the opposite of the previous stimuli, and past responses biasing people to repeat their previous responses.

Next, we used the reduced model to predict the extent of prevalence-induced concept change by fitting behavior when the prevalence was fixed. In particular, we estimated the 6 parameters values for each subject from the first 200 (or 48 in Study 7) trials of the changing condition, i.e. before the prevalence had changed. We then simulated behavior according to these fit parameter values for the entire experiment using the same set of stimuli seen by the subject. Simulations were repeated 100 times to ensure minimal noise in the final result.

Results of this prediction process are shown in Figure 5. As is clear from this figure, there is fairly close correspondence between the simulated behavior (solid yellow line) and the observed behavior (yellow dots) on the later trials. While the agreement is not always perfect (e.g. in Studies 1 and 4 the slope of the predicted curve is too shallow) the correspondence is quite remarkable given that the parameter values were estimated from data where the prevalence was fixed. Of course there is also excellent agreement on the early trials (blue line/dots), but this is to be expected given that the parameters were fit on this data. The worst predictions are for Study 7. This is likely due to the relatively small number of trials in this study (only 48 trials before the prevalence change vs 200 for the other experiments), which was necessary given the complexity of the stimuli in this case, and the overall smaller effect seen in this experiment (Figure 1). Even here, the model gets the direction of effect correct, predicting more unethical judgments as the frequency of unethical proposals is decreased.

Figure 5 – Predicted vs actual prevalence-induced context change in Studies 1-7. Comparison between the actual (dots) and predicted (lines) probability of classifying a stimulus as blue (A-E), threatening (F) or unethical (G) as a function of objective stimulus value, f . For all studies except Study 7, there is close correspondence between the predicted and actual effects.

Discussion

In this paper we reanalyzed the data of Levari and colleagues (2018) to show that prevalence-induced concept change can be accounted for, both qualitatively and quantitatively, by sequential effects. In particular, we found two competing sequential effects that together determine the

magnitude of prevalence-induced concept change. The first was a negative effect of past stimuli, such that people were more likely to choose the opposite of the stimuli they had recently seen (e.g. more likely to choose blue if recent stimuli were purple). The second was a positive effect of past responses, such that people were more likely to repeat a responses they had previously made. Crucially, we found that these order effects were at play even when the prevalence was stable. Thus by fitting a model to participants' behavior before the prevalence changed, we were able to predict the extent of prevalence-induced concept change when the prevalence shifted. This provides strong evidence that the two sequential effects are the cognitive mechanisms behind the observed phenomenon of prevalence-induced concept change.

Of the two sequential effects we found, the negative effect of past stimuli is the one that causes infrequently encountered concepts to expand. This negative effect is consistent with the qualitative proposal of Levari et al. who suggested that the concept judgment is made by comparison to recently encountered stimuli. In our reduced model, this comparison is with an exponentially weighted sum of previous stimuli, such that when one concept is more prevalent, ambiguous stimuli are more likely to be assigned to the less prevalent concept. It is tempting to think of the exponential weighting of previous stimuli as part of a learning process whereby the statistics of the stimuli (e.g. how blue, threatening, or unethical) are implicitly learned (Reber 1967). Such implicit learning has been seen in a number of different situations (Cleeremans et al. 1998; Dienes & Berry 1997; Cox et al. 2012) and it would be interesting to see whether the it is the same process, perhaps even subserved by the same neural systems, here.

In contrast to the negative effect of past stimuli, the positive effect of previous responses works against prevalence-induced concept change. This repetition bias was not included in Levari et al's account, but is consistent with previous work on choice kernels and status quo bias, whereby

a tendency to repeat choices has been observed in a variety of situations in humans (e.g. Samuelson & Zeckhauser 1988; Nicolle et al. 2011) and in other animals (Lau & Glimcher 2005; Padoa-Schioppa, 2013).

The fact that the two sequential effects have opposite effects on prevalence-induced concept change is intriguing and it is tempting to wonder whether these opponent processes can have a stabilizing effect on concepts, at least in the short term. However, it may also be the case that the two processes are unrelated and an artifact of the relatively fast timing of the tasks (where the average response time ranged from around 200-1000 ms for Studies 1-6 to nearly 5 seconds for Study 7). More generally, future experiments will need to investigate whether prevalence-induced concept change as measured in experiments such as these is sensitive to time between trials. This will be particularly important if we are to link the effects observed in the lab to real-world instances of concept change, where the time between “stimuli” may be hours, days or even years. In this regard, it is notable that the smallest effect is observed in Study 7, which has by far the longest reaction times.

Another important avenue for future work, will be to determine whether there is a maximum extent to which changes in prevalence can alter the definition of a concept. Our reduced model makes a clear prediction in this regard: so long as $|\lambda_F| < 1$ and $|\lambda_C| < 1$, then the extent of prevalence induced concept change is bounded. Intuitively such a maximum effect size seems plausible for the perceptual experiments, Studies 1-5, where it seems unlikely that prevalence change alone could cause very purple stimuli to be classified as blue. However, in more complex cases, where the concepts themselves are less well defined to begin with, it may be that over the long term, concepts could change considerably. Clearly more work is required to define the concept of prevalence-induced concept change itself!

Methods

Data and Code availability

All data comes from Levari et al. (2018) and is available at (Levari 2018). We made minor changes to the formatting of the spreadsheets to facilitate reading into Matlab. Specifically we replaced all “NA”s (for not a number) with the Matlab equivalent, “nan,” and all commas (i.e. the “,” symbol) with spaces. This was done using the “Find and Replace” function in Excel.

All analysis code is our own and will be made available on GitHub upon publication. For review purposes, the code and edited spreadsheets are available as Supplementary Material.

The Levari et al. experiments

Data from Levari et al. (2018) was downloaded from (Levari, 2018). This dataset explores prevalence-induced concept change across three different contexts in seven different experiments. These studies are outlined briefly below.

Studies 1-5: Discriminating between blue and purple dots

In Studies 1 through 5, participants were required to decide whether a series of colored dots were either purple or blue. Stimuli were presented one at a time and participants indicated their choice (“purple” or “blue”) via a button press. Stimuli varied in objective color from “Very Purple” (RGB = [100, 0, 155]) to “Very Blue” (RGB = [1, 0, 254]). In the “Stable” prevalence condition, the proportion of blue and purple dots was kept constant over time. While in the “Changing” prevalence conditions the prevalence of blue dots either decreased (Studies 1-4) or increased (Study 5) over time. Apart from the directionality of changing prevalence, the five studies

differed in the number of participants (22 for Study 1, 43 for Study 2, 92 for Study 3, 37 for Study 4 and 23 for Study 5), the instructions, incentives, the rate at which the prevalence changed (gradually for Studies 1, 2, 3, 5 and abruptly for in Study 4), and the total number of trials (1000 for Studies 1, 2, and 5, 800 for Studies 3 and 4). Apart from Study 5, in which the direction of effect is in the opposite direction, there are no major differences in the results of these experiments and in most of our analyses we consider data from all 5 studies together.

Study 6: Discriminating between threatening and non-threatening faces

In Study 6, 49 participants were required to decide whether a series of faces were either threatening or non-threatening. The faces were previously normed such that each face came with an objective threat measure. As with the blue-purple experiments, after 200 trials the prevalence of threatening faces decreased.

Study 7: Discriminating between ethical and unethical research proposals

In Study 7, 84 participants played the role of ethics board members and were required to decide whether a series of written research proposals were ethical or unethical. As with previous experiments, in the Changing condition the prevalence of unethical experiments decreased after 48 trials.

Range normalization of objective stimulus value

To facilitate comparison between the different studies we transformed the objective measure of each stimulus (color, threat level and ethicality) into a variable, f , that ranged from -1 (for very purple, very safe, or very ethical) to +1 (for very blue, very threatening, or very unethical). For Studies 1-5

$$f = 1 - 2 \times \frac{RED - \min (RED)}{\max(RED) - \min (RED)}$$

where RED denotes the red component of the RGB color value of the stimulus and is the name of a variable in the Levvari et al. dataset.

For threat level,

$$f = 2 \times \frac{FACE - \min (FACE)}{\max(FACE) - \min (FACE)} - 1$$

where the FACE variable is large for threatening faces.

For unethicity,

$$f = 2 \times \frac{NORM_MEAN - \min (NORM_MEAN)}{\max(NORM_MEAN) - \min (NORM_MEAN)} - 1$$

where the NORM_MEAN is large for unethical proposals.

Computing the excess choice probabilities

The excess choice probabilities were computed in the following way. First, we computed the raw probability of choosing blue, threatening or unethical as a function of the past stimulus ($f_{t-1} < 0$ or $f_{t-1} > 0$) and past choice, ($c_{t-1} = -1$ or $c_{t-1} = +1$). Next we corrected this measure to take account of the fact that the experiment was not properly balanced for this analysis and that the average stimulus value was different for each of the four different mixtures of past stimulus and past response. To make this correction, we compute the probability that the *objective* stimulus was blue, threatening or unethical, on trial t as a function of past stimulus and past response on trial $t-1$. That is we computed the probability that $f_t > 0$ for each of the four

conditions. Finally, we computed the excess choice probability as the raw probability minus the probability of the objective stimulus.

Fitting the reduced model

Parameters of the reduced model were fit separately to each participant using a standard maximum likelihood approach (Daw 2011). In particular, we computed the probability of each choice as a function of parameters $(\beta_0, \beta_f, \beta_F, \beta_C, \lambda_F, \lambda_C)$ as

$$CP_t(\beta_0, \beta_f, \beta_F, \beta_C, \lambda_F, \lambda_C) = \begin{cases} p_t & \text{if } c_t = +1 \\ 1 - p_t & \text{if } c_t = -1 \end{cases}$$

where p_t is given by equation 2. We then computed the log likelihood as

$$LL(\beta_0, \beta_f, \beta_F, \beta_C, \lambda_F, \lambda_C) = \sum_{t=1}^T \log CP_t(\beta_0, \beta_f, \beta_F, \beta_C, \lambda_F, \lambda_C)$$

where T is either the total number of trials in the experiment for Figure 4, or the number of trials before the prevalence change ($T = 200$ for Studies 1-6 and $T = 48$ for Study 7) for Figure 5.

Maximum likelihood parameter values were then estimated by maximizing the log-likelihood (in practice minimizing the negative log-likelihood) using Matlab's `fmincon` function.

Prediction using the reduced model

Predictions in Figure 5 were made by simulating the reduced model. In particular, we only fit behavior to the trials before the prevalence shift (i.e. 200 trials in Studies 1-6 and 48 trials in Study 7). We then simulated behavior for all trials (including the initial 200 or 48 that were used for fitting) using the same series of stimulus values seen by that person. On each trial of the

simulation, we computed the choice probability according to equation 2 and then sample a particular choice with this probability. F_t and C_t were initialized at zero and C_t was updated according to the model's, rather than the participants, choices. Simulations were repeated 100 times to compute the average behavior of the model.

References

- Cleeremans, A., Destrebecqz, A., & Boyer, M. (1998). Implicit learning: News from the front. *Trends in cognitive sciences*, 2(10), 406-416.
- Cox, W. T., Abramson, L. Y., Devine, P. G., & Hollon, S. D. (2012). Stereotypes, prejudice, and depression: The integrated perspective. *Perspectives on Psychological Science*, 7(5), 427-449.
- Daw, N. D. (2011). Trial-by-trial data analysis using computational models. *Decision making, affect, and learning: Attention and performance XXIII*, 23, 3-38.
- Dienes, Z., & Berry, D. (1997). Implicit learning: Below the subjective threshold. *Psychonomic bulletin & review*, 4(1), 3-23.
- Lau, B., & Glimcher, P. W. (2005). Dynamic response-by-response models of matching behavior in rhesus monkeys. *Journal of the experimental analysis of behavior*, 84(3), 555-579.
- Levari, D. E., Gilbert, D. T., Wilson, T. D., Sievers, B., Amodio, D. M., & Wheatley, T. (2018). Prevalence-induced concept change in human judgment. *Science*, 360(6396), 1465-1467.
- Levari, D. E., Prevalence-induced concept change in human judgment. *Zenodo*, Version v1.0.0 (2018); 10.5281/zenodo.1219833.

- Nicolle, A., Fleming, S. M., Bach, D. R., Driver, J., & Dolan, R. J. (2011). A regret-induced status quo bias. *Journal of Neuroscience*, 31(9), 3320-3327.
- Padoa-Schioppa, C. (2013). Neuronal origins of choice variability in economic decisions. *Neuron*, 80(5), 1322-1336.
- Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of verbal learning and verbal behavior*, 6(6), 855-863.
- Samuelson, W., & Zeckhauser, R. (1988). Status quo bias in decision making. *Journal of risk and uncertainty*, 1(1), 7-59.