

A Bayesian Model of Social Influence under Risk and Uncertainty.

Simon Ciranka (ciranka@mpib-berlin.mpg.de)

Center for Adaptive Rationality, Max Planck Institute for Human Development, Lenzeallee 94, 14195 Berlin, Germany,

Max Planck UCL Centre for Computational Psychiatry and Ageing Research, Lenzeallee 94, 14195 Berlin, Germany,

Wouter van den Bos (vandenbos@mpib-berlin.mpg.de)

Center for Adaptive Rationality, Max Planck Institute for Human Development, Lenzeallee 94, 14195 Berlin, Germany,

Max Planck UCL Centre for Computational Psychiatry and Ageing Research, Lenzeallee 94, 14195 Berlin, Germany,

Department of Developmental Psychology, University of Amsterdam, 1012 WX Amsterdam, Netherlands

Abstract

Humans live in an uncertain world and often rely on social information in order to reduce uncertainty. However, the relationship between uncertainty and social information use is not yet fully understood. In this work we argue that previous studies have often neglected different degrees of uncertainty that need to be accounted for when studying social information use. We introduce a novel experimental paradigm to measure risky decision making, wherein social information and uncertainty are manipulated. We also developed a Bayesian model of social information use. We show that across different levels of uncertainty; social influence follows similar principles. Social information is more impactful when individuals are more uncertain. Notably, this relationship holds for experimental manipulations of uncertainty but also for subjective uncertainty within experimental conditions. We conclude with discussing that social influence can be better understood when paying credit to subjective uncertainties and preferences.

Keywords: Social Influence; Hierarchical Bayes; Uncertainty

Introduction

Some people prefer to take risks while others rather play it safe. Such individual differences in risk-preference influence human behaviour in a variety of different contexts, ranging from problematic gambling over binge drinking up to risky driving (Frey, Pedroni, Mata, Rieskamp, & Hertwig, 2017). Individual differences in risk-preferences are traditionally attributed to differences in weighting potential outcomes or probabilities. However, in the real-world outcomes and probabilities are often not known, and the uncertainty about whether an action will lead to the desired outcome plays an important role in deciding what to do.

Uncertainty arises on different levels of cognitive processing as pointed out by Bach & Dolan, (2012). To illustrate these levels, picture a cyclist who wants to be at work on time on a wet and foggy day. She wants to ride as fast as she can but in order to decide whether she should pick up more speed, some uncertainties need to be resolved. First, there is *sensory uncertainty*, which stems from the fact that our senses never provide us with a perfect representation of our surroundings, but with a noisy estimate of it. Our exemplary cyclist cannot judge precisely whether there will be a sharp turn straight ahead or not. Due to the fog, she will only have a vague notion of the streets' path. This sensory uncertainty converts to *state uncertainty*- a lack of knowledge about the current

state of the world. Turning to our cyclist this means that she likely is uncertain about how slippery the street is; or if she assumes that there is a sharp turn coming up, how far it exactly is until she encounters it. Then there is *rule uncertainty*, which means that decision-makers often find themselves in the situation of not knowing the precise probability with which an action will result in a desired outcome. For example, when the cyclist is uncertain about how slippery the street or how far away the turn is, she will also be uncertain about with what probability speeding up will result in an accident at the turn. Finally, there is *outcome uncertainty*, usually referred to as *risk*. Risk describes the fact that the result of one decision cannot be predicted with certainty, even if all probabilities are known to the decision-maker (i.e. there is no uncertainty on all previous levels). For instance, our cyclist might conclude that there is exactly a 30% chance of crashing at the next sharp turn if she puts up more speed, she still cannot know the outcome for sure. Finally, she has to decide. Her individual preferences for uncertainty (including risk) will guide her choice and it is likely that she is uncertain about these preferences as well (Moutoussis, Dolan, & Dayan, 2016). While we acknowledge that these uncertainties may be qualitatively different from one another, it is important to point out that uncertainties are organised hierarchically. This means that sensory uncertainty can have a downstream effect on the subjective experience of uncertainty about the potential outcomes associated with an action, however these uncertainties can be attenuated by learning or attentional processes. Risk is qualitatively different because it cannot be further reduced by additional information. Therefore, throughout this paper, we will subsume the sensory, state and rule uncertainty under *uncertainty* hereby adopt the widely accepted distinction between *uncertainty* and *risk* (Hertwig, Barron, Weber, & Erev, 2004).

One way by which humans increase their confidence in a choice under uncertainty is by using social information (FeldmanHall & Shenhav, 2019). Imagine our cyclist watching other cyclists failing to take the turn. This will have an effect on her individual decision to accelerate and this effect is likely to be stronger when she is more uncertain about what to do. Indeed, recent research from perceptual judgements has shown that people use more social information when they are presented with more uncertain, or noisy, stimuli. Moreover, there is evidence that the integration of individual and social information follows

Bayesian principles (Bahrami et al., 2010; Toelch & Dolan, 2015). However, it is not known how different degrees of uncertainty affect preference-based choice and social influence therein. The interplay of preferences, uncertainty and social influence has traditionally been studied by using decisions between monetary lotteries, which are solely subject to *risk*. While this research has shown that people *do* integrate social information when making a risky choice (Chung, Christopoulos, King-Casas, Ball, & Chiu, 2015; Ciranka & van den Bos, 2019), it is still unknown to what extent social influence in preference based choice is modulated by experienced uncertainties. Past research in that respect provides heterogenous conclusions which sometimes even are counterintuitive. While some researches find that more uncertainty leads to more social influence (FeldmanHall & Shenhav, 2019; Toyokawa, Whalen, & Laland, 2019), others report the opposite (Delfino, Marengo, & Ploner, 2016). Therefore, it is still unclear to what extent the same principles reported for sensory uncertainty also apply to different kinds of uncertainty and when individual preferences are concerned. In order to provide a more complete picture we aim investigate the mechanisms of social influence by explicitly formulating a Bayesian model of social influence within which subjective uncertainties can be quantified.

Goals and Scope. Our goal in this manuscript is twofold: First, we present a novel experimental paradigm that allows us to study social influence under different uncertainties. We therefore extend psychophysiological work into the value-based domain, focussing on the distinction between *uncertainty* and *risk*. Second, we aim to provide a formal model of how humans integrate social and non-social information in preference-based choice under different uncertainties and show how a Bayesian approach provides novel insights into the mechanisms behind social influence under uncertainty.

Methods

Subjects. We report the data of $n=61$ subjects (aged 18-26, $m=20.82$) who completed the task in our laboratory. The data collected was approved of our institutes' local ethics committee as part of a bigger developmental study, where subjects performed a battery of social decision-making tasks which will be reported elsewhere.

Task and Procedure. In order to experimentally manipulate uncertainty and social influence we developed the marble task. The marble task was written in javascript, using jspsych (de Leeuw, 2015) and presented in a regular browser.

Subjects were instructed to choose between two lotteries in order to accumulate bonus points. The bonus points were transformed into a monetary bonus (factor 0.0025) which was added to the compensation of 10 € at the end of the experiment. The task consisted of 144 trials in each of which, subjects were asked to decide between a risky and a safe option. When they decided to choose safe, they gained five

points. When subjects chose the risky option, they could win either 8, 20 or 50 points. However, the risky points were only probabilistically awarded. The probabilities were 0.125, 0.25, 0.375, 0.5, 0.625 and 0.75. All possible values were combined with all probabilities equally often. This basic task variant was repeated in four different conditions: risk and uncertainty, which were nested within a solo and a social condition.

In the risk condition (figure 1a), probability information was presented by showing an image of a jar, containing 100 marbles. The proportion of blue and red coloured marbles in this jar represented the underlying probabilities, where the number of blue marbles shown indicated the probability of winning. In the uncertainty condition (figure 1b), participants had to learn about the underlying probabilities before deciding. In order to do so first, subjects saw a sequence of 9 pseudo-random draws from the jar instead of being informed about the exact proportions directly. Thus, subjects had to integrate new pieces of information about the outcome probabilities with every draw from the jar. Before making a choice, we asked participants to indicate their estimate of the underlying probability, using a slider. To make sure that all participants were presented with the same information, we sampled binomial sequences before the experiment until the sequences' mean was as representative of the underlying probability as possible. Hence subjects did not see real random draws from the jar.

Social Information. After participants completed a block without social information (solo condition), we computed the percentage of trials at which the current subject chose the risky option. In the next block, a social condition, we assigned an advisor to the participant by finding a subject in our database of participants who completed a similar experiment that used the same probability and combinations. Our criterion for matching participants to advisors was that the advisor chose the risky option on average 20% more frequently than the participant. We chose this threshold for two reasons: First, studies have shown that social information that is too close or too far from individual preferences has little impact (Moussaïd, Herzog, Kämmer, & Hertwig, 2017). And second, this was a way to keep the relationship between social information and individual propensities constant across participants. Social information was presented to the current subject by framing the option that the advisor previously chose on a trial with the same value and probability.

We assumed that people would be more inclined to choose the risky option when it has a higher expected value compared to the safe option. We expected social influence to be greater when uncertainty is higher. That means we expected greater social influence under uncertainty than under risk.

A Bayesian Model of Social Influence. Our main focus was to quantify the mechanisms behind social influence under uncertainty and risk. To this end we developed a




Figure 1: The marble task. A) subjects either decide between a risky or a safe option. Here, outcome probabilities are described to the subjects, thus they make decisions under *risk*. These decisions are either made with (bottom) or without (top) social information. B) Subjects make decisions under *uncertainty*. They see a sequence of draws from the outcome distribution and are asked to indicate their estimate of the underlying probabilities on a slider. Then they decide whether they want to gamble with the distribution they had just experienced or gain a small amount of bonus points with certainty. Subjects either are presented with social information (bottom) or not (top).

computational model which formalizes the assumption that social influence depends on individual uncertainty (Toelch & Dolan, 2015). The model has three building blocks. First, it rests on expected utility theory, in which objective values are non-linearly transformed into subjective utilities which are then multiplied by their probability of occurrence, resulting in the so-called expected utility. Second, we assume that individuals are uncertain about the exact utilities associated with one option, a common assumption in stochastic utility theory (Blavatsky, 2007). In doing so we can elegantly integrate different sources of uncertainty into a probability distribution over expected utility. To model uncertainty, we treat the underlying probability of winning as Beta-distributed and scale the distribution with the subjective utility of a choice option:

$$EU_{risk} \sim \text{Beta}(\alpha_{ij}, \beta_{ij}) * V_k^{\lambda_i},$$

where V is the value associated with one option. The parameter λ shapes the utility function of subject i . λ accounts for individually varying risk-preferences. Thus, a subjective preference for risk is modelled as the extent of diminishing marginal utility of different prospects. Under diminishing marginal utility, risk aversion occurs because potential reward magnitudes are subjectively compressed therefore more similar to each other and in turn it will be less attractive to take a risk in order to obtain a higher reward.

In the risk condition, shape parameters (α_{ij}) and (β_{ij}) of the Beta distribution are estimated per unique probability condition, j and nested within subjects, i . In the uncertainty condition, we assume that with every new piece of information, subjects update their beliefs about the underlying probabilities in a way that can be described as weighted binomial updating. In the uncertainty condition, we

do possess direct information about the optimal Beta distributions' shape parameters, they are represented by the number of blue/red marbles (i.e. observed pieces of evidence) in a given trial, k , for gains (α_{rule}) and losses (β_{rule}):

$$EU_{uncertain} \sim \text{Beta}(\alpha_{rule_{ij}} * \zeta_i, \beta_{rule_{ij}} * \zeta_i) * V_k^{\lambda_i}.$$

The parameter ζ allows us to express "weighted updating" where estimates close to one indicate Bayesian optimality. Values lower than one mean that subjects underweight new information; higher values speak for overweighting of information.

The third part of our model relates to social influence. As we also represent our subjects' uncertainty about the prospects of their decisions, it is now straightforward to model informational social influence as Bayesian inference. The model formalizes the following scenario. If someone observes risky social information, this will make them believe that the risky option is more advantageous than previously thought based on individual information alone.




Figure 2: Schematic of Bayesian updating, risky advice is observed. Left: When there is high prior (black) uncertainty, observing social information favoring a risky option will strongly affect the subjects' posterior belief (green). Right: when there is less prior uncertainty, the same information will lead to a less pronounced shift of the posterior.

Observing social information will also and reduce their uncertainty about how to decide. In other words: observing social information shifts the mean and reduces the variance of the beliefs which have been formed on the basis of individual information alone.

In our model, the prior is the individual's belief about the outcome probabilities and the contribution of social information, x , to our subjects' beliefs comprises the likelihood. Multiplying the beta prior, social likelihood and the options' utility yields the posterior expected utility of one option after integrating social and individual information. Social information can be either risky ($x = 1$) or safe ($x = 0$), scaled by a free parameter, η :

$$EU \sim \underbrace{p_k^{x\eta_{ij}} (1-p_k)^{(1-x)\eta_{ij}}}_{\text{Likelihood (social information)}} * \underbrace{p_k^{\alpha_{ik}-1} (1-p_k)^{\beta_{ik}-1}}_{\text{Prior (individual information)}} * \underbrace{V_k^{\lambda_i}}_{\text{Utility}}$$

Where p is Beta distributed according to equations 1 and 2, depending on whether the current condition demands decisions under risk (Eq. 1) or uncertainty (Eq. 2).

η determines the weight of social information: the higher η , the stronger the influence of social information on the subjects' decisions.

In a last step, the difference between safe and risky utilities is fed into a sigmoid function with temperature τ in order to obtain a trial wise choice probability, π :

$$\pi_{ijk} = \frac{1}{1 + e^{-(EU_{riskijk} - EU_{safeijk}) * \tau_i}}$$

Finally, each trial our subjects' choices, X_{ijk} , are predicted as Bernoulli distributed random variables (figure 3).

While the model itself constitutes a Bayesian model of cognition, we also formulated the model in a hierarchical Bayesian way, which allows us to estimate different hyperpriors for different conditions or groups of subjects in the task (figure 3).




Figure 3: Graphical Model for the marble task. Hierarchy of parameters depend on subjects and condition. Subject level parameters are estimated from a unit normal distribution and are then transformed to their actual value using non-centered parametrization in order to achieve efficient subject level parameter sampling. $\Phi(\cdot)$ denotes the unit normal cumulative distribution function

Model fitting was done using stan (Carpenter et al., 2017), with 4 chains and 20000 iterations per chain, 5000 of which we discarded as warmup. Convergence of the Markov Chains was inspected visually and by consulting $\hat{R}$ (all $\hat{R}=1$).

For inference on the relationship between uncertainty and social influence we estimate Bayesian robust correlations, assuming that measurements follow a student-t distribution. We examine the correlation between our subjects' mean uncertainty about utilities, and the degree of social influence they show, as quantified by parameter η . To compute this uncertainty, we dropped the factor which scales the Beta distribution to the level of utilities from all equations. This was done in order to obtain a comparable measure of uncertainty for all trials. We denote the variance of the Beta distribution as uncertainty.

In a "model free" analysis, we inspected whether the experimental manipulation had an effect on our subjects' decisions to take a risk. For this we constructed a Bayesian generalized linear model; predicting each decision with a logit link function, by coding a risky choice as "1" and a safe choice as "0". We used the uncertainty conditions well as the quality of social information as categorical predictors and the expected value as continuous predictor. We report the odds ratio of each predictor. All statistical inference was performed using the brms package for R (Bürkner, 2017), using brms default priors.

Results

In the present study we examined social influence under different types of uncertainty using a novel experimental paradigm and newly developed hierarchical Bayesian models. We were interested to what extent social influence in preference-based choice is modulated by different kinds of uncertainty (fig 4). In general, social influence followed our expectations. When social information was favouring risky decisions, subjects opted more for the risky option ($m = 51.3\%$, $sem = 0.1\%$) than when making decisions alone ($m = 41.9\%$, $sem = 0.1\%$). When social information favoured safe decisions, subjects chose the risky option less often ($m = 36.8\%$, $sem = 1.2\%$). The odds that have been estimated in order to predict each choice imply as well that more risk taking occurred when social information was risky (odds_socialrisk = 1.58, CI = [1.44, 1.78]). When social information favoured safe decisions, the model did credibly predict more safe choices, (odds_socialsafe = 0.89, CI = [0.76, 1.03]), note that this confidence interval does include a 0, when the raw regression coefficient and not the odds ratio is concerned (odds_socialsafe = 0.89, CI = [0.76, 1.03]). Under uncertainty, subjects were less inclined to take a risk (b_socialsafe = -0.12, CI = [-0.27, 0.03]).

Changes in expected value also affected our subject's decision with higher expected values, of the risky option being predictive of more risk-taking (odds_EV = 1.20, CI = [1.19, 1.21]).

Model based Analysis. Here we report the parameter estimates obtained from fitting the outlined model to our




Figure 4: Social information use in the marble task. a) Mean of risky decisions per condition and the bootstrapped 95% confidence interval of the mean (ci) as error bars, overlaid over subject level means (black) accompanied with simulations made under the multivariate posterior of our model (green). b) Relationship of true probability (x axis; red crosses) and the participants probability estimates (black) and the mean of this estimate (opaque dots) in the uncertainty condition. c) Social information had a higher impact when there was more uncertainty imposed upon our subjects by the experimental conditions. d) This difference was in the range of a medium effect size. e) The more uncertainty one subject experiences within the risk condition, the stronger the same subject weights social information (y axis). f) the same as e) but for the uncertainty condition.

subjects' decisions. Our results show that people use social information more when they are more uncertain of how to decide. This is implied by the higher social influence parameter estimates, when there is uncertainty, as opposed to risk alone (fig 4c). This difference is in the range of a medium effect size, (posterior mean estimate = 0.62; fig 4d). Moreover, the relationship between uncertainty and social influence even holds for subjective uncertainties *within* each condition. This is implied by the correlation between uncertainty about the risky options' utility and mean social influence in the risk and in the uncertainty condition. The Bayesian robust correlation between social influence and uncertainty was substantial in the risk ($\rho = 0.45$ CI = [0.2, 0.65]) and less pronounced and not credibly present in the uncertainty condition ($\rho = 0.26$; CI = [-0.01, 0.49]). Taken together we show behaviourally and computationally that subjects who were subjectively more uncertain, used social information more, despite the fact that every participant observed the same amount of information. Subjects' showed considerable variation in how they weighted each piece of individual information under uncertainty. This weighting is expressed by the parameter ζ , which was on average estimated to be smaller than one ($m = 0.71$, $sd = 0.74$). This is evidence that subjects are underweighting individual information, when comparing their behaviour to how an ideal observer would perform. Underweighting leads to higher uncertainty and this

uncertainty can explain why subjects overestimate low and underestimate high probabilities ($\rho_{\text{true;estimate}} = 0.38$, CI = [0.36, 0.40]). Subjective uncertainty results in the observed regression to the mean of our participants probability estimates (figure 4b). Notably, subjects put a stronger weight on social information as compared to individual information. This is implied by the posterior mean estimates of η , which constitute the weight of social information on our subjects' decisions. η was almost exclusively greater than 1 ($m = 2.40$, $sd = 1.51$), irrespective of whether decisions were made under risk or under uncertainty. This means that one piece of advice had a higher influence on our subject's choice than one piece of individual evidence.

In sum, we show that uncertainty leads to less risky decisions and more social influence. We provide evidence for the influence of subjective uncertainty on decision making.

Discussion

People use social information when they are uncertain of how to decide, therefore it has been argued that social information use follows Bayesian principles (Toelch & Dolan, 2015). In this work we examine to what extent these principles generalize to different kinds of uncertainty and decisions which are subject to preferences. To do so, we developed an experimental paradigm in which subjects are asked to make decisions between two options, where we vary uncertainty in

two conditions. In a *risk* condition we directly inform participants about the outcome distribution. In a *uncertainty* condition, subjects learn about the distribution by observing draws from the distribution prior to their choice. Both conditions were administered with and without social information. We show that when there was more uncertainty, our subjects were more risk averse. This extends previous reports of ambiguity aversion in similar economic games (van den Bos & Hertwig, 2017) and underscores the importance of considering uncertainty as modulator of behaviour. Strikingly, our modelling results show that there is considerable variability in our subjects' tendency to represent uncertainty, which implies a distinction between *objective uncertainties*, which are the consequence of our experimental manipulation and *subjective representations* of these uncertainties (Sharpe, 2018), which are important when subjects make a decision. The subjective nature of uncertainty contributes to the complex dynamics underlying risk taking in real-life because as we have pointed out, humans almost never enjoy the convenience of making decisions in the face of pure risk – decision making is rife with uncertainties. Risk however, is the domain in which decision-making is still usually studied. One consequence of this is that laboratory tasks hardly generalize to real-life. It seems as if tasks featuring multiple sources of uncertainty are stronger related to real life risk-taking (Frey et al., 2017). Uncertainty also plays a key role in social influence. In line with this, we show that social influence is higher when there is more uncertainty in our experimental conditions. To investigate this further, we proposed a model of social influence as Bayesian updating, where uncertainty becomes a mechanism of belief and behaviour change. Crucially, our model captures subjective representations of uncertainty. We show that subjective uncertainty is correlated with the extent of social influence, an inference that could not be made by merely observing aggregate behaviour. We argue that in order to understand social influence, or even use social information to nudge people into one or the other decision policy, more attention needs to be paid on identifying the matters about which an individual is uncertain, instead of attempting to identify when, whether or how individual's decision-making is compromised by social information. For example, others have identified uncertainty about preferences as one factor driving social influence (Moutoussis et al., 2016). Our results are consistent with this idea and targeted experimental paradigms will be insightful to further differentiate these nuances underlying social influence. Note that social information use may stem from different motivations, which we did not consider here. For instance, a prominent distinction between different kinds of social influence is the one between informational and normative influence (Deutsch & Gerard, 1955). Whereas informational social influence reduces individual uncertainty about the world, normative social influence can be used to establish social bonds and belonging. From this point of view, behaviour does not maximise an economic utility but the utility is attributed to social goals per se because any

behaviour also sends a social signal to the observer. This different motivation brings conformity into the equation where using social information can signal agreement to others and is not necessarily related to subjective uncertainties about the objective world. These kinds of behaviours require inference upon norms and are outside of the scope of the model which we have proposed in this paper, but will be interesting to investigate in the future. Another promising application of our task and model is developmental research. For instance, it is known that adolescents are especially sensitive to social information when taking risks (Ciranka & van den Bos, 2019), and we believe that understanding of adolescents propensity to take risks will benefit from the specificity of the approach which we have outlined here.

Summary and Conclusion

People use social information to be less uncertain how to decide. However, it was previously unknown to what extent different types of uncertainty relate to social influence. We therefore developed a novel paradigm and computation model. We show that peoples use of social information can be described as Bayesian updating that follows similar principles under different kinds of uncertainty. We believe it is crucial that future research aims at identifying the uncertainties which drive social influence in human decision making across different domains or periods of life.

Acknowledgments

We are grateful to Chantal Wysocki and Jann Wäscher for collecting the data despite all logistic difficulties. SC is a pre-doctoral fellow of the International Max Planck Research School on Computational Methods in Psychiatry and Ageing Research (IMPRS COMP2PSYCH). The participating institutions are the Max Planck Institute for Human Development, Berlin, Germany, and University College London, London, UK. WB is supported by Open Research Area (ID 176), the Jacobs Foundation, the European Research Council (ERC-2018-StG-803338) and the Netherlands Organization for Scientific Research (NWO-VIDI016.Vidi.185.068).

References

- Bach, D. R., & Dolan, R. J. (2012). Knowing how much you don't know: a neural organization of uncertainty estimates. *Nature Reviews Neuroscience*, 13(8), 572–586. <https://doi.org/10.1038/nrn3289>
- Bahrami, B., Olsen, K., Latham, P. E., Roepstorff, A., Rees, G., & Frith, C. D. (2010). Optimally interacting minds. *Science (New York, N.Y.)*, 329(5995), 1081–1085. <https://doi.org/10.1126/science.1185718>
- Blavatskyy, P. R. (2007). Stochastic expected utility theory. *Journal of Risk and Uncertainty*. Springer. <https://doi.org/10.2307/41761265>
- Bürkner, P.-C. (2017). brms: An R Package for Bayesian Multilevel Models Using Stan. *Journal of Statistical Software*, 80(1), 1–28.

- <https://doi.org/10.18637/jss.v080.i01>
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., ... Riddell, A. (2017). Stan : A Probabilistic Programming Language. *Journal of Statistical Software*, 76(1), 1–32. <https://doi.org/10.18637/jss.v076.i01>
- Chung, D., Christopoulos, G. I., King-Casas, B., Ball, S. B., & Chiu, P. H. (2015). Social signals of safety and risk confer utility and have asymmetric effects on observers' choices. *Nature Neuroscience*, 18(6), 912–916. <https://doi.org/10.1038/nn.4022>
- Ciranka, S., & van den Bos, W. (2019). Social Influence in Adolescent Decision-Making: A Formal Framework. *Frontiers in Psychology*, 10, 1915. <https://doi.org/10.3389/fpsyg.2019.01915>
- de Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior Research Methods*, 47(1), 1–12. <https://doi.org/10.3758/s13428-014-0458-y>
- Delfino, A., Marengo, L., & Ploner, M. (2016). I did it your way. An experimental investigation of peer effects in investment choices. *Journal of Economic Psychology*, 54, 113–123. <https://doi.org/10.1016/J.JOEP.2016.03.005>
- Deutsch, M., & Gerard, H. B. (1955). A study of normative and informational social influences upon individual judgment. *The Journal of Abnormal and Social Psychology*, 51(3), 629–636. <https://doi.org/10.1037/h0046408>
- FeldmanHall, O., & Shenhav, A. (2019). Resolving uncertainty in a social world. *Nature Human Behaviour*, 3(5), 426–435. <https://doi.org/10.1038/s41562-019-0590-x>
- Frey, R., Pedroni, A., Mata, R., Rieskamp, J., & Hertwig, R. (2017). Risk preference shares the psychometric structure of major psychological traits. *Science Advances*, 3(10), e1701381. <https://doi.org/10.1126/sciadv.1701381>
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decision From Experience and the Effect of Rare Events in Risky Choice. *American Psychological Society*, 15(8), 534–539. <https://doi.org/10.1111/j.0956-7976.2004.00715.x>
- Moussaïd, M., Herzog, S. M., Kämmer, J. E., & Hertwig, R. (2017). Reach and speed of judgment propagation in the laboratory. *Proceedings of the National Academy of Sciences*, 114(16), 4117–4122. <https://doi.org/10.1073/pnas.1611998114>
- Moutoussis, M., Dolan, R. J., & Dayan, P. (2016). How People Use Social Information to Find out What to Want in the Paradigmatic Case of Inter-temporal Preferences. *PLOS Computational Biology*, 12(7), e1004965. <https://doi.org/10.1371/journal.pcbi.1004965>
- Sharpe, K. (2018). On Risk and Uncertainty, and Objective versus Subjective Probability. *Economic Record*, 94, 49–72. <https://doi.org/10.1111/1475-4932.12403>
- Toelch, U., & Dolan, R. J. (2015). Informational and Normative Influences in Conformity from a Neurocomputational Perspective. *Trends in Cognitive Sciences*, 19(10), 579–589. <https://doi.org/10.1016/J.TICS.2015.07.007>
- Toyokawa, W., Whalen, A., & Laland, K. N. (2019). Social learning strategies regulate the wisdom and madness of interactive crowds. *Nature Human Behaviour*, 3(2), 183–193. <https://doi.org/10.1038/s41562-018-0518-x>
- van den Bos, W., & Hertwig, R. (2017). Adolescents display distinctive tolerance to ambiguity and to uncertainty during risky decision making. *Scientific Reports*, 7(December 2016), 40962. <https://doi.org/10.1038/srep40962>