Free ranking vs. rank-choosing: New insights on the conjunction fallacy

Torr Polakow

Curiosity Lab, Department of Industrial Engineering,
Tel-Aviv University, Israel

Andrei Teodorescu

Department of Psychology,
Institute for Information Processing and Decision Making (IIPDM),
University of Haifa, Haifa, Israel, 3498838

Jerome R. Busemeyer

Psychological and Brain Science, Indiana University,
Bloomington, Indiana, 47405

Goren Gordon

Curiosity Lab, Department of Industrial Engineering,
Tel-Aviv University, Israel

April 27, 2021

Abstract

It has been consistently shown that when asked to rank options, people often make fallacious judgements. Furthermore, such fallacies can be sensitive to presentation mode. In the first study, we explored a novel type of ranking presentation, namely, choosing between two rank orders of options. To enable a direct comparison of fallacy rates between free ranking and the new presentation mode, we calculated the frequencies of

the two presented rank orders in a free ranking condition. Our analysis shows that people choose the non-fallacious rank order significantly more when asked to choose between two rank orders as compared to freely ranking the possible options. In a second study, we explored whether an agent presenting the rankings to choose from has an effect. To alleviate social biases we used videos of social robots as the presenting agents. We show that rank orders presented by social agents significantly reduce the fallacy rates, compared to rank orders presented without an agent. We discuss the results in view of social decision making theories, wherein a ranking presentation by a social agent is comparable to a consultation with another person regarding free ranking. Our results suggest that fallacious decision making can be mitigated by a social agent presenting rankings to choose from.

Keywords: Decision making, Fallacy, Ranking, Presentation Mode, Social agents

1 Introduction

Decision makers are often required to select from different alternatives. The most prevalent mechanisms are discrete choices between alternatives and ranking. Discrete choices among alternatives result in better differentiation, when compared to rating scales [Wijnen et al.(2015)] , yet are problematic when the number of alternatives is large [Ock, Yi, Ahn, & Jo, (2016)]. Ranking is often based on multiple criteria [Roy(1985)], but is susceptible to various problems, such as weighting issues [Zanakis, Solomon, Wishart, & Dublisch, (1998)] , misinterpretation of personal preferences and internal inconsistencies in the ranking methods [Leoneti and Leoneti(2016)].

In this contribution we address a novel mechanism for decision makers, namely, discrete choice between predetermined rankings. In other words, the decision maker selects between two rank orders of alternatives. This mechanism becomes more useful as different ranking methods appear [Leoneti and Leoneti(2016)], multiple recommendation systems are available [Portugal, Alencar, & Cowan, (2018)] and advisors abound. Thus, decision makers can select not from different alternatives, but from different rankings of alternatives.

We focus on an important aspect of ranking, namely, judgemental fallacies, i.e., judgments that violate classical probability theory. One common fallacy is the conjunction fallacy, which was introduced by Tversky and Kahneman [Tversky and Kahneman(1983)]. They presented to the subjects a hypothetical

personality sketch of a fictitious individual, and asked them to rank order the alternatives according to their probabilities:

Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.

Linda is an elementary school teacher.

Linda is active in the feminist movement. (F)

Linda is a social worker.

Linda is a member of the League of Women Voters.

Linda is a bank teller. (T)

Linda is an insurance salesperson.

Linda is a bank teller and is active in the feminist movement. (T&F)

A conjunction fallacy occurs when people rank option (T&F) higher than (T) or higher than (F) (single conjunction fallacy) or higher than both (double conjunction fallacy). This is a judgemental *fallacy* since in classical probability theory, the probability of the conjunction (T&F) can never be higher than its constituents (T, F). Thus, the actual (real-world) frequency of the conjunction can never be higher than its constituents, yet people do rank the conjunction as more probable. In the original version of the experiment, 85 percent of the respondents ranked option (T&F) as more probable than option (F). Tversky and Kahneman confirmed and replicated these results for different fictional scenarios.

These results have been replicated numerous times with various scenarios, showing that people who freely rank options have high probability of making a judgmental fallacy. However, several studies have shown that the mode of presentation of the scenario and the options can drastically change the fallacy rates. Results by Morier and Borgida [Morier and Borgida(1984)] highlight the importance of the question's framing and its effect on the fallacy rates.

Hertwig and Gigerenzer [Hertwig and Gigerenzer(1999)] replicated the results of Tversky and Kahneman when the subjects were asked to rank or give probabilities to the different options. However, they asked the subject frequency questions such as: "200 women have the same description as Linda above. How many out of the 200 women are bank tellers, bank tellers and feminist etc." In this condition, fallacy rates dropped below 20 percent.

Wedell and Moro [Wedell and Moro(2008)] found that frequency judgments as well as probability judgments produce lower rates of conjunction fallacies than rank orders, and so the effect of using frequency may not be as important as requesting a judgment rather than a rank order.

An interesting study by Charness [Charness(2009)] looked into conjunction fallacy under conditions which allowed a group of individuals to consult with each other. In a treatment with incentives, participants were informed that there was a correct answer and that anyone who chose this correct answer would receive 4\$. In the treatment without incentives, participants were told that they would receive 2\$ for filling out the questionnaire. These experiments were conducted with individuals, pairs, and trios of participants. When no incentive was given, the fallacy rates were 58%, 48% and 26% for single, pairs and trios, respectively. For the incentive condition, the fallacy rates were 33%, 13% and 10% for single, pairs and trios, respectively. These results are particularly interesting showing that incentive and/or cooperation with other participants, reduce the fallacy rates. Thus, the presentation of different rank orders, which is analogous to hearing other participants rank orders, should have a similar effect.

In real world scenarios, another presentation mode is common, namely, discrete choice between two rank orders of options, each having the same alternatives, ranked in different order. This is in contrast to free ranking, probability judgments or direct comparison between two specific alternatives.

We thus set out to compare people’s conjunction fallacy rates of free ranking vs. discrete choice between ranking, where in the latter we mean that the participant has to choose between presentation of two specific rank orders. Thus, the goal of the current studies is to better understand the influence choosing between rankings has on fallacy rates, i.e., whether this prevalent presentation mode can mitigate conjunction fallacies, and what attributes of the options’ presentation are important for this effect.

There are several leading models of conjunction fallacies, such as the averaging model [Juslin, Nilsson, & Winman, (2009)] , Costello’s probability plus noise model [Costello, 2009; Costello, Watts, & Fisher, (2018)] and quantum cognition models [Busemeyer, Wang, & Shiffrin, (2015)]. However, none have specifically addressed the presentation mode of rank choosing. A-priori, these models generate the hypothesis that rank choosing should follow the same fallacy rate as free ranking: participants assign higher probability to rank orders, based on free ranking probabilities, and thus should choose rank order with fallacy more than without.

Thus, the main null hypothesis states that there will not be a significant difference in fallacy rates between free ranking and rank-choosing (H1). Moreover, the presenter of the choices has not been considered before, except from Charness [Charness(2009)], which showed that consulting with people significantly

reduces the fallacy rates. Thus, a second null hypothesis states that a social presenter of the rank-choosing will significantly reduce fallacy rates (H2). We conducted two studies to test these hypotheses.

2 Methods

2.1 Questionnaire

We used the following personality sketches of two fictitious individuals, Emma and Liz, followed by a set of occupations and avocations associated with each. For Emma, the two characteristics that we asked about were “manager” (M) and “likes to exercise” (E) and for Liz they were “programmer” (Pr) and “likes to paint” (Pa). The two questions are given below:

Q1: *Emma lives in an apartment with her two cats in the center of the city and is an outgoing person. She takes yoga classes at the gym three times a week, enjoys reading science-fiction books and volunteers at least once a week in an animal shelter.*

Emma is a manager (M).

Emma likes to exercise (E).

Emma is a manager and likes to exercise (M&E).

Q2: *Liz lives in a Victorian home in Oakland. She is an analytical thinker and works at a start-up company. She tries to attend a few classes at the gym every week. She is very ambitious at her job and also enjoys cooking very much and is quite good at it.*

Liz is a programmer (Pr).

Liz likes to paint (Pa).

Liz is a programmer and likes to paint (Pr&Pa).

In the original Linda story, one of the options was consistently assigned by participants a higher probability than the other ($T > F$) and was thus also ranked higher [Bhatia(2017), Hofmann et al.(2018)]. The conjunction fallacy that ensued was predominantly a “single-conjunction”, i.e., the prevalent conjunction fallacy ranking was $T > T\&F > F$.

In the current study, we deliberately designed the questionnaire to explore two scenarios: Q2 represents a Linda-like scenario in which $Pr > Pa$, whereas Q1 represents a more equi-probable scenario in which $M \simeq E$.

2.2 Participants

In our study, the participants were Amazon Mechanical Turk (AMT) users. Amazon Mechanical Turk is a common method to collect large amounts of data [You & Robert Jr., (2018); Canning, Donahue, & Scheutz, (2014)]. While there

has been some debate about the validation of MTurk users in the scientific community, recent studies validated that MTurk users are an appropriate sample for research, when cautionary actions are taken [Hauser & Schwartz, (2016); Berinsky, Huber, & Lenz, (2012)].

We asked for Amazon Mechanical Turk users that were native English speakers from North America that had completed at least 100 tasks with an approval rate $> 95\%$. First the participants were asked to give their consent, age, gender and education. Then, each participant was presented with the questionnaire on Qualtrics (Qualtrics, Provo, UT). While completing the questionnaire, the participants had to pass a screening check, namely, a trap question, in which we asked “What does Liz enjoy doing?”, instead of ranking or choosing a ranking.

The study was approved by our Institutional Review Board.

3 Study 1: Free ranking vs. Rank-choosing

3.1 Conditions

The study was designed to have two conditions: RANK and CHOOSE. In the RANK condition, the participants were asked to *rank* the statements according to their perceived probability. In the CHOOSE condition, participants were presented with two rankings; one that represented a ranking with no fallacies, e.g., the conjunction statement was ranked lowest and; one that represented a conjunction fallacy, e.g., the conjunction statement was ranked higher than one of the other statements. Participants in the CHOOSE condition were asked to choose which ranking they agreed with. The side on which the fallacious ranking was presented alternated randomly between participants and across questions within subjects.

We administered each condition twice, once for Q1 followed by Q2 and once in the other order. Due to the screening of AMT, the following number of qualified participants was achieved: RANK Q1Q2 ($n = 122$, female=59, age=32 ($s = 0.9$)), RANK Q2Q1 ($n = 90$, female=54, age=36 ($s = 1.1$)), CHOOSE Q1Q2 ($n = 146$, female=77, age=37 ($s = 0.9$)), CHOOSE Q2Q1 ($n = 85$, female=50, age=35 ($s = 1.1$)).

3.2 Results

The analysis showed that while the RANK condition conformed to previous studies, with 61% fallacy rate, the CHOOSE condition had a significantly lower fallacy rate, 32%. However, due to different ex-ante probabilities of each condition, a direct statistical significance test of the two conditions was not possible. Hence, we performed the following procedure to analyze their difference. We first tested whether participants' choices were significantly different than random selections. We then calculated the *rank order* in both conditions, i.e., binary relations between each ranking-pairs, and then directly tested for differences between these rank orders.

3.2.1 RANK condition

We first analyzed the ex-ante probabilities of rankings and then tested for difference from random ranking. In the current setup, there were six possible rankings, from which four were fallacious. Therefore, the ex-ante random ranking probabilities resulted in 2/3 chance of ranking fallaciously (1/6 for each fallacious option).

A Fisher-exact test was performed on the fallacious vs. non-fallacious options and random ranking. The analysis shows that there was no significant difference between participant's ranking and random ranking (Q1: $p = 0.838$, Q2: $p = 0.057$, Fisher's exact test). However, a deeper analysis of the relative frequency of each ranking showed that the rankings were not equi-probable ($\chi^2 = 20.17, p = 0.001$).

Hence, while the fallacy rate was not significantly different than random ranking, the ranking themselves were not random.

3.2.2 CHOOSE condition

In this condition, participants had a binary choice to make, two possible rankings to choose from, one fallacious and one not. Therefore, the ex-ante probabilities of random selections were 1/2. A binomial test was performed on these binary choices to check if they differed from random and revealed that participants chose the non-fallacious ranking significantly more than random selection ($p < 0.0001$).

These results, namely, that the RANK condition *was not* significantly different than random ranking and the CHOOSE condition *was* significantly different

than random choice, do not support hypothesis H1, that the two conditions will not be significantly different.

3.2.3 Ranking order

To enable a direct comparison between the two conditions, we analyzed the *ranking orders*, by investigating the binary relations between every two rankings [Ock et al.(2016)]. In the RANK condition, we compared the ranking frequency, whereas in the CHOOSE condition we analyzed the relative proportion of each rank-pair selection.

For example, in Q1, the options were M (manager), E (exercise) and M & E (manager and exercise). We thus compared between ranking pairs, e.g., $M > E > M \& E$ and $M > M \& E > E$. For each ranking pair in the CHOOSE condition, all of which included both a fallacious and a non-fallacious ranking, we compared the frequency with which those specific rankings were chosen in the RANK condition. We thus treated each ranking pair as a binary decision, even in the RANK condition, and could perform a direct χ^2 test between the two conditions. Our analysis showed that the two conditions were significantly different for the ranking pairs presented ($\chi^2 = 12.18, p < 0.001$).

3.2.4 Questions and order effects

Comparing the two questions, Q1 and Q2, we found that both show approximately the same trend of RANK having a significantly higher fallacy rates than CHOOSE (Q1 RANK 65% and CHOOSE 29%, Q2 RANK 57% and CHOOSE 36%). A deeper analysis showed that while the question itself (Q1,Q2) was not a significant predictor of fallacy rate, there was a significant interaction between condition and question, such that Q1 had higher fallacy rates than Q2 in the RANK condition, whereas the opposite occurred in the CHOOSE condition ($\chi^2 = 79, p < 0.001$, with $\beta(COND) = 0.87, p < 0.001, \beta(Q) = -3.1, p = 0.112, \beta(COND \times Q) = 0.65, p = 0.02$, logistic regression).

Analyzing the specific free ranking the participants performed, we found that for Q1 rankings with $M > E$ (49.5%) were very similar to $E > M$ (50.5%). On the other hand, for Q2 we found a different picture, with $Pr > Pa$ (68%) being much more frequent than $Pa > Pr$ (32%). These results substantiate our initial design of the different questions.

We analyzed the effect of order on the conjunction fallacy, i.e., whether the first encountered question (regardless of type Q1 or Q2) had a different fallacy

rate than the second question, and found no significant effect ($\chi^2 = 0.2, p = 0.64$)

3.3 Discussion

The results of the comparison were highly significant, and did not support the null hypothesis H1. They showed that asking participants to choose from two rankings drastically reduced the conjunction fallacy rate. Current mathematical models, such as the averaging model [Juslin et al.(2009)], the probability theory plus noise [Costello(2009)] and quantum cognition [Busemeyer et al.(2006)], have not addressed this particular presentation mode. However, the a-priori hypothesis, that people would select ranking according to their “simulated” free-ranking probabilities, had been shown to not be supported.

While frequency judgments also lower the fallacy rate [Wedell and Moro(2008)], the currently observed rates are still much higher than the frequency judgments, and do not directly relate to the selection presentation mode. Semantic [Bhatia(2017)] and co-occurrence [Hofmann et al.(2018)] analyses, may explain the difference between the two questions, Q1 and Q2, for example “analytical thinker/ambitious” is semantically more associated with “programmer” than “likes to paint”. Furthermore, representativeness of the options presented has been shown to affect meta-cognitive processes related to conjunction fallacy rates [Dujmović, Valerjev, & Bajšanski, (2020)]. However, these approaches do not directly relate to the presentation mode.

It has been recently shown that group reasoning can lessen the fallacy rates, as well as using paradigmatic reasoning, as compared to narrative reasoning [Donati, Guazzini, Gronchi, & Smorti, (2019)]. This social effect, also reported by Charness [Charness(2009)], does have similar reduction in conjunction fallacy rate, when considering no incentives and a “trio” formation. In other words, one can think of the two rankings as the opinions of two fellow participants, which, together with the actual participant, constitute a “trio” formation”, thus reducing the fallacy rate.

In order to check this hypothesis we conducted a second study in which we used artificial agents, namely, social robots, to present the rankings participants can choose from. Hence, one robot always answered according to the conjunction fallacy, whereas the other robot answered with no fallacy. To maintain consistency and embodiment, for both questions, the robots maintained their side, e.g., for both Q1 and Q2 the robot with fallacies always stayed on the left

side. To control for this consistency-related condition, we ran the CHOOSE condition again, but now without alternating the sides of ranking presentations.

4 Study 2

4.1 Conditions

The study was designed to better understand the attributes that make rank choosing mitigate fallacy rates. In real-life scenarios, the system or agent that presents the ranking to choose from is usually consistent, i.e., one system or agent usually either makes fallacies or not; it does not alternate randomly between fallacious rankings and non-fallacious ones.

For this purpose we introduced the CONSISTENT condition, which differed from the CHOOSE condition from study 1 only in the consistency of the ranking presentation. In other words, the CONSISTENT condition’s presentation of the fallacious rank ordering was the same for both presented questions, i.e., it did not alternate randomly as in the CHOOSE condition. Effectively, the only change between the CHOOSE and CONSISTENT conditions was in the second presented question, in which the fallacious ranking was presented in the same position (left/right) as in the first question for the CONSISTENT condition and random position in the CHOOSE condition.

Another important aspect is the presenter of the ranking, i.e., whether it is a “system” or a “social agent” [Logg, Minson, & Moore, (2019); Kunkel, Donkers, Michael, Barbu, & Ziegler, (2019)]. To study the influence of the presenter, we introduced the SOCIAL condition, in which a video of a social robot presenting the rank ordering came before the participants’ choosing. In each video, the presenter asked the robots a question. Then each robot gave its rank verbally according to its fallacy. For example, a robot with a single conjunction fallacy in question Q1 ranked the options: M highest, then M&E and lastly E, Fig. 1(right). Subtitles were displayed for everything that was said, by the presenter or robots.

In summary, the CONSISTENT condition was contrasted to the CHOOSE condition, to study whether a consistent presentation of ranking influences the fallacy rates, and to the SOCIAL condition to study whether the presenter of the rankings influenced the fallacy rates.

In the SOCIAL condition, in each video two NAO robots (1 blue and 1 red) were presented, one on the left and one on the right with a presenter in the

middle, Fig. 1. We counterbalanced the color and fallacious side of the robots across participants.

Since we have shown that there is no order effect, we administered each condition once for Q1 followed by Q2. We had in total 240 and 274 participants for the SOCIAL and CONSISTENT conditions, respectively. However, due to the screening of AMT, the following number of qualified participants was achieved: SOCIAL ($n = 227$, female=118, age=37 ($s = 0.9$)), CONSISTENT ($n = 160$, female=89, age=38 ($s = 0.9$)).

Figure 1: Screenshot from one of the videos that was presented to the participants. In this image, the left robot behaved according to a non-fallacious ranking, and the right robot behaved according to single conjunction fallacy.

4.2 Results

The analysis showed that the CONSISTENT condition had similar fallacy rates to the CHOOSE condition, with 31% fallacy rate ($\chi^2 = .13, p = 0.71$). However, the SOCIAL condition had a significantly lower fallacy rate, 24%, compared to the CONSISTENT condition ($\chi^2 = 6.79, p = 0.009$).

Questions: Comparing the two questions, we found that both showed approximately the same trend (Q1 CONSISTENT 32% and SOCIAL 28%, Q2 CONSISTENT 30% and SOCIAL 20%), with no significant interaction.

4.3 Discussion

The analysis showed that if a social agent, in this case a social robot, presents the rankings, fallacy rates drop even more than mere rank choosing, thus supporting hypothesis H2. The fallacy rates are now on-par with those reported by Charness [Charness(2009)] for “trio”. Our CONSISTENT control condition, which resulted in almost identical fallacy rates as CHOOSE, rules out an effect due to consistency of presentation.

Since the only difference between the conditions was the actual appearance of a social agent, we can conclude that not only the presentation mode influences fallacy rates, but also the presenter itself, Fig. 2.

One can consider the robots as agents giving advice. It has been shown that people adhere more to advice when they think it comes from an algorithm than from a person, even when the algorithm’s process is unknown [Logg et al.(2019)]. Hence, one can consider the two robots as artificial agents giving advice, with an underlying algorithm. This may influence how participants viewed the social robots, compared to human presenters [Charness(2009)].

Moreover, we did not comment directly why each robot selected its options, which has been shown to increase trust in similar systems [Kunkel et al.(2019)]. The robot also did not explain its own behaviour, which has been shown to improve collaboration between people and robots [Tabrez and Hayes(2019)]. This lack of explanation may have influenced how people perceived the robots’ advice.

5 General discussion

Previously, the two main mechanisms for decision makers, i.e., ranking and discrete choice between alternatives, have been compared. Ock et al. [Ock et al.(2016)] presented a similar conversion from ranking to discrete choice was made, but it compared choices between *alternatives* and not ranking. It was shown that ranking and paired-comparison are consistent.

Conversely, we have shown that presenting two rankings to choose from results in a drastic reduction in conjunction fallacy rates. We also showed that having a social agent presenting the rankings strengthen the mitigation effect of the fallacy.

Hence, there are two cumulative effects when social agents present their options: the first, and vastly more effective, is the fact that there is a choice to make and not ranking. The second effect, which is significant but has a smaller

Figure 2: Fallacy rates as a function of conditions, including those reported in (Charness, 2009) for trio with no incentive. $***p < 0.001$, $**p = 0.004$ (CHOOSE-SOCIAL), $p = 0.0094$ (CONSISTENT-SOCIAL).

effect size, is the appearance of a social agent.

The results reported by Charness [Charness(2009)] can be viewed as presentation of options by social agents. Several important conjectures come from our presented results.

First, that voicing one’s reasons for its opinions is unnecessary for an effective mitigation of the conjunction fallacy [Kunkel et al.(2019), Tabrez and Hayes(2019)]. In Study 2, the robots did not explain their rankings, only expressed them. Furthermore, even without a social agent, the mere presentation of ranking options reduced the fallacy rate.

Second, humans are not the only social agents that can influence one’s opinions. The remarkably similar fallacy rates reported by Charness [Charness(2009)] and in our SOCIAL condition shows that having *any* social agent is enough. While one option for why social agency is important may be its consistency, our CONSISTENT condition results show that it is not the case. Alternatively, the embodiment of the social agent may be the major factor.

Third, Charness [Charness(2009)] has raised the truth-wins standard as a possible explanation for the phenomenon. his explanation claims that simply encountering a more “truthfull”, i.e., non-fallacious, option results in decrease in one’s own fallacious decisions. Our results partly support this explanation, as we have shown that the mere presentation of the rankings, without a social

agent nor reasoning, suffices to reduce fallacy rates.

A complementary explanation maybe due to dual-process or “slow- and fast-brains” [Kahneman(2011), Andersson et al.(2020)]. It has been recently shown that in order to avoid the conjunction fallacy, additional neurocognitive resources, mainly in the fronto-parietal network, are required [Andersson et al.(2020)]. One option is that when asked to freely rank, participants engage fast intuitive processes, wherein the fast and less rational brain takes over. On the other hand, when asked to choose between rankings, participants engage in a more deliberative mode, comparing the two alternatives, possibly also to one they have already formulated. Thus, during this deliberation, their slow and rational brain takes over and results in much smaller fallacy rates. The fronto-parietal network has also been implicated in social interaction [Lahnakoski et al.(2012), Arioli and Canessa(2019)]. Thus, social interaction may increase the effects of the slow and rational brain even more.

6 Conclusions

We have reported on a new type of presentation, wherein participants have to choose between two rankings. Our analysis shows that this mode of presentation drastically reduces conjunction fallacy rates. Furthermore, presentation by an (artificial) social agent reduces the fallacy rates to those reported by conferring with two other human companions.

While we have shown that mere presentation of the rankings is enough to induce reduction in fallacy rates, supporting the truth-wins standard, a cumulative effect of social agents is harder to explain. Future studies will endeavor to unveil the mechanisms behind these effects.

Conflict of Interest Statement

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Author contribution

G. G. and J. R. B. developed the study concept. T. P. and A. T. contributed to the study design. Data collection was performed by T. P. and data analysis was performed by T. P. and G. G. T. P. drafted the manuscript, and A. T., J. R. B. and G. G. provided critical revisions. All authors approved the final version of the manuscript for submission.

Funding

This study was supported by Israel-US Binational Science Foundation grant No. 2016262.

Open Practices Statement

The data and materials for all experiments are available at <https://osf.io/5ta2w/> and none of the experiments was preregistered.

References

- [Andersson et al.(2020)] Andersson, L., Eriksson, J., Stillesjö, S., Juslin, P., Nyberg, L., Wirebring, L.K., 2020. Neurocognitive processes underlying heuristic and normative probability judgments. *Cognition* 196, 104153. URL: doi:10.1016/j.cognition.2019.104153.
- [Arioli and Canessa(2019)] Arioli, M., Canessa, N., 2019. Neural processing of social interaction: Coordinate-based meta-analytic evidence from human neuroimaging studies. *Human Brain Mapping* 40, 3712–3737. URL: doi:10.1002/hbm.24627.
- [Bhatia(2017)] Bhatia, S., 2017. Associative judgment and vector space semantics. *Psychological Review* 124, 1–20. doi:10.1037/rev0000047. place: US Publisher: American Psychological Association.
- [Busemeyer et al.(2006)] Busemeyer, J.R., Matthew, M.R., Wang, Z., 2006. A quantum information processing explanation of disjunction effects, in: *The*

28th Annual Conference of the Cognitive Science Society & the 5th International Conference of the Cognitive Science, Mahwah, NJ: Erlbaum. pp. 131–135.

- [Charness(2009)] Charness, G., 2009. On the Conjunction Fallacy in Probability Judgment: New Experimental Evidence Regarding Linda. Economics Working Paper Archive. The Johns Hopkins University, Department of Economics.
- [Costello(2009)] Costello, F.J., 2009. How probability theory explains the conjunction fallacy. *Journal of Behavioral Decision Making* 22, 213–234. doi:10.1002/bdm.618.
- [Hertwig and Gigerenzer(1999)] Hertwig, R., Gigerenzer, G., 1999. The ‘conjunction fallacy’ revisited: how intelligent inferences look like reasoning errors. *Journal of Behavioral Decision Making* 12, 31.
- [Hofmann et al.(2018)] Hofmann, M.J., Biemann, C., Westbury, C., Murusidze, M., Conrad, M., Jacobs, A.M., 2018. Simple Co-Occurrence Statistics Reproducibly Predict Association Ratings. *Cognitive Science* 42, 2287–2312. URL: doi:10.1111/cogs.12662.
- [Juslin et al.(2009)] Juslin, P., Nilsson, H., Winman, A., 2009. Probability theory, not the very guide of life. *Psychological review* 116, 856.
- [Kahneman(2011)] Kahneman, D., 2011. Thinking, fast and slow. Macmillan.
- [Kunkel et al.(2019)] Kunkel, J., Donkers, T., Michael, L., Barbu, C.M., Ziegler, J., 2019. Let Me Explain: Impact of Personal and Impersonal Explanations on Trust in Recommender Systems, in: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, Association for Computing Machinery, Glasgow, Scotland Uk. pp. 1–12. URL: doi:10.1145/3290605.3300717.
- [Lahnakoski et al.(2012)] Lahnakoski, J.M., Glerean, E., Salmi, J., Jääskeläinen, I.P., Sams, M., Hari, R., Nummenmaa, L., 2012. Naturalistic fMRI Mapping Reveals Superior Temporal Sulcus as the Hub for the Distributed Brain Network for Social Perception. *Frontiers in Human Neuroscience* 6. URL: doi:10.3389/fnhum.2012.00233. publisher: Frontiers.

- [Leoneti and Leoneti(2016)] Leoneti, A.B., Leoneti, A.B., 2016. Considerations regarding the choice of ranking multiple criteria decision making methods. *Pesquisa Operacional* 36, 259–277. doi:10.1590/0101-7438.2016.036.02.0259. publisher: SOBRAPO.
- [Logg et al.(2019)] Logg, J.M., Minson, J.A., Moore, D.A., 2019. Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes* 151, 90–103. URL: doi:10.1016/j.obhdp.2018.12.005.
- [Morier and Borgida(1984)] Morier, D.M., Borgida, E., 1984. The conjunction fallacy: A task specific phenomenon? *Personality and Social Psychology Bulletin* 10, 243–252. doi:10.1177/0146167284102010.
- [Ock et al.(2016)] Ock, M., Yi, N., Ahn, J., Jo, M.W., 2016. How many alternatives can be ranked? A comparison of the paired comparison and ranking methods. *Value in Health* 19, 655–660. ISBN: 1098-3015 Publisher: Elsevier.
- [Roy(1985)] Roy, B., 1985. *Methodologie multicritere d’aide a la decision*. Economica. Issue: BOOK.
- [Tabrez and Hayes(2019)] Tabrez, A., Hayes, B., 2019. Improving Human-Robot Interaction Through Explainable Reinforcement Learning, in: 2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI), pp. 751–753. doi:10.1109/HRI.2019.8673198. iISSN: 2167-2148.
- [Tversky and Kahneman(1983)] Tversky, A., Kahneman, D., 1983. Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review* , 293–315.
- [Wedell and Moro(2008)] Wedell, D.H., Moro, R., 2008. Testing boundary conditions for the conjunction fallacy: Effects of response mode, conceptual focus, and problem type. *Cognition* 107, 105–136. doi:10.1016/j.cognition.2007.08.003.
- [Wijnen et al.(2015)] Wijnen, B.F., van der Putten, I.M., Groothuis, S., de Kinderen, R.J., Noben, C.Y., Paulus, A.T., Ramaekers, B.L., Vogel, G.C., Hiligsmann, M., 2015. Discrete-choice experiments versus rating scale exercises to evaluate the importance of attributes. *Expert review*

of pharmacoeconomics & outcomes research 15, 721–728. ISBN: 1744-8379
Publisher: Taylor & Francis.