Stress, Addiction, and Artificial Intelligence

Mohamed A. Hussein, Stanford University, USA

Szu-chi Huang, Stanford University, USA

90.2% of people with mental illness receive no treatment. Can AI improve access to mental health services? We find that consumers avoid AI (vs. human) advisers and rate them as inferior across mental health domains, such as stress and addictive behaviors (e.g. porn addiction) because AI providers are less warm.

[to cite]:

Mohamed A. Hussein and Szu-chi Huang (2020) ,"Stress, Addiction, and Artificial Intelligence", in NA - Advances in Consumer Research Volume 48, eds. Jennifer Argo, Tina M. Lowrey, and Hope Jensen Schau, Duluth, MN : Association for Consumer Research, Pages: 1155-1159.

[url]:

http://www.acrwebsite.org/volumes/2657172/volumes/v48/NA-48

Subjective to Objective Value of Humans and Algorithms

Chairs: Gizem Yalcin, Erasmus University, The Netherlands
Evan Weingarten, Arizona State University, USA

Paper #1: Thumbs Up or Down: Consumer Reactions to Decisions by Algorithms Versus Humans

Gizem Yalcin, Erasmus University, The Netherlands
Sarah Lim, Cornell University, USA
Stefano Puntoni, Erasmus University, The Netherlands
Stijn van Osselaer, Cornell University, USA

Paper #2: People Prefer Forecasting Methods Similar to the Event Being Predicted

Lin Fei, University of Chicago, USA
Berkeley J. Dietvorst, University of Chicago, USA

Paper #3: Stress, Addiction, and Artificial Intelligence

Mohamed Ayman Hussein, Stanford University, USA
Szu-chi Huang, Stanford University, USA

Paper #4: Experts Outperform Technology in Creative Markets

Evan Weingarten, Arizona State University, USA
Michael Meyer, University of California – San Diego, USA
On Amir, University of California – San Diego, USA
Amit Ashkenazi, Fiverr, USA

SESSION OVERVIEW

With the accelerated demand for new technologies, algorithms / artificial intelligence (AI) have integrated diverse and refined skills. These advancements make them a viable alternative to using humans to complete the same tasks. Algorithms can outperform humans on medical diagnoses (Hutson 2017), taste-based recommendations (Yeomans et al. 2019), predictions (Grove et al. 2000), and games (Hosanagar 2019).

Despite advancements in these new technologies in businesses, however, the literature diverges on the extent consumers embrace or dislike algorithms. It is still unclear from the literature when consumers are open to using algorithms (e.g., Dietvorst et al. 2015; Logg et al. 2019). This disagreement raises important questions about the dynamics between humans and their valuation of technology: What factors affect consumers' reactions towards decisions by algorithms (vs. humans)? Do consumers value decisions differently when generated by algorithms or humans? Does the framing of algorithms change consumers' willingness to use it? Under which situations are consumers more likely to (de)value algorithms? Can algorithms objectively outperform humans in domains involving aesthetics?

This session aims to identify the boundary conditions of when consumers appreciate algorithms more than humans (subjective valuation), and when algorithms do or do not objectively outperform humans (objective). In the first paper, **Yalcin et al.** study consumers' reactions to decisions that are made by algorithms versus humans. Across seven studies, they reveal less positive reactions to favorable outcomes (e.g., acceptances) by algorithms (vs. humans), whereas they find no increased negative responses to algorithms when the outcome is unfavourable (e.g., rejections). They also show that these differences are explained by a shift in perceptions depending on the valence of the decision outcome. **Lin** and **Dietvorst** explore how the way algorithms are described changes consumers' likelihood of using them. Across seven studies and different prediction methods (e.g., navigation), they demonstrate that consumers prefer methods of prediction that replicate the event in question even when doing so is counterproductive. Next, **Hussein** and **Huang** investigate what type of professional help consumers value more when dealing with stress and compulsive behaviors. In four studies, the authors find that consumers perceive AI (vs. human) mental health providers to be less warm, which harms their perceived competence, resulting in higher aversion of AI providers. Finally, the paper by **Weingarten et al.** tests whether AI outperform human experts in creative markets, specifically in logo design. The authors consistently find that the quality of the logos from human experts outperform those from artificial intelligence.

This session provides timely insight into objective and subjective valuation of algorithms and humans, an increasingly important topic given the proliferation of this technology in consumer contexts. Our proposed session directly addresses the ACR 2020 conference theme as it examines the value of deep thinking, rational thought, and reason for humans relative to algorithms in a new era of machine intelligence. We believe that this proposed session should be of interest to a broad audience of scholars and practitioners working on judgment and decision-making, advice/recommendations, aesthetics, and new technologies. All projects have at least four completed studies.

Thumbs Up or Down: Consumer Reactions to Decisions by Algorithms Versus Humans

EXTENDED ABSTRACT

Companies are increasingly adopting algorithms to make decisions that affect existing and potential customers, such as accepting and rejecting applications. Today, algorithms are commonly used to decide who a company should hire (e.g., JetBlue) or provide services to (e.g., rayatheapp.com). This growing trend calls for marketing researchers to gain a better understanding of customers' reactions to decisions made by algorithms and humans. Previous research has predominantly focused on how individuals choose between an algorithmic and human service provider (Castelo et al. 2019; Logg et al. 2019; Longoni et al. 2019). Unlike this line of research, we investigate the responses of individuals as a recipient of decisions made by either an algorithm or a human.

In this research, we propose that customers react to decision-makers (algorithms vs. humans) differently depending on the valence of decision outcomes, namely whether they are accepted or rejected by a firm. Specifically, we hypothesize that customers react less positively (e.g., less perceived self-worth, less positive attitudes towards the firm), to an algorithm than a human decision-maker in the case of favorable decisions (e.g., acceptances), whereas such a negative reaction to an algorithm (vs. a human) would be attenuated in the case of unfavorable decisions (e.g., rejections). Our theorizing is based on attribution theory demonstrating that individuals tend to attribute their successes and failures in a self-serving way (Halperin et al. 1976). Namely, people tend to take credit for their success but blame others for failures so as to defend their self-esteem. Accepted customers would be motivated to view the positive outcome as a result of their individual characteristics, thereby reacting more positively to a human (vs. an algorithmic) decision-maker, who is perceived as more capable of incorporating individuals' uniqueness into the decision. Rejected customers, however, would attribute the unfavorable outcome to decision-makers (regardless of who the decision-maker is), viewing a human decision-maker as less objective and an algorithm as more ignorant of their individual uniqueness.

We tested our predictions across seven experiments (N = 3,535) and showed that customers react differently to favorable (vs. unfavorable) decisions that are told to be made by algorithms and humans across various contexts (e.g., dating websites, bank loans), In studies 1 and 2, we demonstrated that participants felt lower feelings of self-worthiness when favorable decisions were made by an algorithm (vs. a human). This relatively negative reaction to algorithms is, however, mitigated when the decision outcome is unfavorable. Extending these findings, study 3 revealed the same pattern of customers' reactions towards firms: participants demonstrated less positive attitudes towards the firm when an algorithm (vs. a human) accepts them, but such a relatively negative reaction to algorithmic decision-making was mitigated in the case of unfavorable decisions.

In studies 4a, 4b and 5, we aimed to understand possible drivers of such an interaction effect. In studies 4a-b, we revealed the negative impact of disclosing algorithmic decision-makers. Specifically, when the decision-maker was not explicitly mentioned by the company, participants reacted similarly as they did to a human decision-maker. Next, study 5 tested whether participants would react more positively to humans even when they do not actively engage in decision-making but monitoring the process. In line with our theorization, our results demonstrated that the relatively positive effect of human decision-making is driven by knowing that a human *actively made a decision* to accept the applicants instead of passively monitoring the evaluation process. Finally, in study 6, we directly examined the psychological mechanism underlying different reactions to algorithmic versus human decision-makers. In these studies, we showed that the differences in customers' reactions stem from how they perceive these two decision-makers and that their perceptions of algorithms and humans (i.e., perceived objectivity, consideration of applicants' uniqueness).

We believe that our research makes several contributions. Extending the previous work that has predominantly studied situations where people *choose whether to rely on algorithms or humans*, we study situations where *customers are recipients of decisions made by algorithms versus humans*. Furthermore, the current research demonstrates how motivated attribution plays a role in people's perception of algorithmic versus human decision-making. From a practical perspective, our work offers important managerial guidance. Managers are often worried about deploying algorithms in customer-facing functions as they fear algorithms to amplify customers' negative reactions in the case of unfavorable decisions. Our findings, however, suggest that managers should be more concerned about deploying algorithms in the case of favorable experiences as it can result in less positive customer reactions. Together, our work provides valuable insights on how firms can effectively communicate decision outcomes to customers.

People Prefer Forecasting Methods Similar to the Event Being Predicted

EXTENDED ABSTRACT

Consumers face many scenarios where they have to make predictions by choosing between prediction methods. People can check multiple navigation tools for traffic predictions, analysts choose between models to make financial market predictions, and patients can choose between doctors and AI for medical diagnosis. It is clear from past research that people don't always pick the best performing method (Arkes et al. 1986; Dietvorst et al. 2015; Yeomans et al. 2019). However, it is still unclear how consumers choose between prediction methods. What features do people look for in prediction methods?

In this paper, we find that the more similar a prediction method is to the event being predicted (e.g. in its outcome distribution, process, etc.), the more people like it, even when it does not perform as well as alternatives. Humans often learn by mimicking and mirroring others' actions (Meltzoff and Moore 1977), and people tend to "over-imitate" - copying actions that are unnecessary to accomplish the given goal (McGuigan et al. 2011; Hoehl et al. 2019). We propose that this innate behavior transfers to consumers' preferences for prediction methods: consumers prefer prediction methods that best resemble the event being predicted. However, this preference can lead consumers to prefer prediction methods that offer suboptimal performance. For example, when the outcomes of an event are at least partially determined by random chance (Fox and Ülkümen 2011), mimicking the event in question to predict its outcome will result in overfitting of random error. In a set of 7 studies, we investigate consumers' preference for similar prediction methods using both incentivized studies and real-world consumer scenarios.

In Study 1, we present evidence that people like a prediction method to be similar to the event in question by having participants predict the outcome of a die roll. We chose this task because rolling a die has a clear process to be replicated, and its outcome is determined by random chance. Participants were tasked with predicting the outcome of a 7-sided die roll (sides 1,2,3,4,5,6,7). Participants chose between two prediction methods: one that always predicts "4" (constant), and another that chooses a random number between 1 and 7 (similar). Participants learned that they would receive a linearly increasing bonus depending on the difference between their prediction and the actual outcome ($0.21 for a perfect prediction, -$.03 for each unit of error). With this payment scheme, choosing "4" (i.e. the constant method) offers the highest possible return in expectation. Thus, choosing the similar method instead of the constant method is costly. However, we found that a substantial proportion of participants (44.5%) chose the similar method. Even among participants who passed a comprehension check by reporting that choosing "4" produced the highest expected earnings, 36.8% chose the similar method, which suggests that consumers' preference for a similar method is not due to ignorance.

In Studies 2 and 3, we found that the more similar a prediction method is to the event in question, the more people like it. In study 2, we offered participants a choice between two randomly selected methods out of a set of three: the similar and constant options from Study 1, and a "mixed" option that was a combination of the two. The mixed option used the similar method with a probability of 50% and the constant method with a probability of 50%. Participants were more likely to select the similar method both when the alternative was the constant method (58%) and the mixed method (56.2%). In Study 3, we show that method similarity can be increased along multiple dimensions. Specifically, we hypothesized that participants would prefer a method that uses the same process as the event (rolling a die instead of drawing a marble) in addition to matching the outcome distribution. Participants rated the constant method and 4 methods that varied on their proximity to the event being predicted (a fair 7-sided die roll) on 5-point scales: rolling a die that is skewed towards the optimal answer (process resemblance), drawing a marble with values 1-7 from a jar (outcome resemblance), rolling a die with the same numbers as the focal die (perfect resemblance), and drawing a marble with values skewed towards the optimal answer (no resemblance). Participants also rated the similarity of each prediction method to the event being predicted on two 5-point scales. Consistent with our hypothesis, participants rated the prefect resemblance method higher than all other methods (t's$(404) \geq 7.16$, p's$<.001$). Further, participants ratings of the similarity between

each method and the focal die roll mediate the differences in method ratings [10.29,35.99]. These results suggest that the more similar a prediction method is to the event in question, the more people like it.

In the remaining studies, we extend these findings to real-world consumer scenarios. In study 4, participants made an incentivized choice between two prediction methods to predict a consumer's rating of a movie. Participants chose between a similar method that finds the person who gave the most similar responses to the focal consumer in a survey of movie preferences and uses their rating, and an alternative that uses the average rating of the movie among all consumers. 71% of participants chose the similar method even though it performs 1.5 times worse than the alternative. In study 5, participants preferred a navigation tool that based a travel time prediction on the most similar trip in its data over another that averaged among many somewhat similar trips. In follow up studies (Studies 6 and 7), we investigate how companies can frame the prediction methods that they offer consumers to be more similar to the event in question.

People's systematic preference for similar prediction methods provides novel insight into consumer decision making, and suggests interventions to boost consumers' use of a prediction method. For example, it suggests that framing a prediction method as similar to the event in question will make it more palatable to consumers, which can help companies boost the use of prediction methods that they offer consumers (e.g. recommendation systems, calculators, etc.).

Stress, Addiction, and Artificial Intelligence

EXTENDED ABSTRACT

About 90.2% of people suffering from mental illness worldwide do not receive adequate treatment (Alonso et al. 2018). A solution to this problem is therapy powered by artificial intelligence (AI), which is scalable at a low cost, convenient, and private. Are consumers open to adopting AI to manage their mental health? We first investigate if consumers would choose an AI adviser over other types of human advisers present in the marketplace to treat anxiety and stress (Studies 1a-e). We find that consumers are averse to choosing an AI provider to manage their mental health. We replicate our finding in a highly stigmatized context—porn addiction (Study 2)—and using additional dependent variables, such as willingness to pay and sharing intentions (Studies 3 & 4). We argue that perceptions of warmth shape consumers' aversion to AI mental health providers: AI (vs. human) providers were perceived as less warm, which harmed their perceived competence, resulting in higher aversion.

In Study 1a, we asked college students to fill out a survey about stress-related resources. After reflecting on an instance in which they experienced academic stress, participants were offered three types of advisers to choose from to help them improve how they deal with stress. Descriptions for all three advisers were based on real-world descriptions. We scrapped the top 500 Google search results for mental health care providers, and two independent coders catalogued the providers into 5 categories. Three types of advisors—clinical, integrative, and virtual/AI—together form more than 80% of the market. We thus focus on these three types of mental health providers. The first was a clinical adviser with an MD from the university, who has published peer-reviewed articles on stress. The second was an integrative adviser, a life coach who has triumphed over dealing with stress. The third was a virtual adviser powered by artificial intelligence; the order of these advisers was counterbalanced. We find that only 5.4% of the participants chose the AI adviser (χ^2 = *59.14, p-value<0.001*), compared to 64.9% choosing the integrative adviser and 29.7% choosing the clinical adviser.

We replicate this finding in a variety of different stress domains (Studies 1b-e). We consistently find that the AI adviser is dominated by the two other advisers. The percentage of participants choosing an AI adviser ranges from 5.41% (academic/relationship stress) to 20.99% (financial stress).

In Study 2, we investigate consumers' choice of mental health adviser in a high-stigma context: porn addiction, which has harmful consequences such as cognitive decline and isolation (Alarcón et al. 2019). We recruited men who watch at least 5 hours of porn per week (2 SD higher than average). We provided these participants with an article that described the negative consequences of porn overuse and asked them to choose which of the three advisers they would consider discussing their porn consumption habits with. Similar to the results above, we find that the AI adviser is dominated by the two other choices (χ^2 = 12.6, *p-value*<0.001). Interestingly, the percentage of participants willing to choose the AI adviser was higher (23%) than in the stress-related domains we tested.

In Study 3, participants were told about a new service in which customers share symptoms with a (virtual vs. human) doctor who diagnose if they have common mental health disorders, such as depression and anxiety. Participants then reported their willingness to pay for the service. We found that participants in the AI condition reported a lower willingness to pay compared to the human condition (= -2.8, $t(402) = -2.8$, $p = .005$). Importantly, we found support for our proposed serial mediation model, such that an AI (vs. human) provider was associated with lower perceived warmth, which lowered the perceived competence of the provider, thereby leading to lower willingness to pay for the service (indirect effect = −.91, SE = .27, 95% CI [−1.44, −.39]). As a robustness check, we included uniqueness neglect as a parallel mediator in the model and found that, while the indirect effect through warmth and competence remained significant (indirect effect = −.89, SE = .26, 95% CI [−1.40, −.38]), the indirect effect through uniqueness neglect was not significant (indirect effect = −.012, SE = .048, 95% CI [−.11, .082]), suggesting that uniqueness neglect did not affect consumers' willingness to pay for mental health providers in this context.

In Study 4, participants were first asked to reflect on a recent time they experienced financial stress. Participants were then presented with a [virtual] mental health coach [powered by artificial intelligence] that could help them deal with stressful situations. To ensure that consumers' AI aversion generalized across other dependent variables of importance to marketers, we included a measure of participants' intent to share information about the provider with others. Consistent with Study 3, we found that participants in the AI (vs. human) condition reported a lower willingness to pay ($\beta = -2.56$, $t(498) = -3.23$, $p = .001$) and a lower interest in sharing information about this provider with others ($\beta = -.39$, $t(498) = -2.05$, $p = .041$). Importantly, perceptions of warmth again mediated the effect of provider on perceived competence, which in turn influenced participants' willingness to pay (indirect effect = −.83, SE = .30, 95% CI [−1.42, −.23]) and sharing intentions (indirect effect = −.41, SE = .07, 95% CI [−.55, −.28]).

Overall, we found that consumers are reluctant to use AI mental health advisers, and that perceptions of warmth played an important role in shaping consumers' preference for a mental health provider. AI (vs. human) providers were judged as less warm, which harmed their perceived competence, and resulted in higher aversion. We replicated this finding across different types of mental health domains, such as stress and addictive behaviors (e.g. porn addiction), across experimental paradigms (within-subject and between-subject), and across dependent variables (choice, willingness to pay, and sharing intentions).

Experts Outperform Technology in Creative Markets

EXTENDED ABSTRACT

Recent decades have seen an increase in machines and other technological advancements coopting jobs from humans (Autor 2015; Brynjolfsson and McAfee 2014; Hosanagar 2019). Notably, in many cases algorithms or technology based on training from initial human inputs can outperform humans (see Dawes et al. 1989; Dietvorst et al. 2015, 2016). Once thought unlikely, the machines have become dominant in games of human skill such as Jeopardy, chess, and Go (Hosanagar 2019).

Recently, there has been discussion of technology based on artificial intelligence supplanting workers (Lohr 2018; Peiser 2019). While this replacement is argued to occur more for lower-skilled jobs (Autor 2015; Mokyr et al. 2015) and is expected more for jobs involving thinking (Waytz and Norton 2014), there have also been market entrants from artificial intelligence in domains that are typically considered to require the human spark of creativity (Rand 1968). For example, for logos design, sites such as logomaster.ai and Wix offer affordable artificial intelligence builders for logos. Like with human experts, these artificial intelligence makers present people with multiple initial design concepts from which they can choose, and they allow for iteration and revision on the initial concept they like most (Goodwin 2009).

Does artificial intelligence outperform human experts in creative markets, such as in logo design? To explore this question, we run two logo elicitation studies (Studies 1 and 2) in which subjects worked with artificial intelligence and/or human expert designers, and then run seven follow-up studies (Studies 1a-1d and 2a-2c), using both ratings and incentivized choice, in which an external set of participants evaluate the quality of the logos.

Importantly, we align the logo design process in the two logo elicitation studies. In each, participants start with a creative brief, have a designer selected for them or by them, see at least two initial concepts for a logo, and then go through a revision process before arriving at a final logo. In Study 1, we impose more experimental control by providing a fixed set of creative briefs for three (fictitious) companies (the data science company Empirical, the fashion company Forward, and a restaurant Hyperion) and assigning participants to work with either human experts or artificial intelligence logo makers, both of which provide initial concepts and a revision process. In Study 2, with a separate set of participants we allow for more natural design by having participants write creative briefs for a company (ONAK, an origami canoe company), and then let them work both with one human expert designer of their choice and one artificial intelligence logo maker. Study 1 yielded 17 human expert logos and 30 artificial intelligence logos, while Study 2 yielded 62 logos (31 human expert, 31 artificial intelligence). In Study 1, the participants were MBA students enrolled in a new product development course; in Study 2, the participants were managers enrolled in an evening MBA program.

In Study 1a, we first seek overall evaluations on the logos from Study 1 using external raters. Amazon Mechanical Turkers evaluated all of the logos from one of the three companies (Empirical, Hyperion, or Forward) from Study 1 on a seven-point scale (1 = Very Bad, 7 = Very Good). Importantly, this evaluation was blind to whether the logos were produced by artificial intelligence or human experts (see Dietvorst et al. 2015). Across all three companies, relative to artificial intelligence logos, human expert logos were evaluated more favorably ($F(1, 287) = 222.78$, $p < .001$). In Study 1b, human expert designers with at least fourteen years industry experience evaluated all logos and showed directionally consistent results.

In Study 1c, we replicate Studies 1a and 1b but with additional ratings designed to determine what dimensions the human expert logos exceeded artificial intelligence logos on. Undergraduates in a design course completed the procedure from Study 1a and also rated to what extent the logos conveyed the company's industry, were aesthetically pleasing, and were unique. In addition to replicating Study 1a on overall evaluations ($F(1, 144) = 131.92$, $p < .001$), human expert logos were also evaluated, relative to artificial intelligence, to convey the company's industry better ($F(1, 144) = 290.20$, $p < .001$), to be more aesthetically pleasing ($F(1, 144) = 51.85$, $p < .001$), and to be more unique ($F(1, 144) = 57.40$, $p < .001$).

In Study 1d, we conceptually replicated the previous studies with incentivized choice with a set of west coast laboratory participants. That is, participants were informed that those participants who chose the logo that was selected most often would be eligible for a $50 Amazon Gift card as a prize. Consistent with the previous studies, participants chose logos produced by human experts more often than would be expected by chance ($z = 8.49$, $p < .001$).

Study 2a replicates Study 1a on the 62 logos produced for ONAK using Amazon Mechanical Turkers. Again, compared with the artificial intelligence logos, the human expert logos were evaluated more favorably ($F(1, 394) = 114.47$, $p < .001$). Further, Study 2b replicated Study 1c on overall evaluations (i.e., human expert logos were judged to be better than artificial intelligence logos) with a set of west coast laboratory participants ($F(1, 423) = 170.40$, $p < .001$).

Finally, Study 2c attempts to replicate Study 1d on choice with a sample of west coast laboratory participants. When choosing among logos for ONAK, participants were only marginally more likely to select logos produced by human experts ($z = 1.95$, $p = .051$).

Overall, using two types of logo elicitation methods, we find that logos from human experts outperform those from artificial intelligence. This advantage may result from human expert logos being superior on aesthetics and clarity. However, subsequent coding and analyses revealed that those managers in Study 2 who invested more effort into their creative briefs (based on an expert designer coding the strength of each brief) had more favorably evaluated artificial intelligence logos, which might mean that the advantage to human experts may disappear with more experience or clarity of what is desired from the design process.

REFERENCES

Alarcón, Rubén de, Javier de la Iglesia, Nerea Casado, and Angel Montejo (2019), "Online Porn Addiction: What We Know and What We Don't—A Systematic Review." *Journal of Clinical Medicine* 8 (1), 91.

Alonso, Jordi, Zhaorui Liu, Sara Evans-Lacko, Ekaterina Sadikova, Nancy Sampson, Somnath Chatterji, Jibril Abdulmalik, et al. (2018), "Treatment Gap for Anxiety Disorders Is Global: Results of the World Mental Health Surveys in 21 Countries." *Depression and Anxiety* 35 (3), 195–208.

Arkes, Hal R., Robyn M. Dawes, and Caryn Christensen (1986), "Factors influencing the use of a decision rule in a probabilistic task," *Organizational behavior and human decision processes,* 37 (1): 93-110.

Autor, David H. (2015), "Why Are There Still So Many Jobs? The History and Future of Workplace Automation," *Journal of Economic Perspectives*, 29 (3), 3-30.

Blanco, Carlos, Mayumi Okuda, Crystal Wright, Deborah S. Hasin, Bridget F. Grant, Shang-Min Liu, and Mark Olfson (2008), "Mental Health of College Students and Their Non–College-Attending Peers: Results from the National Epidemiologic Study on Alcohol and Related Conditions," *Archives of General Psychiatry* 65 (12), 1429.

Brynjolfsson, Erik, and Andrew McAfee (2014), *The Second Machine Age: Work Progress, and Prosperity in a Time of Brilliant Technologies*.

Castelo, Noah, Maarten W. Bos, and Don R. Lehmann (2019), "Task-Dependent Algorithm Aversion," *Journal of Marketing Research*, 56 (5), 809−25.

Dawes, Robyn M., David Faust, and Paul E. Meehl (1989), "Clinical Versus Actuarial Judgment," *Science*, 243 (4899), 1668-74.

Dietvorst, Berkeley J., Joseph P. Simmons, and Cade Massey (2015), "Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err," *Journal of Experimental Psychology: General*, 144 (1), 114−26.

Dietvorst, Berkeley J., Joseph P. Simmons, and Cade Massey (2016), "Overcoming Algorithm Aversion: People Will Use Imperfect Algorithms If They Can (Even Slightly) Modify Them," *Management Science*, 64 (3), 1155-70.

Ferrari, Francesco, Maria Paola Paladino, and Jolanda Jetten (2016), "Blurring Human–Machine Distinctions: Anthropomorphic Appearance in Social Robots as a Threat to Human Distinctiveness," *International Journal of Social Robotics* 8 (2), 287–302.

Fox, Craig R., and Gülden Ülkümen (2011), "Distinguishing two dimensions of uncertainty," *Perspectives on thinking, judging, and decision making*, 21-35.

Goodwin, Kim (2009), *Designing for the digital age: How to create human-centered products and services*. John Wiley & Sons.

Grove, William M., David H. Zald, Boyd S. Lebow, Beth E. Snitz, and Chad Nelson (2000), "Clinical Versus Mechanical Prediction: A Meta-analysis," *Psychological Assessment*, 12 (1), 19-30.

Halperin, Keith, C. R. Snyder, Randee J. Shenkel, and B. Kent Houston (1976), "Effects of Source Status and Message Favorability on Acceptance of Personality Feedback," *Journal of Applied Psychology*, 61 (1), 85−88.

Hoehl, Stefanie, Stefanie Keupp, Hanna Schleihauf, Nicola McGuigan, David Buttelmann, and Andrew Whiten (2019), "'Over-imitation': A review and appraisal of a decade of research," *Developmental Review*, 51, 90-108.

Hosanagar, Kartik (2019), *A Human's Guide to Machine Intelligence: How Algorithms Are Shaping Our Lives and How We Can Stay in Control*.

Hunt, Justin, and Daniel Eisenberg (2010), "Mental Health Problems and Help-Seeking Behavior Among College Students," *Journal of Adolescent Health*, 46 (1), 3–10.

Leung, Eugina, Gabriele Paolacci, and Stefano Puntoni (2018), "Man versus machine: Resisting automation in identity-based consumer behavior," *Journal of Marketing Research*, 55 (6), 818-31.

Logg, Jennifer M., Julia A. Minson, and Don A. Moore (2019), "Algorithm Appreciation: People Prefer Algorithmic to Human Judgment," *Organizational Behavior and Human Decision Processes*, 151, 90−103.

Longoni, Chiara, Andrea Bonezzi, and Carey Morewedge (2019), "Resistance to Medical Artificial Intelligence," *Journal of Consumer Research,* 46 (4), 629–50.

McGuigan, Nicola, Jenny Makinson, and Andrew Whiten (2011), "From over-imitation to super-copying: Adults imitate causally irrelevant aspects of tool use with higher fidelity than young children," *British Journal of Psychology*, 102 (1), 1-18.

McQueen, Amy, and William M. P. Klein (2006), "Experimental Manipulations of Self-Affirmation: A Systematic Review," *Self and Identity* 5 (4), 289–354.

Meltzoff, Andrew N., and M. Keith Moore (1977), "Imitation of facial and manual gestures by human neonates," *Science*, 198 (4312), 75-78.

Mokyr, Joel, Chris Vickers, and Nicolas L. Ziebarth (2015), "The History of Technological Anxiety and the Future of Economic Growth: Is This Time Different?," *Journal of Economic Perspectives*, 29 (3), 31-50.

Rand, Paul (1968), The Designer's Problem. *A Designer's Art*. Yale University Press.

Sontag-Padilla, Lisa, Michael S. Dunbar, Feifei Ye, Courtney Kase, Rebecca Fein, Sara Abelson, Rachana Seelam, and Bradley D. Stein (2018), "Strengthening College Students' Mental Health Knowledge, Awareness, and Helping Behaviors: The Impact of Active Minds, a Peer Mental Health Organization," *Journal of the American Academy of Child & Adolescent Psychiatry* 57 (7), 7500–7.

Steele, Claude M. (1988), "The Psychology of Self-Affirmation: Sustaining the Integrity of the Self." In *Advances in Experimental Social Psychology*, edited by Leonard Berkowitz, 21, 261–302.

Waytz, Adam, and Michael I. Norton (2014), "Botsourcing and Outsourcing: Robot, British, Chinese, and German Workers Are for Thinking—Not Feeling—Jobs," *Emotion*, 14 (2), 434−44.

Yeomans, Michael, Anuj Shah, Sendhil Mullainathan, and Jon Kleinberg (2019), "Making Sense of Recommendations," *Journal of Behavioral Decision Making,* 32, 403-414.