Article

AMERICAN MARKETING ASSOCIATION

Journal of Marketing Research
1-17

DOI: 10.1177/0022243719851788
journals.sagepub.com/home/mrj
SAGE

Task-Dependent Algorithm Aversion

Noah Castelo, Maarten W. Bos, and Donald R. Lehmann

Abstract

Research suggests that consumers are averse to relying on algorithms to perform tasks that are typically done by humans, despite the fact that algorithms often perform better. The authors explore when and why this is true in a wide variety of domains. They find that algorithms are trusted and relied on less for tasks that seem subjective (vs. objective) in nature. However, they show that perceived task objectivity is malleable and that increasing a task's perceived objectivity increases trust in and use of algorithms for that task. Consumers mistakenly believe that algorithms lack the abilities required to perform subjective tasks. Increasing algorithms' perceived affective human-likeness is therefore effective at increasing the use of algorithms for subjective tasks. These findings are supported by the results of four online lab studies with over 1,400 participants and two online field studies with over 56,000 participants. The results provide insights into when and why consumers are likely to use algorithms and how marketers can increase their use when they outperform humans.

Keywords

algorithms, new products, technology adoption

Online supplement: https://doi.org/10.1177/0022243719851788

Algorithms—a set of steps that a computer can follow to perform a task—increasingly outperform humans at many tasks. Pioneering literature from the 1950s demonstrated that very simple algorithms such as linear regression could outperform expert humans on tasks such as diagnosing medical and psychological illnesses (Dawes, Faust, and Meehl 1989; Grove et al. 2000; Meehl 1954). Since then, rapid progress in the field of artificial intelligence has endowed algorithms with the abilities to understand and produce natural language, learn from experience, and even understand and mimic human emotions. Today, algorithms can outperform even expert humans at an increasingly comprehensive list of tasks, from diagnosing complex diseases (Simonite 2014) to driving cars and providing legal advice (Krasnianski 2015). Algorithms can also perform seemingly subjective tasks such as detecting emotion in facial expressions and tone of voice (Kodra et al. 2013). Algorithms thus offer enormous potential for improving outcomes for consumers and firms, including the automation of a large proportion of marketing decisions (Bucklin, Lehmann, and Little 1998). The rise of algorithms means that consumers are increasingly presented with a novel choice: should they rely more on humans or on algorithms? Research suggests that the default option in this choice is to rely on humans, even when doing so results in objectively worse outcomes.

Algorithm Aversion

Table 1 summarizes the primary results of empirical studies that have examined trust in and use of algorithms compared with humans, ordered by publication date. The dominant theme is that consumers prefer humans over algorithms (for an exception, see Logg, Minson, and Moore [2019]). For example, after seeing an algorithm err, people prefer to rely on humans for forecasting student performance, even when doing so results in suboptimal forecasts (Dietvorst, Simmons, and Massey 2014). People also trust medical recommendations less when they come from an algorithm than from a human doctor (Promberger and Baron 2006; Longoni, Bonezzi, Morewedge 2019). Scholars have argued that consumers are averse to relying on automated medical care because they believe that such care will not fully take into account their unique individual circumstances (Longoni, Bonezzi, and Morewedge 2019). In the same domain, Shaffer et al. (2013) found that participants rated physicians who made an unaided diagnosis significantly more positively than a physician who used an algorithm to assist

Noah Castelo is an Assistant Professor of Marketing, University of Alberta, Canada (email: ncastelo@ualberta.ca). Maarten W. Bos is Senior Research Scientist, Snap Inc., USA (email: maarten.w.bos@gmail.com). Donald R. Lehmann is George E. Warren Professor of Business, Columbia University, USA (email: drl2@columbia.edu).

Table 1. Summaries of Relevant Research on Perceptions and Use of Algorithms.

Research	Sample	Main Independent Variable	Main Dependent Variable	Main Findings
Promberger and Baron (2006)	Online U.S. panel; two-study N = 166	Medical recommendation from "physician" vs. "computer program"	Acceptance of recommendation and trust in recommender	Acceptance and trust are higher for humans vs. computer programs (algorithms).
Önkal et al. (2009)	Turkish economics students; two-study N = 130	Financial forecasts from "human expert" vs. "statistical forecasting method"	Weight of advice (how much participants adjusted own forecast after receiving advice)	Human advice was given more weight than algorithmic advice.
Diab et al. (2011)	Online U.S. and non-U.S. panel; one-study N = 462	Employee selection via "thorough discussion" vs. "mathematical formula"	Perceived usefulness, fairness, and flexibility of selection method.	Thorough discussions (humans) were viewed as more useful, fair, and flexible than formulae (algorithms).
Eastwood, Snook, and Luther (2012)	Canadian undergraduate students and shopping mall attendees; three-study N = 235	Financial and medical advice based on "intuition and personal experience" vs. "a statistical formula"	Preference, accuracy, fairness, ethicality of advice methods	Intuition and experience (humans) were preferred and seen as more accurate, ethical, and fair than formulae (algorithms).
Shaffer et al. (2013)	U.S. undergraduate students; three-study N = 732	Doctor who makes an unaided diagnosis or solicits aid from either computer program or from human colleague	Doctor's perceived diagnosis ability, professionalism, thoroughness	Soliciting aid from a computer but not from human decreases perceived ability, professionalism, and thoroughness.
Dietvorst, Simmons, and Massey (2014)	U.S. students and MTurk workers; five-study N = 2,367	Observing vs. not observing an algorithm perform (and err) at forecasting tasks	Choice to rely on algorithm vs. oneself or algorithm vs. other participants when making incentivized forecasts	Participants exhibited reduced reliance on the algorithm after seeing it err, even when it outperforms humans.
Dietvorst, Simmons, and Massey (2016)	U.S. students and MTurk workers; three-study N = 1,922	Being able vs. unable to modify an algorithm's forecasts	Choice to rely on algorithm vs. oneself when making incentivized forecasts	Participants showed increased reliance on the algorithm when its output was modifiable.
Logg, Minson, and Moore (2019)	Students, MTurk workers, academic researchers, and national security professionals; seven-study N = 1,464	Various forecasts from either "another person" or an "algorithm"	Weight of advice (how much participants adjusted their own forecast after receiving advice)	Nonexperts relied more on advice from algorithms than from other people. Experts relied less on algorithms than on themselves.
Yeomans et al. (2019)	MTurk workers and museum visitors; six-study N = 3,647	Receiving joke recommendations from either a human or an algorithm	Preference for source of recommendations and likelihood of using algorithm to make recommendations for others	People prefer to receive jokes from humans and to recommend jokes themselves (vs. using an algorithm).
Longoni, Bonezzi, and Morewedge (2019)	MTurk workers and U.S. students; nine-study N = 2,596	Receiving medical treatment from a human or from artificial intelligence	Preference for using, reservation prices, and actual use of treatment provider	People prefer to receive medical treatment from humans because they think artificial intelligence neglects their unique circumstances.

with the diagnosis, but no differently than a physician who consulted a colleague to assist with the diagnosis. Participants who had a greater internal locus of control had more negative evaluations of algorithmic diagnoses.

Önkal et al. (2009) found that participants relied less on advice from an algorithm than from a human when forecasting stock prices. Like Shaffer et al. (2013), they argued that participants find it easier to shift responsibility or blame to other humans. They also noted that humans, unlike most algorithms, can provide explanations for their decisions (Armstrong 1980), are perceived to have high confidence (Sniezek and Buckley 1995), have a reputation to maintain (Eisenhardt 1989), and have information about future events (Blattberg and Hoch 1990). In contrast, algorithms are thought to possess none of these qualities.

In the domain of employee selection and hiring decisions, Diab et al. (2011) found that participants thought that human interviews were more useful, professional, fair, personal, flexible, and precise than algorithms. In the domain of student performance forecasting, Dietvorst, Simmons, and Massey (2014) found that participants preferred to make their own forecasts rather than relying on an algorithm after seeing the algorithm err, and that while algorithms were seen as better than humans at avoiding obvious mistakes, appropriately weighing attributes, and consistently weighing information, humans were seen as better than algorithms at learning from mistakes, getting better with practice, finding "diamonds in the rough," and detecting exceptions to the rule. Dietvorst, Simmons, and Massey (2016) also found that allowing participants to slightly modify the output of an algorithm makes them feel more satisfied with the forecasting process, more tolerant of errors, more likely to believe that the algorithm is superior, and more likely to choose an algorithm to make subsequent forecasts.

Finally, Yeomans et al. (2019) found that participants relied less on an algorithm than on humans for the task of predicting joke funniness, while Logg, Minson, and Moore (2019) found that participants relied *more* on algorithms than on humans for numerical tasks with an objectively correct answer. These findings suggest that whether consumers are averse to algorithms or appreciate them varies significantly depending on the type of task for which the algorithm is used and how that task is perceived.

There are two notable gaps in this literature. First, there has not been a systematic exploration of how and why consumers' willingness to use algorithms varies across the many different types of tasks for which algorithms can be used. Second, there has been little exploration and validation of practical interventions that marketers can use to increase consumers' willingness to rely on algorithms instead of humans, especially in cases where the algorithm outperforms expert humans.

We therefore make two primary contributions in this article by addressing these gaps. First, we examine how willingness to trust and use algorithms varies by the characteristics of the task. We identify a robust effect that algorithms are trusted and used less for tasks that are perceived as subjective in nature, and we show that this effect occurs primarily because of a belief that algorithms are ineffective at subjective tasks. Trust involves both cognitive and affective dimensions (Johnson and Grayson 2005). Cognitive trust involves confidence in another agent's performance or reliability, whereas affective trust is based on one's feelings and can be independent from performance. In the context of our research, we suggest that consumers' overall trust in algorithms is affected by both performance-based, cognitive beliefs about the algorithm's performance and by feelings stemming from consumers' comfort with the use of algorithms for tasks typically done by humans, which can be independent of performance-related beliefs. We therefore explore both consumers' beliefs about algorithmic performance as well as their comfort with the use of algorithms as mechanisms of our main effect.

Second, we explore approaches for making algorithms more attractive to potential users when use is low despite algorithmic superiority over expert humans. We show that the perceived objectivity of a task is malleable and that reframing tasks as being relatively objective increases trust in and willingness to rely on algorithms. Furthermore, we show that the belief in algorithm ineffectiveness for subjective tasks is also malleable. Specifically, increasing the affective human-likeness of algorithms by providing real examples of algorithms with affective abilities, such as understanding emotion and creating art, can make algorithms seem more effective at performing subjective tasks, which ultimately increases reliance on algorithms for such tasks.

Hypothesis Development

There are many potentially relevant dimensions along which tasks vary that can influence consumers' use of algorithms. For example, consumers are already familiar with the use of algorithms for certain tasks such as recommending movies on Netflix or filing taxes on TurboTax. Familiarity with algorithms for a given task is likely to increase trust in and willingness to rely on algorithms for that task. Similarly, some tasks are more consequential than others, in the sense that performing the task poorly will have more serious consequences. Consumers may be less willing to trust and rely on algorithms for more consequential tasks because doing so poses greater risks. Our primary focus, however, is the perceived *objectivity* of the task. We define an objective task as one that involves facts that are quantifiable and measurable, compared with subjective tasks, which we define as being open to interpretation and based on personal opinion or intuition. Research has shown that lay people see objective tasks as requiring logical, rule-based analysis and subjective tasks as requiring intuition and "gut instincts" (Inbar, Cone, and Gilovich 2010). Importantly, the objectivity of a task is not completely inherent in a given task but may be a malleable perception with heterogeneity both among different people and over time. We therefore exploit this heterogeneity to develop manipulations of perceived task objectivity.

The impact of task objectivity on consumers' trust and use of algorithms likely depends on the kinds of abilities consumers typically believe that algorithms possess. One relevant conceptual distinction is between cognitive and emotional abilities. For example, research on dehumanization has shown that people perceive two categories of human abilities. First are "human uniqueness abilities," which distinguish humans from other animals but can be shared with machines. These tend to be cognitive in nature (such as logic and rationality). Second are "human nature abilities," which may be shared with other animals but not with machines. These tend to be affective or emotional in nature (such as warmth and intuition; Haslam 2006; Loughnan and Haslam 2007). Importantly, research has shown that machines such as robots are viewed as lacking human nature abilities (which are emotional) but not human uniqueness abilities (which are cognitive; Haslam et al. 2008).

Research on mind perception has focused on two similar dimensions: agency, the ability to engage in intentional planning and action, and experience, the ability to subjectively experience emotions and sensations. Mirroring research on humanness, machines (such as robots) are seen as having some degree of agency but no experience (Gray, Gray, and Wegner 2007), and endowing robots with experience creates more negative reactions than endowing robots with agency (Gray and Wegner 2012). These streams of research demonstrate that consumers perceive human abilities as either cognitive or emotional and are willing to grant machines more cognitive than emotional abilities. Integrating these streams of research with the distinction between objective tasks, which are typically associated with more "cognitive" abilities, and subjective tasks, which are typically associated with more "emotional" abilities (Inbar, Cone, and Gilovich 2010), suggests that consumers will believe that algorithms will be less effective at subjective tasks because they are believed to lack the affective or emotional abilities typically associated with these tasks. Beliefs about whether a technology is effective are fundamental determinants of whether that technology is ultimately adopted (Davis, Bagozzi, and Warshaw 1989; Rogers 1976). Therefore,

> **H_1:** Consumers trust and rely on algorithms less for subjective (vs. objective) tasks.

We measure trust in algorithms in several studies because research has shown that trust in a technology is an important determinant of the decision to use it (Komiak and Benbasat 2006; Li, Hess, and Valacich 2008; Pavlou and Gefen 2004). Trust is relevant in situations in which one party is somehow dependent on the actions of another party and this dependence involves risk (Chopra and Wallace 2003), which is the case when consumers use an algorithm to perform a task typically done by a human. Higher trust in an algorithm should therefore lead to greater willingness to use the algorithm. Thus, while the majority of our studies focus specifically on the actual or intended use of algorithms, we also measure trust in algorithms (both in terms of affective and cognitive trust) as an important factor that contributes to use.

It follows from H_1 that one way of increasing the use of algorithms for a given task is to increase the degree to which the task is perceived as objective. Most tasks can be perceived as more or less objective depending on how they are framed and which task components are emphasized. Specifically, a given task can be approached either by measuring and analyzing relevant quantitative variables or by using intuition or gut feelings. For example, evaluating and hiring a job candidate could be done by using standardized psychometric tests and measures or by conducting informal interviews and relying on one's gut feeling or intuition. Importantly, it is not always clear which of these approaches is superior for many important tasks, as psychologists continue to debate the relative merits of more "deliberate" versus more "intuitive" approaches to different tasks (Bear and Rand 2016; Dijksterhuis et al. 2006; Gigerenzer and Brighton 2009; Slovic et al. 2004). Furthermore, consumers also differ in terms of their preferences for and tendencies to rely on more analytical versus intuitive approaches to decision making (Greifeneder, Bless, and Pham 2011; Inbar, Cone, and Gilovich 2010). This uncertainty provides an opportunity to frame tasks typically viewed as subjective as being more objective. Therefore,

> **H_2:** Describing a task as benefiting from quantitative analysis (relative to intuition) increases perceived task objectivity and consumers' trust in and reliance on algorithms.

In addition to changing how the task is perceived, a second approach for increasing the use of algorithms involves changing how the algorithm itself is perceived. As mentioned earlier, consumers tend to believe that machines lack fundamentally human capabilities that are emotional or affective in nature (i.e., that they lack affective human-likeness) (Gray, Gray, and Wegner 2007; Haslam et al. 2008). However, this belief is increasingly inaccurate. Algorithms can already create paintings that sell for hundreds of thousands of dollars (Quackenbush 2018), write compelling poetry and music (Deahl 2018; Gibbs 2016), predict which songs will be hits (Herremans, Martens, and Sörensen 2014), and even accurately identify human emotion from facial expressions and tone of voice and respond accordingly (Goasduff 2017; Kodra et al. 2013; McDuff et al. 2013). Algorithms are therefore increasingly capable of performing the kinds of tasks typically associated with subjectivity and emotion. Note that even though algorithms may accomplish these tasks using very different means than humans do (i.e., using predetermined computer programming rather than intuitions or gut feelings) we show that the fact that algorithms can accomplish such tasks at all makes them seem more human-like.

We expect that increasing algorithms' perceived human-likeness in this way will moderate the effect of task objectivity. This moderation, however, could plausibly either increase or decrease the effect of objectivity. On one hand, increasing affective human-likeness is likely to increase the perceived effectiveness of algorithms for subjective tasks, because consumers believe that subjective tasks require affective abilities that algorithms are typically thought to lack. Making algorithms seem more human-like could therefore decrease or eliminate the main effect of task objectivity. This result would indicate that cognitive trust (i.e., beliefs about algorithm effectiveness) is more important than affective trust (i.e., feelings that are independent from beliefs about effectiveness) in shaping consumers' use of algorithms.

On the other hand, increasing affective human-likeness may also produce discomfort with the use of algorithms by challenging the belief that humans are distinct from machines. Indeed, social identity theory posits that people derive meaning and satisfaction from membership in distinct groups, and that when an out-group threatens the sense of distinctiveness of their in-group, they react negatively toward the threatening out-group (Tajfel 1982). In other words, people like to believe that their in-group is unique, and when an out-group begins to challenge that perceived uniqueness, the out-group will be evaluated negatively (Brewer 1991; Ferrari, Paladino, and Jetten 2016). Increasing the affective human-likeness of algorithms could therefore represent an intergroup challenge in the sense of algorithms as an out-group challenging the distinctiveness of humans (as an in-group) from machines. According to social identity theory, this challenge to in-group distinctiveness could in turn lead to negative evaluations of the challenging out-group (algorithms in this case), ultimately decreasing the use of algorithms, including for subjective tasks. This result would suggest that affective trust is more important than cognitive trust in determining consumers' use of algorithms.

Measuring consumers' reliance on algorithms that vary in affective human-likeness can thus help tease apart these two competing hypotheses. Whether increasing affective human-likeness reduces or exacerbates the main effect of task objectivity may ultimately depend on whether cognitive factors (i.e., beliefs about effectiveness) or affective factors (i.e., feelings of discomfort potentially stemming from intergroup challenges) have a stronger impact on consumers' use of algorithms. In other words, endowing algorithms with affective abilities may negatively affect consumers' affective reactions to the use of algorithms while also positively affecting consumers' cognitive reactions, and both cognitive and affective reactions should influence consumers' willingness to use algorithms. In place of a third hypothesis, we thus posit a final research question:

> Will increasing algorithms' perceived affective human-likeness decrease or increase the effect of perceived task objectivity on consumers' use of algorithms?

We test these hypotheses and answer this research question using a variety of field and lab studies and several different dependent measures. To increase generalizability and demonstrate the robustness of our effects, we operationalize our dependent variable in multiple ways, including self-reported trust in and preference for algorithms relative to humans, clicks on online advertisements for algorithm- and human-based services, and actual reliance on algorithms in the context of an incentivized task.

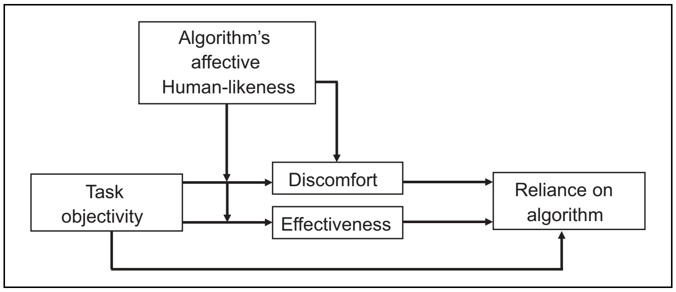

Figure 1. Conceptual model.
Notes: Task objectivity decreases discomfort with using algorithms for the task, increases algorithms' perceived effectiveness, and increases reliance on algorithms, but these effects are reduced when the algorithm has high affective human-likeness.

To summarize, existing explanations of algorithm aversion suggest that it is largely driven by a perception that algorithms lack various human abilities. We propose that it is specifically affective abilities that algorithms are seen as lacking and that this belief decreases willingness to rely on algorithms for tasks that seem subjective. Consequently, emphasizing that a given task benefits from a more quantitative approach can increase the perceived objectivity of that task, ultimately increasing trust and use of algorithms. Finally, increasing algorithms' perceived affective human-likeness may either reduce or increase the effect of perceived task objectivity on consumers' use of algorithms, depending on whether algorithmic effectiveness or discomfort with human-like algorithms have a stronger impact on use.

We explore these questions in 6 studies. Study 1 shows that trust in algorithms varies substantially depending on the task, and that trust is lower for more subjective tasks. Study 2 replicates this effect in an applied field study. Study 3 shows that providing evidence that algorithms are superior to humans for a specific task is less effective at increasing consumers' preference for using algorithms when the task is relatively subjective. Study 4 shows that reframing a subjective task as being more objective increases trust in algorithms. Study 5 replicates the task-framing effect in a field study. Finally, Study 6 shows that increasing algorithms' perceived affective human-likeness by providing examples of algorithms with affective abilities *increases* the use of algorithms for subjective tasks, thus eliminating the effect of task objectivity. Put simply: consumers have strong preconceptions about what algorithms are good at, so two ways to increase reliance on algorithms are to change the way (1) the task and (2) the algorithm's abilities are perceived.

Figure 1 depicts the conceptual model that we test in this research. The main effect that we demonstrate is that perceived task objectivity reduces consumers' trust in and willingness to rely on algorithms (Studies 1–6). We provide evidence that this effect is explained partially by the perceived effectiveness of algorithms for subjective tasks (Studies 4 and 6) and partially by consumers' discomfort using algorithms for subjective tasks

(Study 6). Furthermore, we study the effects of algorithms' affective human-likeness, finding both direct effects on discomfort and interactions between human-likeness and task objectivity in shaping discomfort, perceived effectiveness of the algorithm, and reliance on the algorithm (Study 6). Specifically, we find that the effects of task objectivity on perceived algorithm effectiveness, consumers' discomfort, and reliance on algorithms are each reduced when the algorithm has high affective human-likeness. The primary results of these studies are summarized in Table 2.

Study 1

To gain an initial understanding of how trust in algorithms varies by task, in Study 1 we examined 26 tasks that vary along several dimensions. The primary goal of this study was to test H_1, that trust in algorithms would be lower for more subjective tasks. We also measured task consequentialness and how familiar participants were with the use of algorithms for each task, two other potentially relevant dimensions. Furthermore, we also measured trust in well-qualified humans for each task so that we could compare trust in algorithms to trust in humans for each task. This study was conducted in two parts, with one sample of participants rating the tasks along the dimensions of objectivity, consequentialness, and familiarity with the use of algorithms, and a second sample rating their trust in algorithms or in humans for each task.

Method

Participants and design. For part 1, we recruited 250 participants (M_{age} = 37 years, 41% female) from Amazon's Mechanical Turk (MTurk), who rated tasks along several dimensions. For part 2, we recruited 387 participants (M_{age} = 36 years, 45% female) from MTurk, who were randomly assigned to one of two conditions (trust in humans vs. trust in algorithms).

Procedure. For part 1, participants rated each of 26 tasks on how objective versus subjective it seemed, how consequential versus inconsequential it seemed, and how familiar they were with the use of algorithms for each task, using scales from 0 ("not at all") to 100 ("completely"). The tasks as well as the dimensions being rated were presented in random order. For part 2, participants indicated how much they would trust *either* an algorithm or a "very well qualified person" for each of 26 tasks also on a scale from 0 ("not at all") to 100 ("completely"). For example, for the task of diagnosing a disease, the person was described as a doctor. The tasks are shown in Table 3.

Results and Discussion

Averaged across tasks, trust in a qualified human was higher than trust in algorithms (M_{human} = 70.2 vs. $M_{algorithm}$ = 52.8; $t(385) = 5.75$, $p < .001$). However, trust in algorithms was higher than in humans for certain tasks (predicting stock market outcomes, predicting the weather, analyzing data, and giving directions, ts > 4.70, *p*s < .001; see Table 3).

To test the effects of the three dimensions of consequentialness, familiarity, and objectivity on trust in algorithms, we conducted a regression in which these three dimensions were simultaneously used to predict trust in algorithms. This revealed that trust in algorithms was lower for tasks that seemed more consequential ($\beta = -.56$, $p < .001$), higher for tasks for which consumers were more familiar with the use of algorithms ($\beta = .42$, $p < .001$), and most importantly higher for tasks that seemed more objective ($\beta = .46$, $p = .004$). The adjusted R^2 for this regression was .54. These results provide initial support for H_1, suggesting that trust is higher for more objective tasks. The next study will corroborate this finding in an applied field study.

Study 2

While our first study provided initial support for H_1, it is possible that participants' self-reported trust is not a reliable indicator of their actual behavior. To address this concern, in Study 2 we examined the role of task objectivity in a field study in which participants' behavior was directly observed.

Method

Participants and design. We created four advertisements, organized in a 2 (human vs. algorithmic advisor) $\times$ 2 (dating vs. financial advice) design, and displayed them on Facebook (for the exact stimuli used in this and all other studies, see Web Appendix A). The ads portrayed either a human or an algorithm providing either dating advice (rated as highly subjective in Study 1) or financial advice (rated as highly objective in Study 1). These four ads were seen by 41,592 unique Facebook users (40% female, mean age not observed) on their Facebook newsfeed (i.e. the stream of posts that users see on Facebook from their friends and advertisers).

Procedure. Participants who see ads on their Facebook newsfeed can click on those ads to learn more about them. Our ads were clicked on 604 times in total. Participants who clicked on our ads were taken to a page explaining that we were studying consumers' trust in algorithms. Our dependent variable was the click-through rate (CTR) of the ads, which is the number of clicks a given ad received divided by the number of times that ad was seen (Facebook shares this information with the creator of the advertising campaign). We expected participants to be more likely to click on an ad for dating advice when it was advertised as coming from a human versus an algorithmic advisor, but that CTRs would not differ between human and algorithmic financial advice, because the latter is a more objective task.

Results and Discussion

We conducted a logistic regression to estimate the effects of human versus algorithm (human = 1, algorithm = 0), finance versus dating ad (finance = 1, dating = 0), and their interaction on the CTR (click = 1, no click = 0). This revealed that the

Table 2. Summary of Main Findings from Six Studies and Within-Article Meta-Analysis.

Study 1: N = 250; MTurk; measured task objectivity, trust in humans, and trust in algorithms; between-subjects (0–100 scales)		
	Subjective Tasks	Objective Tasks
Trust in algorithm	47.3	56.8
Trust in human	73.3	68.7
Main finding	Perceived task objectivity is positively related to trust in algorithms ($\beta = .46$, $p < .001$)	

Study 2: N = 41,592; Facebook; manipulated task objectivity and source of advice (human vs. algorithm); measured ad CTRs		
	Subjective Task	Objective Task
Algorithm ad CTR	.6%	1.6%
Human ad CTR	2.1%	1.8%
Main finding	Higher CTR for human versus algorithm ad only for subjective task (interaction $\beta = -.99$, $p < .001$)	

Study 3: N = 201; MTurk; manipulated awareness of algorithm's performance; measured task objectivity and preference for algorithms relative to humans (0–100 scales)		
	Subjective Tasks	Objective Tasks
Preference for algorithm (known performance)	47.3	55.1
Preference for algorithm (unknown performance)	31.7	34.5
Main finding	Awareness of algorithm's performance increases preference for algorithm more when task is objective (interaction $\beta = -.20$, $p = .013$)	

Study 4: N = 201; Prolific Academic; manipulated perceived task objectivity; measured preference for algorithms relative to humans (0–100 scale)		
	Subjective Tasks	Objective Tasks
Preference for algorithm	33.9	44.5
Main finding	Reframing the same task as being more objective increases preference for algorithms ($p < .001$).	

Study 5: N = 13,621; Facebook; manipulated perceived task objectivity (0–100 scale); measured ad CTRs		
	Subjective Framing	Objective Framing
Algorithm ad CTR	.4%	.9%
Main finding	Reframing the same task as being more objective increases CTR for algorithm-based advice ads ($p = .038$).	

Study 6: N = 399; Prolific Academic; manipulated perceived task objectivity, manipulated algorithm's affective human-likeness; measured reliance on algorithm in an incentivized forecasting task (0–1 scale)		
	Subjective Task	Objective Task
Reliance on algorithm (low human-likeness)	.22	.39
Reliance on algorithm (high human-likeness)	.35	.40
Main finding	Effect of perceived objectivity eliminated when algorithm has high affective human-likeness (interaction $\beta = -.15$, $p = .060$)	

Six-study meta-analysis: N = 56,264; effects of task objectivity on algorithm appreciation		
	Cohen's d	95% CI
Study 1 (effect of task objectiveness on trust)	.32	[.02, .61]
Study 2 (effect of task objectiveness on CTR, algorithm condition)	.32 (1.80)	[.28, .36]
Study 3 (effect of task objectiveness on preferred use, no-performance-information condition)	.17	[.11, .24]
Study 4 (effect of task objectiveness on preferred use)	.56	[.27, .84]
Study 5 (effect of task objectiveness on CTR)	.44 (2.21)	[.39, .49]
Study 6 (effect of task objectiveness on use, low-human-likeness condition)	.52	[.24, .84]
Weighted average	.36	[.35, .37]

Notes: We obtained Cohen's d in Studies 2 and 5 by converting ORs (shown in parentheses).

Table 3. Trust in Algorithms Versus Qualified Humans (Study 1).

	Trust Human	Trust Algorithm	Human–Algorithm Gap	Task Objectiveness
Predicting joke funniness	65	30	**35**	27
Hiring and firing employees	72	34	**38**	49
Recommending a romantic partner	59	37	**22**	26
Writing news article	79	37	**42**	48
Predicting recidivism	54	42	**12**	45
Composing a song	81	43	**38**	30
Driving a truck	81	43	**38**	70
Recommending a gift	75	46	**29**	26
Predicting student performance	63	46	**17**	52
Piloting a plane	79	47	**32**	78
Driving a car	81	47	**34**	69
Recommending disease treatment	73	48	**25**	69
Diagnosing a disease	73	48	**25**	77
Predicting employee performance	61	50	**11**	51
Driving a subway	77	52	**25**	73
Predicting an election	51	54	−3	57
Recommending a marketing strategy	70	56	**14**	55
Recommending music	75	59	**16**	22
Recommending a movie	76	59	**17**	23
Buying stocks	62	60	2	56
Playing a piano	84	61	**23**	48
Predicting stocks	55	63	**−8**	58
Predicting weather	57	67	**−10**	68
Scheduling events	78	69	**9**	62
Analyzing data	69	80	**−11**	73
Giving directions	70	82	**−12**	75

Notes: Tasks are listed in increasing order of trust in algorithms. The human–algorithm gap for a task is statistically significant ($p < .001$) when the corresponding number is in boldface.

CTR was higher for human ads ($\beta = 1.14$, $p < .001$) and finance ads ($\beta = .59$, $p = .001$). We also found a significant interaction between these factors ($\beta = -.99$, $p < .001$). As predicted, the CTR for the dating advice ads was significantly higher for the human advisor (2.1%) than for the algorithm advisor (.6%, $\chi^2(1) = 29.10$, odds ratio [OR] = 3.14, $p < .001$). In contrast, for the financial advice ads, the CTR was only marginally significantly higher for the human advisor (1.8%) than for the algorithm advisor (1.6%, $\chi^2(1) = 3.26$, OR = 1.16, $p = .092$; see Figure 2). This range of CTRs is comparable to other recent studies using Facebook advertising campaigns (Matz et al. 2017). These results provide further support and external validity to the notion that trust in algorithms is low primarily for tasks that are seen as subjective.

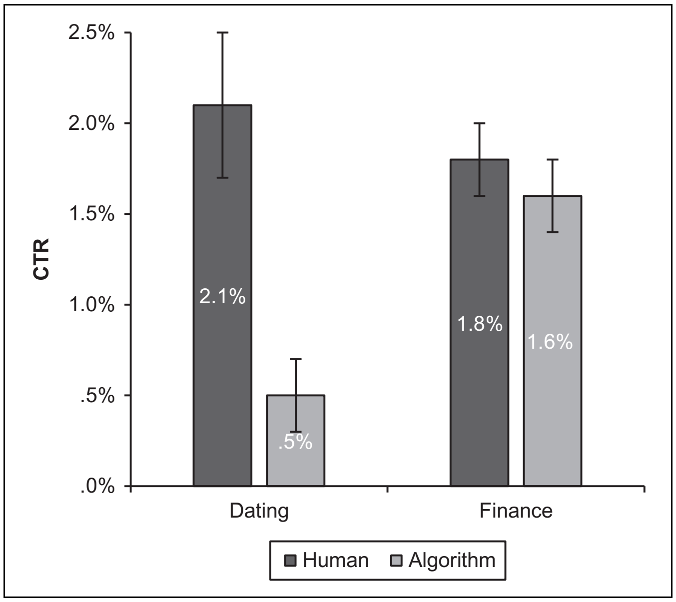

Figure 2. Algorithm aversion for a subjective (but not objective) task.
Notes: Error bars represent standard errors.

Study 3

Our first two studies have provided support for H_1, showing that consumers are less willing to trust and use algorithms for tasks that seem subjective. Nevertheless, algorithms are often highly effective at such tasks. Interventions that can increase trust and use of algorithms for such tasks could therefore be helpful to both consumers and firms. One of the most intuitive approaches for increasing consumers' willingness to use algorithms is to provide them with empirical evidence of the algorithms' superior performance relative to humans for the specific task in question. However, given the effect demonstrated in the previous studies, we anticipated that this evidence would be less effective at increasing willingness to use algorithms for tasks that seem subjective because consumers may be less likely to believe that algorithms can perform subjective tasks better than humans even when provided relevant evidence. Note that Study 3 does not manipulate the perceived human-likeness of the algorithm, which itself may have an impact on the perceived effectiveness of algorithms in general, but instead manipulates the perceived effectiveness of the algorithm for the specific task in question.

Method

Participants and design. Two hundred one MTurk users (M_{age} = 36 years, 49% female) reported their preference for using an algorithm relative to a qualified human for nine tasks (see Table 4). These nine tasks varied in terms of both trust in algorithms and perceived objectivity as measured in Study 1. Importantly, there is published research available documenting the superiority of algorithms over qualified humans for all nine tasks. Participants were assigned to one of two performance conditions: known performance or unknown performance. In

Table 4. Trust in Humans Versus Algorithms.

	Preference for Algorithm Relative to Human			
	Without Performance Data	With Performance Data	Δ	Task Objectivity
Predict student performance	**39**	**57**	18	52
Predict employee performance	**27**	52	25	51
Recommend disease treatment	**31**	**59**	28	69
Predict recidivism	**24**	52	28	45
Drive a car	**26**	53	27	69
Recommend a movie	**33**	52	19	23
Diagnose a disease	**23**	46	23	77
Predict personality traits	**35**	**40**	5	41
Predict joke enjoyment	**19**	**35**	16	27
Average	**29**	50	21	50

Notes: Higher numbers indicate greater trust in algorithms relative to humans. Boldfaced means are significantly different from the scale midpoint (50, labeled as "trust both equally"). The next-to-last column is the change in preference between condition (significant for all tasks), and the last column is the tasks' rated objectivity, taken from Study 1.

the known performance condition, we told participants that the algorithm outperformed the human and described the results of a real study that had demonstrated the algorithm's superior performance. In the unknown performance condition, this information was omitted, and the performance of the algorithm was not mentioned.

Procedure. Participants read about and rated each of the nine tasks individually. In the known performance condition, participants read about a published academic study for each task which demonstrated that an algorithm could outperform qualified humans. We provided links to each study and reported how much better the algorithm performed compared with the humans in the study. In the unknown performance condition, this information was not provided, and participants simply reported their preference without learning how the performance of algorithms compared with the performance of humans.

Participants reported whether they would rather use an algorithm or a qualified human for each task, with responses entered on a 0 to 100 scale, with 0 labeled as the relevant qualified human, 50 labeled as no preference, and 100 labeled as algorithm. For example, the relevant qualified human was "human doctor" for the task of diagnosing a disease and "human judge" for the task of deciding a parole case.

Results and Discussion

Preference for using algorithms was higher when performance data was provided ($M = 50.4$) compared with when it was not ($M = 28.6$; $t(199) = 17.5$, $p < .001$). The effect of performance data was highly significant for each of the nine tasks (see Table 4). We assigned each task the objectivity score that it received in Study 1 and then conducted an analysis of variance (ANOVA) in which performance condition, task objectivity, and their interaction were used to predict preference for using an algorithm relative to a human. This revealed main effects for performance condition ($F(1, 183) = 208.23$, $p < .001$), task objectivity ($F(8, 183) = 4.94$, $p = .026$), and a significant interaction ($F(8, 183) = 6.12$, $p = .013$). The effect of providing performance information was significant for each task. To explore the interaction, we divided tasks into "objective tasks" (rated as greater than 50, or the midpoint used to measure objectivity), and "subjective tasks" (rated as less than 50). The effect of providing performance data on preference for using an algorithm was significantly greater for objective tasks ($M_{perf\_data} = 55.1$, $M_{no\_perf\_data} = 34.5$; $t(199) = 14.37$, $d = .64$, $p < .001$) than it was for subjective tasks ($M_{perf\_data} = 47.3$, $M_{no\_perf\_data} = 31.7$; $t(199) = 10.35$, $d = .51$, $p < .001$).

The results of this study suggest that consumers' willingness to use an algorithm instead of a qualified human can be increased simply by demonstrating that the algorithm outperforms the human, although the increase is significantly smaller for subjective tasks. This study therefore provides further support for H_1 by showing that task objectivity reduces willingness to use algorithms even when consumers are made explicitly aware that the algorithm outperforms humans. However, participants remained roughly indifferent between humans and superior algorithms even for several of the most objective tasks. Indifference is insufficient for marketers interested in selling algorithm-based products and services. Importantly, this indifference is suboptimal for consumers when algorithms outperform humans (as is the case in each of the nine tasks in this study), because consumers who are indifferent will tend to select the default or status quo option (i.e., the worse-performing human; Dinner et al. 2011). These findings emphasize the need for other approaches to increase trust and use of algorithms. The next set of studies test H_2, which posits that reframing the task for which an algorithm is used can be one such approach.

Study 4

In Study 4, we attempt to increase consumers' trust in algorithms for subjective tasks by emphasizing that such tasks can benefit from quantitative analysis (relative to intuition), thus providing a test of H_2. This study also tests both the cognitive and affective dimensions of trust, confirming that our effects are driven by the cognitive dimension.

Method

Participants and design. Two hundred one Prolific Academic users ($M_{age} = 33$ years, 47% female) were randomly assigned to one of two conditions in which two tasks were described in such a way that emphasized either their quantitative components (the objective condition) or their intuitive components (the subjective condition). The tasks were recommending a movie and recommending a romantic partner, two consumer-relevant tasks that were rated as highly subjective in Study 1. Prolific Academic is a crowdsourcing website in which participants are less familiar with common experimental paradigms and more honest than participants on MTurk (Peer et al. 2017).

Procedure. In the subjective condition, participants read that, according to previous studies, the tasks were best accomplished by focusing on one's moods, emotions, and intuitions. In the objective condition participants read that, according to previous studies, the tasks were best accomplished by focusing on quantifiable data such as measured personality traits. Participants reported how much they would trust an algorithm relative to a human for the two tasks. Responses were entered on 0–100 scales, with the scale anchored at 0 ("trust human more"), 50 ("trust both equally"), and 100 ("trust the algorithm more"). The qualified human was specified as a friend for the movie task and a professional matchmaker for the dating task. Participants also reported how objective the tasks seemed.

To measure both the affective and cognitive components of consumers' trust in algorithms, we asked participants how much they agreed with the following questions: for cognitive evaluations, "I can see the benefits in algorithms that can perform this kind of task better than humans," "Algorithms that can perform this kind of task could be useful," and "I believe this kind of algorithm can perform well," and for affective evaluations, "Algorithms that can perform this kind of task better than humans make me uncomfortable," "Algorithms that can perform this kind of task go against what I believe computers should be used for," and "Algorithms that can perform this kind of task are unsettling." Alphas were .96 and .92, respectively, and all items were anchored at 0 ("not at all") and 100 ("completely"). We refer to the two kinds of evaluation as "effectiveness" and "discomfort" in the following analyses.

Results and Discussion

Participants trusted humans more than algorithms in both conditions for both tasks (ts > 1.89, ps $< .062$; comparing each mean to 50 ["trust both equally"]). However, emphasizing the quantitative approach to accomplishing the tasks succeeded at increasing trust in algorithms for both tasks (movie recommendation: $M_{subjective} = 31.5$, $M_{objective} = 45.1$; $t(199) = 4.04$, $p < .001$; romantic partner recommendation: $M_{subjective} = 36.4$, $M_{objective} = 43.8$; $t(199) = 2.09$, $p = .038$). Looked at another way, 36% of participants in the quantitative framing condition reported trusting an algorithm more than a human (>50 on the scale, collapsing across both tasks), compared with only 19% of participants in the intuitive framing condition ($\chi^2(1) = 6.60$, $p = .010$).

Emphasizing the quantitative approach to the tasks also increased the degree to which the tasks were viewed as objective (collapsing across the two tasks: $M_{subjective} = 32.1$, $M_{objective} = 40.7$; $t(199) = 3.11$, $p = .002$) and made algorithms seem more effective at the tasks ($M_{subjective} = 70.7$, $M_{objective} = 76.9$; $t(199) = 1.96$, $p = .052$). The manipulation had no effect, however, on participants' discomfort with the use of algorithms for the tasks (reverse-coded: $M_{subjective} = 72.0$, $M_{objective} = 72.1$; $t(199) = .27$, $p = .795$). Discomfort with the use of algorithms did have a significant negative effect on trust in the algorithm ($\beta = -.19$, $p < .001$), while perceived effectiveness had a significant positive effect on trust ($\beta = .34$, $p < .001$). However, a regression using both discomfort and perceived effectiveness simultaneously to predict trust showed that effectiveness remained a significant predictor ($\beta = .34$, $p < .001$) while discomfort did not ($\beta = -.004$, $p = .948$), suggesting that any initial effect of discomfort can be reduced if the algorithm is viewed as being effective.

A mediation analysis with 5,000 bootstrapped replications confirmed that perceived task objectivity and perceived effectiveness of algorithms sequentially mediated the relationship between task framing and trust in algorithms. As reported previously, task framing affected perceived task objectivity, which in turn affected the perceived effectiveness of algorithms for the tasks ($\beta = .20$, $p = .007$). The direct effect of task framing on trust in algorithms ($\beta = -10.49$, $p < .001$) was reduced but still significant when accounting for the mediators ($\beta = -6.04$, $p = .012$), and the indirect effect was significant ($\beta = -.002$, 95% confidence interval = [–.0002, –.0117]).

These results demonstrate that the perceived objectivity of a given task is malleable and that objectivity influences both the perceived effectiveness of algorithms for a task and self-reported trust in the algorithm for that task. These findings therefore suggest a practical marketing intervention that can be used to increase trust in and use of algorithms for tasks that are typically seen as subjective. We test this intervention using a field study in Study 5.

Study 5

In Study 5, we increase the external validity of the findings from Study 4 and test whether they can be practically useful for marketers. In this second Facebook advertising study, we manipulate the perceived objectivity of a subjective task in the context of a Facebook advertising campaign for algorithm- and human-based dating services.

Method

Participants and design. We created two advertisements for an algorithm-based dating service that either highlighted the benefits of using a quantitative approach to choosing a romantic partner (i.e., "studies show that using objective, quantifiable data is the best way to choose who to date") or did not (using the more neutral ad for algorithm-based dating advice from Study 2). We displayed these ads on Facebook, and they were seen by 13,621 Facebook users (39% female, mean age not observed).

Procedure. Our ads were clicked on 110 times in total. As in Study 2, participants who clicked on our ads were taken to a page explaining that we were studying consumers' trust in algorithms, and our dependent variable was the CTR of the ads.

Results and Discussion

Replicating the results of Study 4, logistic regression revealed that framing dating advice as benefiting from a quantitative approach increased the CTR (.87%) compared with the neutral ad (.39%, $\beta = -.004$, $\chi^2(1) = 3.74$, OR $= 2.21$, $p = .038$). Framing a task that is typically seen as highly subjective as in fact benefiting from quantitative data thus provides marketers with a practical tool for increasing consumers' willingness to use algorithm-based products for tasks in which algorithm aversion might otherwise occur.

Study 6

Study 6 attempts to increase the use of algorithms for subjective tasks in a different way. Instead of providing data regarding the algorithm's performance at the specific task in question or reframing the task itself as benefiting from quantification, we instead attempt to increase the perceived affective human-likeness of the algorithm by providing real examples of algorithms performing tasks that are typically thought of as requiring emotional and intuitive abilities (i.e., the kinds of abilities that machines are thought to lack and that are seen as necessary for subjective tasks; Gray, Gray, and Wegner 2007; Haslam et al. 2008; Inbar, Cone, and Gilovich 2010). Although social identity theory suggests that this approach may also make consumers less comfortable with the use of algorithms by challenging the belief in human distinctiveness from machines, the results of Study 4 suggest that the perceived effectiveness of an algorithm is a stronger influence than discomfort on trust in algorithms. We therefore expect that increasing affective human-likeness will make algorithms seem more effective at subjective tasks and thus increase the use of algorithms for such tasks, despite potentially also creating discomfort with the idea of a human-like algorithm.

Recall that prior research has consistently identified two dimensions of human-likeness, corresponding to cognitive and affective abilities (also called "agency" and "experience" or "human uniqueness" and "human nature"; Gray, Gray, and Wegner 2007; Haslam 2006). We chose to manipulate the affective dimension of human-likeness in this study because that dimension is the most relevant to subjective tasks, and because that dimension is the one commonly viewed as distinguishing humans from machines. Algorithms with affective human-likeness are therefore likely to be seen as both more useful for subjective tasks as well as more threatening to human distinctiveness from machines. This study therefore further helps determine the relative effects of perceived effectiveness (i.e., cognitive trust) and discomfort (i.e., affective trust) on the use of algorithms.

Method

Participants and design. Three hundred ninety-nine participants from Prolific Academic (49% female, $M_{age} = 35$ years) were assigned to one of four conditions in a 2 (affective human-likeness: high vs. low) $\times$ 2 (task framing: subjective vs. objective) between-subjects design.

Procedure. In the high-human-likeness conditions, participants read that algorithms can perform a range of tasks that are typically thought of as contributing to affective human-likeness, including creating music and art, predicting which songs will be popular, and understanding people's emotions. In the low-human-likeness conditions, participants read that algorithms cannot perform these kinds of tasks. In reality, algorithms can in fact perform these tasks. We chose these specific tasks because both creativity and emotional sensitivity are seen as fundamental components of human nature (i.e., the affective dimension of human-likeness; Haslam et al. 2005). Participants were asked to summarize the information they read in this section to ensure that they paid attention to the material. As a manipulation check, participants reported how much they agreed with the statement, "Algorithms that can perform this kind of task make humans seem less distinct from machines" on a 0 ("not at all") through 100 ("completely") scale.

Participants were then shown a graph of the value of the S&P 500 stock market index over the past year and were asked to estimate its value 30 days in the future. Before providing their initial estimate, participants were informed that the 5% most accurate estimates would be rewarded with a bonus payment six times larger than their base compensation, to incentivize accuracy and encourage serious engagement with the task.

After making their initial estimate, participants were told that an algorithm designed by an expert financial advisor had also made an estimate, were shown the algorithm's estimate, and were given the opportunity to revise their initial estimate. This paradigm is known as the judge–advisor system and is commonly used to measure reliance on advice by computing how much participants revise their initial estimate in response to external advice (Logg, Minson, and Moore 2019; Sniezek and Buckley 1995). Reliance on advice is measured as the difference between the final and initial estimates produced by each participant divided by the difference between the advice and the initial estimate.

The algorithm's estimate was accompanied by a manipulation of task objectivity. Specifically, in the objective framing condition, participants were told that there are clear mathematical relationships between economic measures such as supply and demand and the price of a stock, and that relying on these objective indicators is therefore the best way to estimate a stock's future value. In contrast, in the subjective framing condition, participants were told that human feelings and intuition are the primary drivers of stock prices, and that relying on these subjective factors is therefore the best way to estimate a stock's future value.

Finally, to measure perceived effectiveness of the algorithm and discomfort with the use of algorithms, participants were

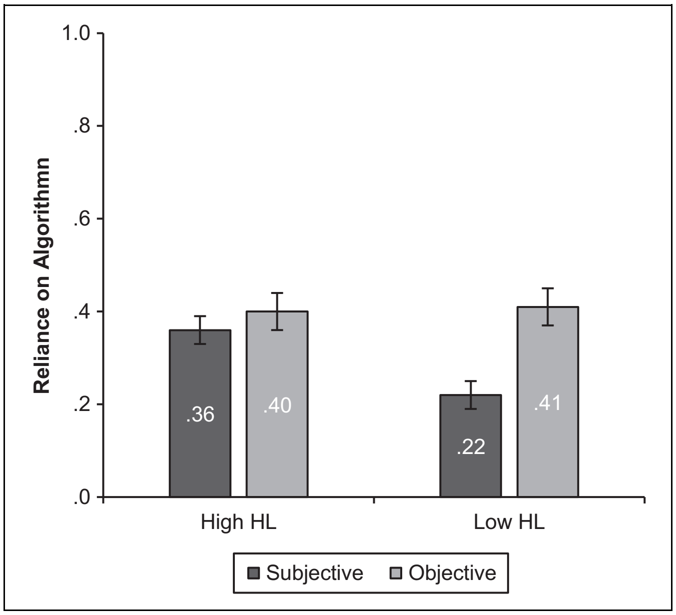

Figure 3. Algorithm aversion for a subjective task is eliminated for algorithms with high human-likeness.
Notes: Error bars represent standard errors. HL = human-likeness.

also asked how much they agreed with the following statements: "I believe this kind of algorithm can perform well," and "Algorithms that can perform this kind of task better than humans make me uncomfortable." These items were measured on a 0 ("not at all") through 100 ("completely") scale.

Results and Discussion

We excluded 33 participants (8.27% of the sample) whose summaries of the human-likeness manipulation indicated that they had not read the information, resulting in a final sample of 366. The results reported here are nearly identical if we include these participants. We first confirmed that our manipulation of human-likeness was effective: participants thought that humans were less distinct from machines in the high-human-likeness condition ($M = 45.8$) compared with the low-human-likeness condition ($M = 37.8$; $t(364) = 2.75$, $p = .006$). We computed reliance on the algorithm's advice by dividing the difference between participants' final and initial estimates by the difference between the advice and their initial estimate. This produces a measure that ranges, in most cases, from 0 (complete discounting of the advice) to 1 (complete reliance on the advice). A value of .30 thus corresponds to a 30% reliance on advice, which is the typical value found in the advice-taking literature (Soll and Larrick 2009).

A 2 × 2 ANOVA revealed that reliance on the algorithm was significantly affected by the task framing (objective vs. subjective; $F(1, 363) = 10.01$, $p = .002$), nonsignificantly affected by providing information about algorithms' human-likeness ($F(1, 363) = 2.24$, $p = .134$), and marginally affected by the interaction between these two factors ($F(1, 363) = 3.52$, $p = .060$). When participants were told that algorithms have low human-likeness, the effect of task framing was significant, as in prior studies ($M_{subjective} = .22$, $M_{objective} = .39$; $t(186) = 3.30$, $p = .001$). However, when participants were told that algorithms have high human-likeness, the effect of task framing was no longer significant ($M_{subjective} = .35$, $M_{objective} = .40$; $t(176) = 1.03$, $p = .303$). For a visual depiction of these results, see Figure 3.

Furthermore, we conducted another 2 × 2 ANOVA with belief in the algorithm's effectiveness as the dependent variable. This revealed no main effect of task framing ($F(1, 363) = .94$, $p = .339$), a marginally significant main effect of the algorithms' human-likeness ($F(1, 363) = 2.75$, $p = .099$), and a significant interaction between these two factors ($F(1,363) = 4.10$, $p = .044$). The interaction pattern was the same as the previous interaction: when participants were told that algorithms have low human-likeness, the effect of task framing on perceived effectiveness of the algorithm was significant ($M_{subjective} = 73.1$, $M_{objective} = 79.1$; $t(186) = 2.10$, $p = .037$). However, when participants were told that algorithms have high human-likeness, the effect of task framing was not significant ($M_{subjective} = 80.9$, $M_{objective} = 78.5$; $t(176) = .91$, $p = .366$).

Looked at another way, the effect of algorithm human-likeness on perceived effectiveness was significant in the subjective condition ($M_{high\ HL} = 80.9$, $M_{low\ HL} = 73.1$; $t(172) = 2.54$, $p = .014$), but not in the objective condition ($M_{high\ HL} = 78.5$, $M_{low\ HL} = 79.1$; $t(192) = .29$, $p = .834$). This indicates that increasing affective human-likeness increases the perceived effectiveness of algorithms for exactly the kind of task for which algorithms are typically not trusted or used.

The same ANOVA with discomfort as the dependent variable revealed no main effects but a significant interaction ($F(1,363) = 4.06$, $p = .045$). In the low-human-likeness condition, discomfort with algorithms was higher in the subjective task condition ($M_{subjective} = 28.8$, $M_{objective} = 21.0$; $t(186) = 2.02$, $p = .044$). In the high-human-likeness condition, discomfort was roughly equivalent across both task framings ($M_{subjective} = 26.2$, $M_{objective} = 30.3$; $t(176) = .91$, $p = .366$). Furthermore, the belief that algorithms make humans less distinct from machines was positively associated with discomfort ($\beta = .42$, $p < .001$) but was not associated with the perceived effectiveness of the algorithm ($\beta = .04$, $p = .321$). Finally, as in Study 4, discomfort on its own was negatively associated with reliance on the algorithm ($\beta = -.001$, $p = .040$), while perceived effectiveness on its own was positively associated with reliance ($\beta = .003$, $p < .001$). However, discomfort became nonsignificant when controlling for effectiveness ($\beta = -.0006$, $p = .486$), whereas effectiveness remained significant ($\beta = .003$, $p = .002$). This result suggests that consumers' discomfort with algorithms, stemming partly from decreasing human distinctiveness, has effects on consumers' willingness to use algorithms ceteris paribus, but that the effects of this discomfort are diminished when the algorithm is perceived as being highly effective.

These results also provide further evidence that algorithms are relied on less for tasks that seem subjective and suggest an additional method of eliminating this effect by increasing awareness of algorithms' affective human-likeness in terms of abilities typically perceived as distinguishing humans from machines. Importantly, these results also help tease apart the competing hypotheses regarding the role of algorithms' human-likeness in shaping consumers use of algorithms. Whereas social identity theory would suggest that algorithms that challenge the distinctives of humans from machines would create negative evaluations of those algorithms, a more cognitive perspective based on consumers' beliefs about algorithms' effectiveness suggests that decreasing human–machine distinctiveness would increase consumers' use of algorithms by making algorithms seem more useful. The fact that we found support for the second of these hypotheses suggests that effectiveness beliefs are stronger determinants of reliance on algorithms than discomfort stemming from intergroup challenges.

General Discussion

As algorithms become increasingly capable of outperforming humans at tasks ranging from making recommendations (for, e.g., music, movies, stocks) to diagnosing diseases and driving cars, a key issue is whether (or at least when or how quickly) and for what purposes humans will trust and use them. This research explored several aspects of this question. Specifically, in a series of six experiments with over 56,000 participants in total, we have studied how trust in and use of algorithms varies depending on how both the task at hand and algorithms are perceived.

Study 1 found that trust in algorithms for a given task is negatively related to perceived objectivity. Study 2 replicated these findings in a field study, finding that consumers click on ads for algorithm-based advice less than on ads for human-based advice when the task is subjective (dating advice), but not when the task is objective (financial advice). Study 3 tested the effect of making consumers explicitly aware of an algorithm's superior performance compared with humans, finding that such awareness is indeed a powerful influence on willingness to use algorithms. However, awareness of superior performance is not sufficient for creating a true preference for algorithms over humans—only indifference between the two—and is particularly ineffective for tasks that are more subjective. Study 4 manipulated perceived task objectivity, finding that reframing subjective tasks as being amenable to quantification and measurement increases trust in algorithms for those tasks. Study 5 replicated this finding in a field study. Finally, Study 6 showed that actual reliance on algorithms in an incentivized task is also lower when the task is seen as subjective but that this effect can be eliminated by providing real examples of algorithms with affective abilities. These findings support the conceptual model proposed in Figure 1, demonstrating that consumers perceive algorithms as less useful and are less comfortable with them for subjective tasks but that these effects are reduced when algorithms are seen has having high affective human-likeness, suggesting that the perceived effectiveness of algorithms is a stronger influence than discomfort on willingness to use algorithms.

Limitations and Directions for Future Research

Our work has several limitations. First, half of the studies relied on participants' reports of what they intended to do rather than direct evidence of what they did. More studies like Studies 2, 5, and 6, which used real behaviors as dependent variables, are needed to further calibrate real-world reactions to algorithms. Related to this, attempts to assess the costs of choosing not to rely on algorithms whose performance is superior to humans may be a valuable direction for future research. Nevertheless, it is encouraging that our self-report findings were replicated in field settings. Importantly, concerns about our results being driven by demand effects are reduced by our use of both field studies and incentive-compatible behaviors to replicate the effects observed in the self-report studies.

Our manipulation of human-likeness in Study 6 also leaves open the possibility of alternative manipulations of this construct in future research. We chose to manipulate specifically the affective dimension of human-likeness because this is the most closely relevant dimension to the performance of subjective tasks. Indeed, we expected (and found) that increasing affective human-likeness would make algorithms seem more effective and useful for subjective tasks. In this sense, affective human-likeness is inextricably linked to the effectiveness of the algorithm for subjective tasks, because performing such tasks requires affective abilities. While we manipulated the perceived effectiveness of the algorithm for the tasks in question directly in Study 3, the manipulation of human-likeness in Study 6 is not specific to the task in question (i.e., forecasting stock prices) but is more domain-general. This had the intended effect of making the algorithm seem more useful for specific subjective tasks. However, this manipulation also had the expected effect of increasing discomfort with the algorithm. The purpose of Study 6 was therefore to pit the effects of effectiveness and discomfort against each other, because both are affected by human-likeness and both were expected to affect reliance on algorithms, but in opposite directions. Nevertheless, alternative operationalizations of affective human-likeness are worth testing as well, as is increasing algorithms' cognitive human-likeness, or even their physical human-likeness to test how these other dimensions affect reliance on algorithms for different kinds of tasks.

Another limitation is that the descriptions of the algorithms were quite basic. More realistic presentations might include brochures, ads, websites, and videos. In addition, different presentation modes (in terms of both what and how information is presented) could be examined: use of algorithms may depend as much or more on how they are presented and accessed as on what is said about them. One aspect of this is the role of social influences on both individual decisions and overall adoption patterns (e.g., how do the p and q coefficients of the Bass diffusion model, which are related to advertising and word-of-mouth effects on adoption, differ between

Table 5. The Effects of Additional Factors on Preference for Using Humans Relative to Algorithms for Different Tasks.

	Preference for Using Humans For:			
	Personality Prediction	Joke Funniness Prediction	Wedding Planning	Stock Price Prediction
Privacy concerns	−.05	.05	.09	.22
Marketing concerns	.02	.05	.10	.27
Enjoy doing the task myself	.08	**.36**	**.47**	.27
Feed bad about myself	.22	**.36**	**.47**	.29
Feel less control	.28	**.44**	**.29**	**.34**
Requires emotion	**.35**	**.54**	**.48**	**.51**
Related to being human	.16	**.39**	**.65**	**.34**

Notes: Standardized regression coefficients are shown. Significant effects are in boldface. Positive coefficients indicate that greater concern with the factor in question is associated with greater preference for using a human (lower preference for using an algorithm).

algorithm adoption and adoption of other consumer and industrial products?; Bass 1969). Similarly, while descriptions and advertising of algorithms is undoubtedly important for initial adoption decisions, how actual experience with algorithms over extended periods of time influences their use is also an important topic for future research. For example, while we found that providing evidence about an algorithm's performance has a smaller effect on preference for using the algorithm when the task seems subjective, repeated exposure to algorithms effectively performing such tasks may convince more consumers of their value.

Our finding that consumers are relatively averse to algorithms that are used for subjective tasks is particularly relevant in light of the current trend toward affective computing, which is a growing industry intent on creating explicitly emotional algorithms and building them into products from driverless cars to refrigerators to digital personal assistants (Goasduff 2017; Kodra et al. 2013). Our results suggest that consumers will likely be skeptical about the emotional abilities of such algorithms but that convincing demonstrations of their effectiveness may ultimately increase willingness to use them for tasks typically thought to be "incompatible" with algorithms or computers. Future research should explore how different ways of presenting emotional algorithms to consumers affects their acceptance.

Future research could also explore additional factors that shape use of algorithms. For example, several factors that are not explicitly related to the nature of the algorithm or to the algorithm's performance might affect use, such as concerns about privacy or simply the enjoyment of performing a task oneself. While it is beyond the scope of this research to explore all such factors, we conducted a small survey to begin exploring such beliefs. Specifically, we asked 80 MTurk participants to report to what extent they had several potential concerns about using an algorithm for each of four tasks, as well as whether they would prefer to use an algorithm or a human for the tasks (predicting personality on the basis of Facebook likes, predicting joke funniness, planning a wedding, and predicting the price of stocks). For example, participants were asked whether using an algorithm for each task would involve privacy concerns, or whether they would feel badly about themselves for using an algorithm for each task. Table 5 depicts the results of this survey. Interestingly, concerns about privacy implications and companies using algorithms for marketing purposes were not significant predictors of willingness to use an algorithm for any task. In contrast, the enjoyment of doing the tasks oneself, feeling less control over the task, feeling bad about oneself if using an algorithm, the belief that the task is related to what it means to be human, and the belief that the task requires emotion were each significant predictors of preference for using an algorithm for at least two of these tasks. Future research is therefore needed to explore the role of these and other non-performance-related concerns in shaping consumers' use of algorithms.

One additional variable we explored in more depth is the degree to which algorithms can diminish the belief that humans are distinct from machines. Specifically, we hypothesized that the use of algorithms for subjective tasks could challenge the belief in human distinctiveness, which in turn would produce discomfort and negative evaluations of the algorithms. We found some support for the link between challenging human distinctiveness and discomfort, and in turn between discomfort and decreased trust in and use of algorithms (without controlling for perceived effectiveness). However, the link between task objectivity and belief in human distinctiveness was not consistently significant. Interestingly, algorithms that challenge the belief in human distinctiveness often have the simultaneous effects of producing discomfort and making the algorithms seem more effective. In our studies, however, effectiveness was a stronger predictor of actual use of algorithms. Nevertheless, the belief in human distinctiveness remains an interesting and potentially important factor in consumers' use of algorithms and other human-like technologies and should be explored more in future research. Web Appendix B describes our studies focusing on this factor.

Our findings are broadly consistent with the vast majority of existing research on consumers' perceptions and use of algorithms, insofar as the dominant theme of this research (summarized in Table 1) is algorithm aversion, which we also observe for many tasks. Our findings are inconsistent, however, with a recent article showing algorithm appreciation (i.e., greater reliance on advice from an algorithm than from a human; Logg, Minson, and Moore 2019). That article focuses on preferences for algorithms versus humans when making decisions for other people, whereas we focus on preferences for algorithms versus humans when making decisions for oneself. This difference may contribute to the divergent findings, because making decisions for other people involves greater psychological distance and therefore more abstract and logical thinking (Trope, Liberman, and Wakslak 2007). Similarly, consumers are more optimistic and focused on a choice's potential benefits when choosing for others versus choosing for themselves (Beisswanger et al. 2003; Polman 2012). These differences in choosing for self versus others may thus increase consumers' willingness to choose an algorithm, which represents a more abstract and logical approach than humans, and to focus on the potential gains rather than losses associated with using an algorithm. Future research could directly test whether algorithm aversion is more likely when choosing for oneself rather than for others.

Additional broader questions revolve around the potentially detrimental effects of increasing reliance on algorithms. For example, could an increased reliance on algorithms diminish people's capacity to think on their own and solve problems creatively or to perform the tasks that have been outsourced to algorithms? Could it diminish the utility and satisfaction that people receive from accomplishing tasks on their own? Furthermore, some have argued that the increasing use of algorithms in society can entrench economic and social inequalities by building discrimination into inflexible models applied on a large scale in contexts such as parole, hiring, and credit decisions (O'Neil 2016). Scholars and marketers who attempt to increase the use of algorithms to improve outcomes for consumers and society should be mindful of these concerns, striving to ensure that the promoted algorithms are both effective and fair.

Conclusion

We examined consumers' willingness to use algorithms to perform tasks in several areas and examined ways to increase their use. Our work highlights the important role of perceived task objectivity and of algorithmic human-likeness in shaping consumers' trust in and use of algorithms. While there appears to be an inexorable trend toward increasing use of algorithms, the pace at which they are adopted—as well as the areas where they will be adopted first—depends on several interrelated factors including in which areas companies develop algorithms for general use, how they market them, and how soon customers become comfortable with the idea of using algorithms to outsource decisions that affect their lives, in ways large and small. Our research suggests that marketers face a challenge in balancing the increasing capabilities of algorithms in subjective domains while also addressing consumers' lay beliefs that algorithms are ineffective at such tasks. Our results provide several practical strategies for achieving this balance.

Acknowledgment

The authors are grateful for methodological advice in conducting the field studies from Sandra Matz.

Associate Editor

Vikas Mittal

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

References

Armstrong, J. Scott (1980), "The Seer-Sucker Theory: The Value of Experts in Forecasting," *Technology Review*, 82 (7), 16–24.

Bass, Frank M. (1969), "A New Product Growth for Model Consumer Durables," *Management Science*, 15 (5), 215–27.

Bear, Adam and David G. Rand (2016), "Intuition, Deliberation, and the Evolution of Cooperation," *Proceedings of the National Academy of Sciences*, 113 (4), 936–41.

Beisswanger, Amy H., Eric R. Stone, Julie M. Hupp, and Liz Allgaier (2003), "Risk Taking in Relationships: Differences in Deciding for Oneself Versus for a Friend," *Basic and Applied Social Psychology*, 25 (2), 121–35.

Blattberg, Robert C. and Stephen J. Hoch (1990), "Database Models and Managerial Intuition: 50% Model + 50% Manager," *Management Science*, 36 (8), 887–99.

Bucklin, Randolph E., Donald R. Lehmann, and John D.C. Little (1998), "From Decision Support to Decision Automation: A 2020 Vision," *Marketing Letters*, 9 (3), 235–46.

Chopra, Kari and William A. Wallace (2003), "Trust in Electronic Environments," in *Proceedings of the 36th Annual Hawaii International Conference on System Sciences, HICSS 2003*. New York: Association for Computing Machinery, 1–10.

Davis, Fred, Richard Bagozzi, and Peter Warshaw (1989), "User Acceptance of Computer Technology: A Comparison of Two Theoretical Models," *Management Science*, 35 (8), 982–1003.

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

Deahl, Dani (2018), "How AI-Generated Music Is Changing the Way Hits Are Made," *The Verge* (August 31), https://www.theverge.com/2018/8/31/17777008/artificial-intelligence-taryn-southern-amper-music.

Diab, Dalia L., Shuang-Yueh Pui, Maya Yankelevich, and Scott Highhouse (2011), "Lay Perceptions of Selection Decision Aids in US and Non-US Samples," *International Journal of Selection and Assessment*, 19 (2), 209–16.

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.

Dietvorst, Berkeley J., Joseph P Simmons, and Cade Massey (2014), "Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err," *Journal of Experimental Psychology: General*, 143 (6), 1–13.

Dijksterhuis, A.P., Maarten W. Bos, Loran F. Nordgren, and Rick B. Van Baaren (2006), "On Making the Right Choice: The Deliberation-Without-Attention Effect," *Science*, 311 (5763), 1005–07.

Dinner, Isaac, Eric J. Johnson, Daniel G. Goldstein, and Kaiya Liu (2011), "Partitioning Default Effects: Why People Choose Not to Choose," *Journal of Experimental Psychology: Applied*, 17 (4), 332–41.

Eastwood, Joseph, Brent Snook, and Kirk Luther (2012), "What People Want from Their Professionals: Attitudes Toward Decision-Making Strategies," *Behavioral Decision Making*, 25 (5), 458–68.

Eisenhardt, Kathleen M. (1989), "Agency Theory: An Assessment and Review," *Academy of Management Review*, 14 (1), 57–74.

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.

Gibbs, Samuel (2016), "Google AI Project Writes Poetry Which Could Make a Vogon Proud," *The Guardian* (May 17), https://www.theguardian.com/technology/2016/may/17/googles-ai-write-poetry-stark-dramatic-vogons.

Gigerenzer, Gerd and Henry Brighton (2009), "Homo Heuristicus: Why Biased Minds Make Better Inferences," *Topics in Cognitive Science*, 1 (1), 107–43.

Goasduff, Laurence (2017), "Emotion AI Will Personalize Interactions," *Gartner* (January 27), https://www.gartner.com/smarterwithgartner/emotion-ai-will-personalize-interactions/.

Gray, Heather M., Kurt Gray, and Daniel M. Wegner (2007), "Dimensions of Mind Perception," *Science*, 619 (5812), 10–11.

Gray, Kurt and Daniel M. Wegner (2012), "Feeling Robots and Human Zombies: Mind Perception and the Uncanny Valley," *Cognition*, 125 (1), 125–30.

Greifeneder, Rainer, Herbert Bless, and Michel Tuan Pham (2011), "When Do People Rely on Affective and Cognitive Feelings in Judgment? A Review," *Personality and Social Psychology Review*, 15 (2), 107–41.

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.

Haslam, Nick (2006), "Dehumanization: An Integrative Review," *Personality and Social Psychology Review*, 10 (3), 252–64.

Haslam, Nick, Paul Baine, Lauren Douge, Max Lee, and Brock Bastian (2005), "More Human Than You: Attributing Humannness to Self and Others," *Journal of Personality and Social Psychology*, 89 (6), 937–50.

Haslam, Nick, Yoshihisa Kashima, Stephen Loughnan, Junqi Shi, and Caterina Suitner (2008), "Subhuman, Inhuman, and Superhuman: Contrasting Humans with Nonhumans in Three Cultures," *Social Cognition*, 26 (2), 248–58.

Herremans, Dorien, David Martens, and Kenneth Sörensen (2014), "Dance Hit Song Prediction," *Journal of New Music Research*, 43 (3), 291–302.

Inbar, Yoel, Jeremy Cone, and Thomas Gilovich (2010), "People's Intuitions About Intuitive Insight and Intuitive Choice," *Journal of Personality and Social Psychology*, 99 (2), 232–47.

Johnson, Devon and Kent Grayson (2005), "Cognitive and Affective Trust in Service Relationships," *Journal of Business Research*, 58 (4), 500–07.

Kodra, Evan, Thibaud Senechal, Daniel McDuff, and Rana El Kaliouby (2013), "From Dials to Facial Coding: Automated Detection of Spontaneous Facial Expressions for Media Research," in *Proceedings of the 10th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition*. New York: IEEE, 15–21.

Komiak, Sherrie and Izak Benbasat (2006), "The Effects of Personalization and Familiarity on Trust and Adoption of Recommendation Agents," *Management Information Systems Quarterly*, 30 (4), 941–60.

Krasnianski, Adriana (2015), "Meet Ross, the IBM Watson-Powered Lawyer," *PSFK* (January 29), http://www.psfk.com/2015/01/ross-ibm-watson-powered-lawyer-legal-research.html.

Li, Xin, Traci J. Hess, and Joseph S. Valacich (2008), "Why Do We Trust New Technology? A Study of Initial Trust Formation with Organizational Information Systems," *Journal of Strategic Information Systems*, 17 (1), 39–71.

Logg, Jennifer M. (2017), "Theory of Machine: When Do People Rely on Algorithms?" Working Paper No. 17-086, Harvard Business School, Harvard University.

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 (March), 90–103.

Longoni, Chaira, Andrea Bonezzi, and Carey Morewedge (2019), "Resistance to Medical Artificial Intelligence," *Journal of Consumer Research* (published online May 3), DOI:10.1093/jcr/ucz013.

Loughnan, Stephen and Nick Haslam (2007), "Animals and Androids: Implicit Associations Between Social Categories and Nonhumans," *Psychological Science*, 18 (2), 116–22.

Matz, Sandra, Michael Kosinski, Gideon Nave, and David Stillwell (2017), "Psychological Targeting as an Effective Approach to Digital Mass Persuasion," *Proceedings of the National Academy of Sciences*, 114 (48), 12714–19.

McDuff, Daniel, Rana El Kaliouby, Thibaud Senechal, May Amr, Jeffrey F. Cohn, and Rosalind Picard (2013), "Affectiva-MIT Facial Expression Dataset (AM-FED): Naturalistic and Spontaneous Facial Expressions Collected 'In-the-Wild,'" in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. New York: IEEE, 881–88.

Meehl, Paul E. (1954), *Clinical vs. Statistical Prediction: A Theoretical Analysis and a Review of the Evidence*. Minneapolis: University of Minnesota Press.

O'Neil, Cathy (2016), *Weapons of Math Destruction*. New York: Crown.

Önkal, Dilek, Paul Goodwin, Mary Thomson, Sinan Gönül, and Andrew Pollock (2009), "The Relative Influence of Advice from Human Experts and Statistical Methods on Forecast Adjustments," *Behavioral Decision Making*, 22 (4), 390–409.

Pavlou, Paul A. and David Gefen (2004), "Building Effective Online Marketplaces with Institution-Based Trust," *Information Systems Research*, 15 (1), 1–108.

Peer, Eyal, Laura Brandimarte, Sonam Samat, and Alessandro Acquisti (2017), "Beyond the Turk: Alternative Platforms for Crowdsourcing Behavioral Research," *Journal of Experimental Social Psychology*, 70, 153–63.

Polman, Evan (2012), "Effects of Self-Other Decision Making on Regulatory Focus and Choice Overload," *Journal of Personality and Social Psychology*, 102 (5), 980–93.

Promberger, Marianne and Jonathan Baron (2006), "Do Patients Trust Computers?" *Journal of Behavioral Decision Making*, 19 (5), 455–68.

Quackenbush, Casey (2018), "Painting Made by Artificial Intelligence Sells for $432,500," *Time* (October 26), https://time.com/5435683/artificial-intelligence-painting-christies/.

Rogers, Everett M. (1976), "New Product Adoption and Diffusion," *Journal of Consumer Research*, 2 (4), 290–301.

Shaffer, Victoria A., C. Adam Probst, Edgar C. Merkle, Hal R. Arkes, and Mitchell A. Medow (2013), "Why Do Patients Derogate Physicians Who Use a Computer-Based Diagnostic Support System?" Medical Decision Making, 33 (1), 108–18.

Simonite, Tom (2014), "IBM Watson's Plan to End Human Doctors' Monopoly on Medical Know-How," *MIT Technology Review* (July 21), https://www.technologyreview.com/s/529021/ibm-aims-to-make-medical-expertise-a-commodity/.

Slovic, Paul, Melissa L. Finucane, Ellen Peters, and Donald G. MacGregor (2004), "Risk as Analysis and Risk as Feelings: Some Thoughts about Affect, Reason, Risk, and Rationality," *Risk Analysis*, 24 (2), 311–22.

Sniezek, Janet A. and Timothy Buckley (1995), "Cueing and Cognitive Conflict in Judge-Advisor Decision Making," *Organizational Behavior and Human Decision Processes*, 62 (2), 159–74.

Soll, Jack B. and Richard P. Larrick (2009), "Strategies for Revising Judgment: How (and How Well) People Use Others' Opinions," *Journal of Experimental Psychology: Learning Memory and Cognition*, 35 (3), 780–805.

Trope, Yaacov, Nira Liberman, and Cheryl Wakslak (2007), "Construal Levels and Psychological Distance: Effects on Representation, Prediction, Evaluation, and Behavior," *Journal of Consumer Psychology*, 17 (2), 83–95.

Yeomans, Michael, Anuj Shah, Sendhil Mullainathan, and Jon Kleinberg (2019), "Making Sense of Recommendations," *Journal of Behavioral Decision Making*, (published online February 14), DOI:10.1002/bdm.2118.