

Resistance to Medical Artificial Intelligence

CHIARA LONGONI

ANDREA BONEZZI

CAREY K. MOREWEDGE

Author Note

Chiara Longoni (clongoni@bu.edu) is Assistant Professor of Marketing at the Questrom School of Business, Boston University, Boston MA. Andrea Bonezzi (abonezzi@stern.nyu.edu) is Associate Professor of Marketing, at the Stern School of Business, New York University, New York, NY. Carey K. Morewedge (morewedg@bu.edu) is Professor of Marketing and Everett W. Lord Distinguished Faculty Scholar at the Questrom School of Business, Boston University, Boston MA. Correspondence: Chiara Longoni. The authors gratefully acknowledge the helpful input of the editor, associate editor, and the reviewers. Supplementary materials are included in the web appendix accompanying the online version of this article.

Abstract

Artificial intelligence (AI) is revolutionizing healthcare, but little is known about consumer receptivity toward AI in medicine. Consumers are reluctant to utilize healthcare provided by AI in real and hypothetical choices, separate and joint evaluations. Consumers are less likely to utilize healthcare (study 1), exhibit lower reservation prices for healthcare (study 2), are less sensitive to differences in provider performance (studies 3A-3C), and derive negative utility if a provider is automated rather than human (study 4). *Uniqueness neglect*, a concern that AI providers are less able than human providers to account for their unique characteristics and circumstances, drives consumer resistance to medical AI. Indeed, resistance to medical AI is stronger for consumers who perceive themselves to be more unique (study 5). Uniqueness neglect mediates resistance to medical AI (study 6), and is eliminated when AI provides care (a) that is framed as personalized (study 7), (b) to consumers other than the self (study 8), or (c) only supports, rather than replaces, a decision made by a human healthcare provider (study 9). These findings make contributions to the psychology of automation and medical decision making, and suggest interventions to increase consumer acceptance of AI in medicine.

Keywords: automation, artificial intelligence, healthcare, uniqueness, medical decision making

Artificial Intelligence (AI) is revolutionizing healthcare. Medical AI applications are manifold and pervasive. IBM's Watson diagnoses heart disease (Hutson 2017), chatbots dispense medical advice for the United Kingdom's National Health Service (O'Hear 2017), apps like SkinVision detect skin cancer with expert accuracy (Haenssle et al. 2018) and algorithms identify eye diseases such as diabetic retinopathy just as well as specialized physicians (Gulshan, Peng and Coram 2016; Abramoff et al. 2018). Medical AI is forecasted to become a \$10 billion market in the United States by 2025 (Bresnick 2017), pervade 90% of hospitals (Das 2016), and replace as much as 80% of what doctors currently do (Khosla 2012).

This revolution is predicated on the premise that AI can perform with expert-level accuracy (Gallagher 2017; Leachman and Merlino 2017) and deliver cost-effective healthcare at scale (Esteva et al. 2017). In some cases, AI even outperforms human healthcare providers. When the performance of IBM Watson was compared to human experts for 1,000 cancer diagnoses, Watson found treatment options that doctors missed in 30% of the cases (Lohr 2016). When U.K. researchers compared the accuracy of triage diagnoses made by doctors to those made by AI, doctors were found to be correct 77.5% of the times, whereas AI reached an accuracy rate of 90.2% (Donnelly 2017).

Whereas much is known about medical AI's accuracy, cost-efficiency, and scalability, little is known about patients' receptivity to medical AI. Yet, patients are the ultimate consumers of medical AI, and will determine its adoption and implementation both directly and indirectly (Israyelyan 2017). Medical AI empowers patients to interact with autonomous tools that collect information and make decisions without the intervention of a physician and outside of clinical settings (Yu, Beam and Kohane 2018). In this context, consumers will *directly* drive medical AI's adoption. In conventional clinical settings, where patients' interactions with AI may be still

mediated by physicians, consumers will *indirectly* determine medical AI's adoption. How medical AI is integrated into healthcare delivery may profoundly influence patients' satisfaction with medical services, a metric that determines key outcomes for healthcare institutions (Centers for Medicare and Medicaid Services 2019), ranging from federal reimbursements to long-term financial sustainability (Mehta 2015).

In our research, we explore consumers' receptivity to medical AI. We use the term AI to refer to any machine that uses any kind of algorithm or statistical model to perform perceptual, cognitive, and conversational functions typical of the human mind, such as visual and speech recognition, reasoning, and problem solving. We suggest that consumers may be more reluctant to utilize medical care delivered by AI providers than comparable human providers, because the prospect of being cared for by AI providers is more likely to evoke a concern that one's unique characteristics, circumstances, and symptoms will be neglected. We refer to this concern as *uniqueness neglect*.

Our research makes two fundamental contributions. First, we contribute to research on the psychology of automation (Castelo, Bos and Lehman 2018; Dawes, Faust and Meehl 1989; Dietvorst, Simmons and Massey 2014; Meehl 1954; Granulo, Fuchs and Puntoni 2018) by extending its scope to the study of consumers in medical settings. Second, we advance the theoretical understanding of consumer decision making in the medical domain (Bolton et al. 2007; Botti and Iyengar 2006; Botti, Orfali and Iyengar 2009; Kahn and Baron 1995; Kahn et al. 1997; Keller and Lehmann 2008) by identifying a novel psychological mechanism that shapes decisions related to consumer utilization of innovative healthcare services.

THEORETICAL BACKGROUND

In 1954, Paul Meehl published the first systematic review of studies that compared forecasts made by humans (i.e., clinical judgments) to forecasts made by statistical models (i.e., statistical judgments). Meehl's review showed that statistical models were more accurate than humans at predicting various outcomes. Since then, a vast body of empirical evidence has corroborated the notion that statistical models outperform human intuition when making predictions about disparate outcomes such as student grades (Dawes 1979), employee performance (Kuncel et al. 2013), parole bail violations (Kleinberg et al. 2017), and market demand (Sanders and Manrodt 2003; for a review, see Dawes et al. 1989).

Given the superior accuracy of statistical models over human intuition, people should prefer to follow the advice of statistical models than human intuition. Yet, in most cases people do not (Castelo et al. 2018; Logg 2018). Consumers prefer to rely on friends rather than on computerized recommendation systems for suggestions about books, movies, and jokes (Sinha and Swearingen 2001; Yeomans et al. 2017), lose confidence in statistical models more quickly than in humans after seeing both make the same mistake (Dietvorst et al. 2014), and place greater weight on the same advice given by human experts than statistical models (Önkal et al. 2009). Professionals exhibit similar preferences. Recruiters trust their own judgment more than the recommendations of algorithms (Highhouse 2008), and auditors tend to ignore fraud predictions made by computerized decision aids (Boatsman, Moeckel and Pei 1997).

Resistance to statistical models has been documented in diverse non-medical settings, including predicting employee performance (Kuncel et al. 2013), market demand (Sanders and Manrodt 2003), fraud (Boatsman et al. 1997), crime (Kleinberg et al. 2017; Wormith and Goldstone 1984), and consumer preferences (Sinha and Swearingen 2001; Yeomans et al. 2017).

In the medical context, however, existing research is limited and has mostly focused on physicians' receptivity to statistical models. This research shows that although statistical models reliably outperform doctors (Dawes et al., 1989; Eastwood, Snook and Luther 2012; Grove and Meehl 1996; Grove et al. 2000; Sawyer 1966), doctors generally prefer to rely on their own intuition rather than on statistical models (Böckenholt and Weber 1992; Keefe et al. 2005), and are evaluated as less professional and competent if they do rely on computerized decision-aids (Palmeira and Spassova 2013; Shaffer et al. 2013).

Adoption of medical AI, however, ultimately depends on consumer receptivity to this new technology. To date, only one research endeavor has empirically examined consumer receptivity to automated medical providers. Promberger and Baron (2006) found that people are more likely to follow the recommendation of a physician than the recommendation of a computer. In this research paradigm, however, participants were given no information about the performance of either the physician or the computer, thus they may have assumed that the computer's performance was inferior to the doctor's. If a difference in perceived performance is the driver of consumer preferences, then consumers would not truly exhibit resistance to automated medical providers. Consumer preferences would be better characterized as driven by a perception that automated providers deliver objectively inferior care than physicians. To address this possibility, in our research we explicitly inform consumers of the performance of AI and human providers. Providing explicit performance information helps us identify whether consumer receptivity toward medical AI is simply due to different performance perceptions, or to a less normative psychological driver.

THE CURRENT RESEARCH

We propose that consumers may be more reluctant to utilize healthcare delivered by AI providers than healthcare delivered by comparable human providers, even when explicitly informed about the performance of the providers. We further propose that this occurs because the prospect of being cared for by an automated provider evokes a concern that one's unique characteristics, circumstances, and symptoms will be neglected. We refer to this concern as *uniqueness neglect*. We argue that uniqueness neglect emerges out of a mismatch between two fundamental beliefs. Whereas consumers view themselves as unique and different from others (Snyder and Fromkin 1980; Brewer 1991), consumers view machines as capable of operating only in a standardized and rote manner that treats every case the same way (Haslam 2006; Montagu and Watson 1983).

The belief in self-uniqueness is central to many social psychological theories (Baumeister 1998; Brown 1998; Dunning 1993; Epstein 1990; Greenwald 1980; Steele 1988; Taylor and Brown 1988). Social comparison research is replete with theories and data indicative of people's tendency to perceive their opinions, beliefs, attitudes, skills and characteristics as idiosyncratic to themselves. People rate themselves as different from others on driving skills, managerial shrewdness, coping with the death of a loved one, juggling, computer programming, and susceptibility to bias in their judgments (Blanton et al. 2001; Kruger 1999; Larwood and Whittaker 1977; Scopelliti et al. 2015; Svenson 1981; Windschitl et al. 2003). Perceptions of self-uniqueness even pervade predictions for future events. Undergraduate students who were given insurance company longevity data, predicted they would live ten years longer than the age predicted actuarially (Snyder 1978). Some psychological theories view the pursuit of self-uniqueness as a motivational drive, and suggest that people have a need to see themselves as

unique and distinct from others (i.e., uniqueness theory, Snyder and Fromkin 1980; optimal distinctiveness theory, Brewer 1991). Situations that threaten self-uniqueness, result in feelings of anxiety (Fromkin 1968; Horney 1937; Maslow 1962) and drive people to engage in behaviors aimed at restoring their desired self-distinctiveness (Sensening and Brehm 1968).

The belief that machines treat every case in the same way derives from heuristics people use to distinguish inanimate objects, such as machines and computers, from humans. One such heuristic is the absence, in machines, of abilities indicating cognitive flexibility. Cognitive flexibility refers to the capacity to adapt cognitive processes to new and unexpected conditions in the environment (Cañas et al. 2003). It is typically captured by imagination, creativity, and openness (Halsam 2006; Leyens et al. 2001). Indeed, people use perceptual cues and former experiences with inanimate objects to characterize computers as rote, rigid, and inflexible (Loughnan and Haslam 2007; Montagu and Watson 1983). Computers are perceived as capable of operating only in a standardized and repetitive fashion, based on a pre-programmed set of procedures (Nissenbaum and Walker 1998).

We propose that in a healthcare context, these beliefs manifest in consumers viewing their health-related characteristics—whether biological, psychological, or behavioral—as unique and distinct from those of an average person. Indeed, an empirical illustration ($N = 149$) found that people view their health-related circumstances and medical conditions to be unique and depart from standard clinical criteria more than they view health-related circumstances and medical conditions of other people. For example, when given a description of insomnia and the causes typically implicated in causing insomnia, people thought that episodes of insomnia occurring to them had more unique causes ($M_{\text{self}} = 3.60$, $SD_{\text{self}} = 1.69$) than episodes of insomnia occurring to other people ($M_{\text{other}} = 2.65$, $SD_{\text{other}} = 1.83$, $t(147) = 3.29$, $p = .001$). When given a

description of the typical symptoms of a headache, people thought that their headaches had more unique symptoms ($M_{\text{self}} = 2.87$, $SD_{\text{self}} = 1.59$) than a headache of other people ($M_{\text{other}} = 2.32$, $SD_{\text{other}} = 1.54$, $t(147) = 2.11$, $p = .03$). When given a description of the typical symptoms of anxiety, people thought that their anxiety had more unique symptoms ($M_{\text{self}} = 3.73$, $SD_{\text{self}} = 1.64$) than anxiety of other people ($M_{\text{other}} = 2.72$, $SD_{\text{other}} = 1.66$, $t(147) = 3.77$, $p < .001$; details in web appendix A). In sum, people seem to perceive their health-related characteristics to be more unique than the same characteristics in other people (e.g., “I caught *a* cold. You caught *the* cold.”).

In contrast to the belief that one’s health characteristics tend to be unique, we argue that consumers view medical care delivered by AI providers as standardized and calibrated for an average patient. Consequently, we propose that consumers will be more resistant to utilize healthcare delivered by AI providers than healthcare delivered by comparable human providers out of a concern that the unique facets of their case will be neglected. In a diagnostic context, for example, consumers might believe that an automated dermatologist will not take into account, to the same extent that a human dermatologist would, skin characteristics that they regard as unique. Similarly, when considering a recommendation on whether to undergo surgery, consumers might believe that a computer might not take into account what they regard as unique symptoms, to the same extent that a doctor would.

Uniqueness neglect builds upon what is anecdotally referred to as the “broken leg” hypothesis (Meehl 1954; Grove and Meehl 1994). According to the broken leg hypothesis, people’s distrust statistical judgments because they fear that a statistical model might omit key evidence. The classic example used by Meehl and colleagues is that of a sociologist who uses a regression model for predicting whether people will go to the movies on a certain night. Taking

into account demographic information and past behavior, the model predicts that a certain person is very likely to go to the movies on a specific Friday. Yet, the model overlooks the fact that the person has a broken leg, thus is unable to leave home. Consistent with the broken leg hypothesis, our uniqueness neglect account hinges on people's beliefs about how statistical models operate. However, we argue that the mere belief that statistical models might be mis-specified cannot fully explain resistance to medical AI. Resistance to medical AI originates in a mismatch between a belief about how automated providers operate, and a belief about the self. Only when consumers view themselves as unique and different from others will they exhibit resistance to medical AI.

We propose that uniqueness neglect is an important driver of consumer resistance to medical AI, but acknowledge that this is a complex phenomenon, most likely multiply determined. To illustrate, consumers may exhibit resistance to medical AI because they believe that AI provides objectively inferior, or more expensive care. They may eschew AI providers because they lack characteristics they deem a medical provider ought to have, such as warmth and conscientiousness (Haslam 2006; Haslam, Bain, Douge, Lee and Bastian 2005). Or they may seek to avoid medical AI because they feel less able to offload the responsibility for a potentially negative outcome onto a machine, as compared to a human being (Promberger and Baron 2006). Throughout our studies, we aim to show that uniqueness neglect is a powerful and distinct driver of resistance to medical AI that operates above and beyond these other factors.

OVERVIEW OF THE STUDIES

We tested our uniqueness neglect hypothesis across eleven studies¹. Studies 1 to 4 provided a systematic examination of consumer preferences for healthcare delivered by AI versus human providers. In the first two studies, we tested resistance to medical AI by examining how the likelihood of utilizing healthcare (study 1) and willingness to pay for healthcare (study 2) vary across automated versus human providers with the same performance. In studies 3A, 3B, and 3C, we examined whether consumer resistance to medical AI persists when automated providers offer superior performance than human providers. In study 4 we used a choice-based conjoint exercise to estimate the implicit (dis)utility associated with an automated compared to a human provider, and relative to other attributes of the providers.

Studies 5 to 9 focused on the role of uniqueness neglect as an underlying driver of resistance to medical AI via mediation and moderation. Study 5 provided process evidence by use of a measured moderator, testing whether participants with a greater sense of uniqueness are more reluctant to utilize an automated than a human provider. Study 6 provided process evidence via mediation, testing whether uniqueness neglect mediates the effect of provider on likelihood to utilize healthcare, vis-à-vis other potential mechanisms. Studies 7 and 8 provided process evidence via moderation, testing for reduced resistance toward medical AI when uniqueness neglect is attenuated, such as when care is framed as personalized and when choice to pursue medical care affects an average person rather than the self. Finally, study 9 tested whether resistance to medical AI is mitigated if an automated provider supports, rather than replaces, a human provider who remains in charge of making the critical medical decision. Figure 1 presents an overview of the empirical package.

¹ For all experiments, we determined sample size in advance and analyzed the data only after all responses had been collected. We report all conditions, manipulations, data exclusions, and critical measures related to our hypothesis testing. In some experiments, we collected additional measures after the key hypothesis-related measures for exploratory purposes. Discussions of exploratory measures are available from the authors upon request.

* Insert Figure 1 about here *

STUDY 1: LIKELIHOOD TO UTILIZE HEALTHCARE BY PROVIDER

Study 1 investigated whether consumers are less likely to pursue medical care administered by automated versus human providers with comparable performance. We explored this question in a real-world setting, offering college students the opportunity to schedule an appointment for a diagnosis of their stress level. Stress is a ubiquitous and relevant health concern for college students. Over 85% of students report feeling overwhelmed by stress during college (National College Health Assessment 2015; Center for Collegiate Mental Health 2016). In all conditions, in addition to providing information about the provider's performance, we specified that there would be no interaction with the medical provider, we informed participants of the cost for the procedure (i.e., \$0), and we indicated the content and timeline of the diagnosis. Thus, any resistance to utilize medical AI in this context would be related to its decision making, and not to a preference for interacting with a human rather than an automated provider, or to different perceptions of healthcare costs, content, or timeframes.

Procedure

Two-hundred and twenty-eight students ($M_{\text{age}} = 19.6$, $SD = 1.1$, 54.9% female) from two large universities in the United States participated in the study in exchange for course credit. Participants read that the study examined how students manage stress. To ensure that knowledge about the medical domain was consistent across conditions, students first read a leaflet by the

American Institute of Stress that defined stress, and described its signs and symptoms. Next, students learned that the first step toward stress management is measuring one's stress level. A detailed explanation of how stress would be measured followed.

To measure a person's level of stress, assessment tools typically combine a Salivary Cortisol Test with a Stress Questionnaire. The Salivary Cortisol Test measures the levels of the stress hormone ACTH and cortisol in your saliva. The test is based on a saliva swab you can take on your own using the swab kit shown below [*photo*]. The Stress Questionnaire consists of a survey that includes questions designed to detect behavioral symptoms of stress, measure the impact of stress on your mental health, and assess how you cope with stress. The data from the Salivary Cortisol Test and the data from the Stress Questionnaire are then combined to obtain a person's stress profile.

We then offered participants the opportunity to have their stress level measured. We specified that measurement would entail no interaction with the healthcare provider that would analyze their data. We also specified the content of the diagnosis and how the diagnosis would be delivered.

The data from the Salivary Cortisol Test and from the Stress Questionnaire are sent to the American Institute of Stress to be analyzed. A typical report from this assessment include: an easy to read chart, a qualitative assessment, and a recommended course of action to help you manage stress. The report can be received via email, downloaded online via a passport-protected web site, or mailed through regular mail.

The rest of the script differed between participants. In the *human provider* condition, students read that the data collected would be analyzed by a physician. In the *automated provider* condition, students read that the data would be analyzed by a computer. Human and automated provider were described as having the same accuracy rate. Diagnosis by either provider was described as free of charge for college students.

The data is analyzed by a physician [*computer*]. The physician uses her own judgment [*computer uses an algorithm*] developed based on prior cases, to assess your level of stress and make recommendations. In the past, physicians [*computers*] making these kinds of diagnoses have been accurate 82-85% of the times. That is, in the past, physicians [*computers*] correctly identified the right level of stress and made appropriate recommendations 82-85% of the times. This service is provided to college

students free of charge.

After reading this information, students indicated whether they wanted to schedule an appointment to complete the stress assessment and measure their level of stress on campus. (“Are you interested in scheduling a time in the next two weeks to come back to the lab and measure your level of stress? Please click “yes” to make an appointment, “no” if you are not interested”). Participants then indicated age and gender. Finally, we debriefed participants, informed them about the aim of the research, and referred them to further information about stress management (stimuli in web appendix B).

Results and Discussion

We compared the proportion of students who indicated that they wanted to schedule an appointment to complete the stress assessment as a function of provider (human vs. automated). Whereas 39.8% of students signed up for the service when the analysis would be carried out by a human provider, only 26.7% signed up for the service when the analysis would be carried out by an automated provider ($\chi^2(1) = 4.4; p = .03$). Age and gender did not significantly influence any of the analyses in this and in all subsequent studies, thus we refrain from discussing these variables further.

The results of study 1 offered initial evidence of reluctance to utilize healthcare delivered by AI as compared to human providers. Students were less likely to schedule a diagnostic stress assessment when the provider was automated than when the provider was human, even though both providers were equally accurate. However, because of the between-subject nature of the design, participants only received information about one of the providers, thus they could have

inferred that accuracy rates were at ceiling for the human provider, and at floor for the automated provider. In the next three studies, we addressed this potential inference by informing participants of the performance of both providers and asking them to express a relative preference in a joint evaluation paradigm.

STUDY 2: WILLINGNESS-TO-PAY FOR HEALTHCARE BY PROVIDER

In Study 2 we made the relative performance of the providers salient and evaluable (Hsee et al. 1999), to provide a more conservative test of whether assumptions about differences in performance drives resistance to medical AI. Also, to probe the robustness of the effect observed in study, we explored how much consumers are willing to pay to receive medical care from a human versus an automated provider, and vice-versa. We set a price for a certain medical service provided by a human (vs. automated) provider, and measured how much consumers would be willing to pay to switch to an equally accurate automated (vs. human) provider.

Procedure

One hundred and three respondents ($M_{\text{age}} = 33.6$, $SD = 10.1$, 55.3% female) from Amazon Mechanical Turk participated in exchange for monetary compensation. Participants read about the same diagnostic tool used in study 1 (i.e., assessment of a person's level of stress). Different from study 1, in study 2 all participants were informed that the analysis of their data could be carried out either by a doctor (*human provider* condition) or by a computer (*automated provider* condition). We specified that the two providers had the same performance (i.e., an

accuracy rate of 89%), that there would be no interaction with the provider performing the analysis, and that the medical diagnosis would entail the same information and would be delivered according to the same timeline. Between subjects, we specified the price (i.e., \$50) to have the analysis performed by a doctor in one condition and by a computer in the other, and asked participants to indicate their willingness to pay to switch to the computer or to the doctor, respectively.

[*Human provider*] Analysis by the computer costs \$50. How much would you be willing to pay to have the analysis performed by the doctor?

[*Automated provider*] Analysis by the doctor costs \$50. How much would you be willing to pay to have the analysis performed by the computer?

Participants reported their willingness to pay to switch provider on an 11-point scale ranging from \$0 to \$100, with the mid-point anchored at the reference price of \$50 (stimuli in web appendix B). Participants ended the survey by indicating their age and gender.

Results and Discussion

Participants were willing to pay less to use an automated provider when the default was a human provider ($M = \$31.22$, $SD = 18.89$) than to use a human provider when the default was an automated provider ($M = \$45.00$, $SD = 20.62$, $t(101) = 3.52$, $p = .001$). Compared to the reference price of \$50, the decrease in WTP for the alternative provider was greater when switching from a human to an automated provider ($-\$18.78$) than when switching from an automated to a human provider ($-\$5.00$; $t(101) = 3.52$, $p = .001$).

Together, the results of studies 1 and 2 showed that consumers are more resistant to utilize healthcare delivered by medical AI than comparable care delivered by a human provider.

Participants were less likely to utilize (study 1), and exhibited a lower reservation price for a diagnostic service when the analyses it provided were performed by AI rather than by a human (study 2). In both studies 1 and 2, human and automated providers had the same performance, suggesting that inferences about differences in performance could not explain the results. In the next study, we probed the role of performance information in an even more conservative way, by manipulating the performance of the automated provider to be either *equal* or *superior* to that of the human provider.

STUDIES 3A-3C: RELATIVE PREFERENCE FOR MEDICAL PROVIDER

In Study 3, we sought to test the boundaries of the role of performance information in driving resistance to medical AI. In each of three separate studies (3A-3C), we systematically varied the performance of an automated provider to be either equal or superior to the performance of a human provider. If consumers exhibit resistance to medical AI even when AI's performance is clearly superior to that of human providers, then objective performance, alone, cannot fully explain consumer preferences for providers.

To further probe the generalizability of resistance to medical AI, we assessed participants' preference in three different domains: prevention (3A), diagnosis (3B), and treatment (3C). In each domain, we identified an application of AI currently available to consumers and pitted it against its human counterpart: a software for early detection of skin cancer against a dermatologist (3A; based on the applications *DermaCompare* and *SkinVision*), a conversational robot that triages potential emergencies against a nurse (3B; based on the *Babylon* chatbot used in the U.K. National Health Service), and a robot for cardiac surgery against a

surgeon (3C; based on the *DaVinci* surgical system).

Across the three studies, we systematically varied the following factors: (a) *framing* of the performance of each provider (i.e., focusing on success by use of accuracy rates vs. focusing on failure by use of complication rates); (b) *degree of agency* (i.e., consumers maintain a high level of agency in the interaction with the provider vs. consumers have no interaction with the provider); and (c) *format* used to indicate the past performance of each provider (i.e., numeric percentages only vs. numeric percentages and color-coded visual categories; table 1).

* Insert Table 1 About Here *

Procedure

A total of 744 participants recruited on Amazon Mechanical Turk completed these three studies in exchange for money: 250 in study 3A ($M_{\text{age}} = 35.5$, $SD = 11.0$; 47.6% females), 253 in study 3B ($M_{\text{age}} = 35.4$, $SD = 11.2$; 43.9% females), and 241 in study 3C ($M_{\text{age}} = 37.0$, $SD = 11.5$; 44.6% females). The experimental procedure was identical in all three studies. Participants read that we were interested in understanding their preferences for medical care providers.

Participants read detailed information about a specific medical procedure that varied by study: screening for early detection of skin cancer in study 3A; triaging of a potential emergency in study 3B; pacemaker implant surgery in study 3C. We provided participants with verbatim descriptions of these services that are currently available to patients and consumers. Participants then reviewed information about two different providers, labeled “provider X” and “provider Y.”

Participants were then randomly assigned to a 2x3 between-subjects design. The two-level factor was the composition of the choice-set. Half of the participants were assigned to choose between two human providers (both provider X and Y were human); the other half was assigned to choose between a human and an automated provider (provider X was human and

provider Y was automated). The three-level factor was the relative performance of the two providers. The two providers were either equivalent in terms of their performance (provider X = provider Y), or one provider was slightly better than the other provider (provider X < provider Y), or one provider was better than the other provider to a larger extent (provider X << provider Y). In studies 3A and 3B, performance was expressed as success (i.e., past accuracy rate). In study 3C, performance was expressed as failure (i.e., past adjusted rate of complications). Participants read precise instructions on how to interpret these performance metrics, and completed comprehension checks to indicate that they correctly understood how to interpret the performance metrics before answering the primary dependent variable. Table 1 summarizes the differential levels of performance.

Participants indicated their preferred medical care provider on a 7-point scale (1 = *definitely provider X*, 4 = *indifferent*, 7 = *definitely provider Y*). The word “provider” was substituted with the relevant term for the provider (whether human or automated) in each medical domain (i.e., “dermatologist,” “nurse,” “surgeon”). To conclude, participants indicated their age and gender. Stimuli in web appendix B.

Results and Discussion

We performed all analyses with and without those participants who responded incorrectly to the comprehension checks. Here we report the results for the subset of participants who responded correctly to the comprehension checks, corresponding to the following sample sizes: 205 in study 3A, 227 in study 3B, and 235 in study 3C. Results were identical, both in terms of direction and statistical reliability, if no exclusion criterion was applied.

Our main hypothesis was that participants would be more reluctant to choose an automated provider if a human provider was available, even when the automated provider performed better than the human provider. To test this hypothesis, we conducted 2x3 ANOVAs on provider preferences. The analyses revealed significant main effects of choice-set (study 3A: $F(1, 199) = 30.07, p < .001, \eta_p^2 = .131$; study 3B: $F(1, 221) = 62.94, p < .001, \eta_p^2 = .222$; study 3C: $F(1, 229) = 40.31, p < .001, \eta_p^2 = .150$), significant main effects of provider performance (study 3A: $F(2, 199) = 81.94, p < .001, \eta_p^2 = .452$; study 3B: $F(2, 221) = 86.82, p < .001, \eta_p^2 = .440$; study 3C: $F(2, 229) = 105.14, p < .001, \eta_p^2 = .479$), and no significant interaction. We then conducted pairwise comparisons at each level of performance. To illustrate, in study 3A and with respect to the condition $X = Y$, we compared preference for provider Y in the human-automated condition ($M = 3.11$) with preference for provider Y in the human-human condition ($M = 4.00$); with respect to the condition $X < Y$, we compared preference for provider Y in the human-automated condition ($M = 4.64$) with preference for provider Y in the human-human condition ($M = 6.06$). Finally, with respect to the condition $X \ll Y$, we compared preference for provider Y in the human-automated condition ($M = 5.97$) with preference for provider Y in the human-human condition ($M = 6.62$). These analyses resulted in three pairwise comparisons in each study. Across all three studies, participants were more reluctant to choose provider Y when provider Y was automated than when it was human (all F 's ≥ 4.29 , all p 's $\leq .04$; table 2; figure 2). Reluctance to choose provider Y held not only when provider Y delivered the same performance as provider X, but even held when provider Y delivered superior performance to provider X.

* Insert Table 2 About Here *

* Insert Figure 2 About Here *

Together, studies 3A-3C provided evidence that consumer resistance to medical AI emerges across a variety of medical domains. Resistance to medical AI was robust across different types of medical care (i.e., prevention, diagnosis, treatment), framing of the providers' performance rates (i.e., accuracy/success versus complications/failure), degree of agency (i.e., consumers maintain a high level of agency in the interaction with the provider vs. consumers have no interaction with the provider), and format of performance information (i.e., numeric vs. numeric and categorical). Furthermore, the results provide strong evidence that consumer resistance to medical AI is not merely driven by the belief that the performance of an automated provider is objectively inferior to that of a human provider. Participants were resistant to medical AI even when the performance of AI providers was explicitly specified to be superior to that of human providers.

STUDY 4: UTILITIES BY PROVIDER IN CHOICE-BASED CONJOINT

Study 4 used a conjoint-based exercise to examine consumer resistance to medical AI from a different angle. This approach allowed us to estimate the implicit (dis)utility associated with an automated (vs. a human) provider, relative to other attributes of the providers. We estimated (dis)utilities from a series of choices that respondents made between human and automated providers that also varied on performance and cost.

Procedure

One hundred respondents from Amazon Mechanical Turk ($M_{age} = 33.9$, $SD = 10.6$; 47.0% female) participated in exchange for monetary compensation. The study consisted of a choice-based conjoint exercise conducted using Sawtooth software. Participants read that we were interested in their opinion about a new skin cancer screening promoted by the American Academy of Dermatology. The screening consisted of an analysis of high-resolution photographs (i.e., scans) of a person's body. Participants read that the analysis could be performed by different providers which varied on three attributes: (a) the *provider* performing the analysis, (b) the *accuracy* of the analysis, and (c) the *cost* of the analysis. The first attribute, provider performing the analysis, had two levels: human or robotic. The second attribute, accuracy of the analysis, had three levels: 81%, 84% and 87%. The third attribute, cost of the analysis, had three levels: \$20, \$40 and \$60. Participants read detailed descriptions of each attribute. Then, they read that they would be making seven consecutive choices among three options differing in terms of provider, accuracy, and cost. An example of this choice task followed (figure 3).

* Insert Figure 3 About Here *

Next, the choice-based conjoint exercise began. We optimized the experimental design following the settings recommended by the Sawtooth software. Specifically, participants completed seven consecutive choice tasks, with three choice options each. Stimuli in web appendix E.

Results and Discussion

We estimated attribute partworth utilities using the Hierarchical Bayesian (HB) module in Sawtooth. We specified accuracy and cost to be linear attributes. Thus, for these two

attributes, we estimated a single partworth parameter that captured the change in utility due to a unit change in the level of the attribute (e.g., each additional dollar provides a negative utility of 2.24). Provider was a dichotomous attribute, with human serving as the baseline. Thus, the estimated partworth corresponded to the (dis)utility associated with an automated provider relative to a human provider. As illustrated in table 3, the automated provider had a negative partworth utility, suggesting that, controlling for accuracy and cost, participants preferred a human to an automated provider.

* Insert Table 3 About Here *

The results of study 4 provided additional evidence for resistance to medical AI. In a choice-based conjoint task, in which respondents traded off provider, accuracy, and cost of healthcare, the partworth utility associated with an automated rather than a human provider was negative, suggesting that respondents were reluctant to choose an automated medical provider.

Together, the first four studies suggest that consumers are resistant to AI healthcare providers relative to human healthcare providers. Consumers less likely to utilize healthcare, exhibit lower reservation prices for healthcare, are less sensitive to differences in provider performance, and derive negative utility if a provider is automated rather than human. In the studies that follow (5-9), we focus on the role of uniqueness neglect as an underlying driver of this resistance to medical AI, empirically testing the explanatory power of this process account via mediation and moderation.

STUDY 5: SENSE OF UNIQUENESS AS MEASURED MODERATOR

Study 5 served as an initial test of our process account. Specifically, we examined whether resistance to medical AI is moderated by consumers' perception that their characteristics are unique and different from those of others (Simsek and Yalınçetin 2010). After measuring the extent to which consumers perceived themselves to be unique, participants expressed a relative preference between a human and an automated medical provider. If uniqueness neglect drives resistance to medical AI, consumers who consider themselves to be more different from others should exhibit a more pronounced resistance to an automated provider compared to consumers who consider themselves to be more similar to others.

Procedure

Two-hundred and eighty-six respondents ($M_{age} = 36.9$, $SD = 12.5$; 48.1% female) recruited from Amazon Mechanical Turk participated in the study in exchange for monetary compensation. Participants read that they would be participating in a series of studies. The first purportedly asked participants some background information, including three questions aimed to assess their perceived sense of uniqueness, adapted from the personal sense of uniqueness scale developed by Simsek and Yalınçetin (2010). Participants indicated the extent to which they agreed with the statements: "I feel that some of my characteristics are completely unique to me," "I think that the characteristics that make me up are completely different from others," and "I feel unique" (1 = *strongly disagree*, 7 = *strongly agree*). We averaged the scores on these three items to obtain a sense of uniqueness index, with higher scores indicating higher sense of uniqueness ($\alpha = .86$).

Participants were then presented with an ostensibly unrelated survey, described as a study about medical decision making. First, participants read the following extract about skin cancer and the importance of early detection:

Skin cancer is the most common cancer in the United States. 1 in 5 Americans will develop skin cancer in their lifetime. People of all colors and races can get skin cancer. There are many different types of skin cancer, including carcinoma and melanoma. Carcinoma is the most common form of skin cancer; melanoma is the deadliest. With early detection and proper treatment, both carcinoma and melanoma have a high cure rate. Skin screenings help people detect skin cancers early on.

Then, participants imagined that they had decided to get a skin cancer screening and would choose a dermatologist to perform this screening. Prior to choosing, participants reviewed accuracy information about two different dermatologists that could perform the screening. The two dermatologists had the same accuracy rate of 90%. Half of the participants were asked to choose between two human providers, dermatologist X and dermatologist Y. The other half of the participants were asked to choose between a human provider, dermatologist X, and an automated provider, dermatologist Y. Therefore, dermatologist X was always human, and dermatologist Y was either human or automated depending on the condition. Participants indicated which dermatologist they would choose to perform the screening (1 = *definitely dermatologist X*, 7 = *definitely dermatologist Y*). Stimuli in web appendix F.

Results and Discussion

We regressed provider preference on choice-set, sense of uniqueness (mean-centered), and their interaction. The analysis revealed a significant main effect of choice-set ($\beta = -1.382$, $t(282) = 11.12$, $p < .001$), and a significant two-way interaction between choice-set and sense of

uniqueness, ($\beta = -.363$, $t(282) = 3.92$, $p < .001$). There was no significant main effect of sense of uniqueness ($p = .97$). Because sense of uniqueness was a continuous measure, we explored the interaction further using the Johnson-Neyman floodlight technique (Spiller et al. 2013). The results revealed a negative and significant effect of choice-set on provider preference for levels of sense of uniqueness (mean-centered) greater than -2.4 ($\beta_{JN} = -.51$, $SE = .26$, $p = 0.050$; figure 4).

* Insert Figure 4 About Here *

These results provided correlational evidence for uniqueness neglect as a driver of resistance to utilize medical AI. The more participants perceived their characteristics to be unique and different from those of others, the more they exhibited resistance to an automated provider. These results were consistent with the notion that consumers believe that AI providers are not able to account for one's uniqueness as well as human providers. In the studies that follow, we sought converging evidence of uniqueness neglect via mediation (study 6) and moderation (studies 7-9).

STUDY 6: UNIQUENESS NEGLECT AS MEDIATOR

Study 6 tested the role of uniqueness neglect as driver of resistance to medical AI, vis-à-vis two potential alternative mechanisms: perceptions of responsibility (Promberger and Baron (2006), and ascriptions of human nature and human uniqueness (Haslam 2006).

Promberger and Baron (2006) argue that consumers do not want to feel responsible for difficult medical decisions, thus seek to shift responsibility onto someone else. Because humans are seen as capable of carrying responsibility, whereas machines are not, consumers are more

likely to follow the recommendation of a physician than the recommendation of a computer, because they think that they can offload responsibility onto a human but not onto a machine. Based on this logic, consumers may exhibit resistance to medical AI because they think they will not be able to shift responsibility to an automated provider, whereas they can shift responsibility to a human provider. In study 6 we measured perceptions of responsibility and pitted it against our uniqueness neglect account.

Consumers might also eschew medical AI because it lacks fundamental human characteristics that consumers deem important in a medical provider. Prior research differentiates between human nature (characteristics such as warmth and emotionality), and human uniqueness (characteristics such as conscientiousness and competence; Haslam 2006; Haslam et al. 2005). Absence of human nature might make the automated provider be perceived as less empathetic and caring than the human provider. Absence of human uniqueness might make the automated provider seem less capable and invested in the decision than the human provider. In study 6, we measured ascriptions of human nature and human uniqueness and pitted them against our uniqueness neglect account².

Procedure

Two-hundred and forty-three respondents ($M_{\text{age}} = 40.5$, $SD = 13.19$; 45.7% female) recruited from Amazon Mechanical Turk participated in exchange for monetary compensation. Participants read the same scenario used by Promberger and Baron (2006), with one critical difference: we specified that physician and computer had the same accuracy rate. Specifically,

² We thank the Associate Editor for this suggestion.

participants imagined that they were medical patients and that they were considering undergoing coronary by-pass surgery. They read that a physician had conducted an examination and collected their health-relevant data. The results of this examination would be analyzed, for the purpose of coming up with a treatment recommendation (i.e., undergoing by-pass surgery or not), either by a physician or by a computer, specified to be equally accurate. Specifically, participants read the following:

Imagine that you are considering undergoing coronary bypass surgery. This is how this particular decision (to have coronary bypass surgery) goes. Imagine that a physician has examined you and measured some health data (which we will tell you in detail later), and now you have to decide whether you want to have coronary bypass surgery or not. To make this decision, you will get a recommendation—about whether you should have this operation or not. There are two possibilities as to where this recommendation comes from. Either a physician evaluated the results of your examination and used his/her judgment and experience to compare your case to patients who faced the same decision and the results of their treatment. Or a computer program evaluated the results of your examination and used an algorithm to compare your case with patients who faced the same decision and the results of their treatment. In the past, the physician and the computer showed the same accuracy in making recommendations.

These are the health-data that a physician has measured: your blood pressure, your ECG, your LDL blood cholesterol, and the blockage in your coronary artery. These are the results. Blood pressure can have the values low, normal, high, very high. Yours is: high. Electrocardiogram can have the values very good, normal, poor, very poor. Yours is: normal. LDL blood cholesterol levels can have the values low, normal, high, very high. Yours is: very high. Coronary artery examination can have the values no blockage, light blockage, medium blockage, severe blockage. Yours is: no blockage.

Between-subjects, we then manipulated the provider that generated the recommendation.

Specifically, participants read that the analysis of their health-data was carried out either by a human provider (i.e., a physician) or by an automated provider (i.e., a computer).

[*Human provider*] In this case, it is a physician who looks at your examination, analyses your data, and recommends: You should have the operation.

[*Automated provider*] In this case, it is a computer that processes your examination, analyses your data, and recommends: You should have the operation.

Participants indicated the extent to which they would follow the recommendation (1 =

definitely not follow, 7 = *definitely follow*). We then measured uniqueness neglect, perceptions of responsibility, and humanness ascriptions.

We measured uniqueness neglect by asking participants to rate the extent to which, in the situation described, they would be concerned that, when analyzing the results of their examination, the recommender “would not recognize the uniqueness of your medical condition,” “would not consider your unique circumstances,” and “would not tailor the recommendation to your unique case” (1 = *not at all*, 7 = *very much*; $\alpha = 0.96$).

We measured perceptions of responsibility by asking participants to rate the extent to which, in the situation described, they would feel responsible for the outcome of the surgery (1 = *not at all*, 7 = *very much*). We also utilized two items from Promberger and Baron (2006, study 1): “If you followed the recommendation, how much would you feel responsible for the decision that you made?” and “If you did not follow the recommendation, how much would you feel responsible for the decision that you made (1 = *not at all*, 7 = *very much*). In keeping with Promberger and Baron (2006), we computed a variable capturing *responsibility reduction*: this variable was simply the difference between how responsible participants said they would feel had they followed the recommendation and how responsible participants said they would feel had they not followed the recommendation.

Finally, we measured ascriptions of human nature and human uniqueness by asking participants to rate the extent to which, in the situation described, they would be concerned that, when analyzing the results of their examination, the recommender was warm, emotional, active (characteristics distinctive of human nature; $\alpha = .79$), intelligent, competent and conscientious (characteristics distinctive of human uniqueness; $\alpha = .90$; 1 = *not at all*, 7 = *very much*). Stimuli in web appendix G.

Results and Discussion

T-tests on propensity to follow the recommendation revealed that participants were more reluctant to follow the recommendation when the provider was automated ($M = 3.57$; $SD = 1.83$) than human ($M = 4.61$, $SD = 1.81$; $t(241) = 4.46$, $p < .001$, $d = .076$). Additionally, participants reported higher concern of uniqueness neglect when the provider was automated ($M = 5.27$, $SD = 1.72$) than human ($M = 3.71$; $SD = 1.79$; $t(241) = 6.95$, $p < .001$, $d = .167$). Participants also reported lower responsibility reduction when the provider was automated ($M = -0.18$, $SD = 1.46$) than human ($M = -0.62$; $SD = 1.55$; $t(240) = 2.28$, $p = .02$, $d = .021$ [1 missing data]). Perceptions of responsibility ($p = .48$), ascriptions of human nature ($p = .08$) and ascriptions of human uniqueness were not significantly affected by the type of provider ($p = .38$).

Mediation. We conducted a mediation analysis to test whether uniqueness neglect mediated the observed differences in propensity to follow the recommendation. Along with uniqueness neglect, we entered, as simultaneous mediators, the two variables that were significantly affected by provider: responsibility reduction and human nature ascriptions (for a conservative test, we included human nature ascriptions even though the effect was marginal). We examined confidence intervals (CI) using 5,000 bootstrap iterations (Hayes 2013) coding provider as 0 when human and 1 when automated. The direct effect of provider was not significant ($-.16$, 95% CI $[-.61, .30]$). The indirect effect through uniqueness neglect was significant ($-.83$, 95% CI $[-1.16, -.57]$), and the direction of the effects confirmed that an automated provider led to greater uniqueness neglect, which in turn contributed to higher reluctance to follow the recommendation. The indirect effects through responsibility reduction (-

.03, 95% CI [-.11, .02] and human nature ascriptions (-.03, 95% CI [-.14, .01) were not significant, indicating that none of these variables accounted for the observed difference in propensity to follow the recommendation. A mediation model that did not include human nature ascriptions as a mediator showed the same results, as did other specifications of the mediation model (details in web appendix H).

These results provided evidence for the hypothesized role of uniqueness neglect as driver of resistance to medical AI. Participants were less likely to follow a treatment recommendation offered by an automated rather than a human provider, and uniqueness neglect—concern that an automated provider cannot account for a person’s unique symptoms, circumstances, and characteristics—mediated the effect. Furthermore, we did not find evidence that resistance to medical AI was driven by perceptions of responsibility, or by ascriptions of humanness, suggesting that the effect was neither due to perceived differences in the providers’ ability to shoulder the responsibility for an outcome, nor to perceived differences in humanness. In study 7 and 8 we employed similar experimental procedures and sought to provide converging evidence for the role of uniqueness neglect in driving resistance to medical AI via moderation.

STUDY 7: PERSONALIZED CARE CURBS UNIQUENESS NEGLECT

Study 7 tested a theoretically driven and practically relevant moderator that should curb resistance to medical AI: personalization of healthcare. If resistance to medical AI manifests because consumers view themselves as unique and different from others, yet believe that AI treats everyone the same way and neglect a person’s unique characteristics, circumstances, and symptoms, then framing care offered by AI providers as personalized should assuage uniqueness

neglect and curb resistance to medical AI. In study 7, we manipulated the degree of personalization offered by a healthcare provider and tested its effect on resistance to medical AI.

Procedure

Two-hundred and ninety-four respondents ($M_{\text{age}} = 39.2$, $SD = 12.4$; 43.7% female) recruited from Amazon Mechanical Turk completed this study in exchange for money. Respondents were randomly assigned to one of the conditions in a 2 (provider: human versus automated) x 2 (personalization: not salient vs. salient) between-subjects design.

Participants read the same medical scenario as in study 6. Between subjects, we manipulated whether a patient's health-related data would be analyzed by a human provider (i.e., a physician) or by an automated provider (i.e., a computer), specified to have the same accuracy. Also between-subjects, we manipulated the salience of the provider's ability to personalize and tailor the analysis, and consequently the recommendation, to a patient's unique characteristics, or not. Specifically, participants read:

[Personalization not salient] A [physician/computer] has analyzed these data. The [physician/computer] recommends: You should have the operation.

[Personalization salient] A [physician/computer] has analyzed these data and, when analyzing these data [they/it] has ensured that the analysis was personalized and tailored to your unique characteristics. The [physician/computer] recommends: You should have the operation.

Participants indicated how likely they would be to follow the recommendation (1 = *definitely not follow*, 7 = *definitely follow*). Stimuli in web appendix I.

Results and Discussion

A 2x2 ANOVA on likelihood to follow the recommendation revealed a significant main effect of provider ($F(1, 290) = 10.99, p = .001, \eta_p^2 = .037$), a significant main effect of personalization ($F(1, 290) = 5.72, p = .017, \eta_p^2 = .019$), and a significant 2-way interaction between the two factors ($F(1, 290) = 6.72, p = .01, \eta_p^2 = .023$). Planned contrast revealed that when the recommendation was not explicitly personalized, participants were more reluctant to follow it if the provider was automated ($M = 3.61, SD = 1.89$) than human ($M = 4.88; SD = 1.78, p < .001$). As predicted, however, there was no significant difference in propensity to follow the recommendation when it was explicitly personalized ($M_{automated} = 4.68, SD_{automated} = 1.75, M_{human} = 4.83; SD_{human} = 1.94, p = .61$). These results appear to suggest that participants assumed that care delivered by the human provider was personalized, as propensity to follow the recommendation was the same irrespective of whether the recommendation was explicitly personalized or not ($p = .88$). By contrast, participants appeared to believe that care delivered by the automated provider was personalized only if explicitly told so, as propensity to follow the recommendation was higher when the recommendation was explicitly personalized than when it was not ($p < .001$; figure 5).

* Insert Figure 5 About Here *

The results of study 7 provided additional evidence that uniqueness neglect drives resistance to medical AI. When there was no mention of the provider's ability to personalize care, consistent with our previous results, participants were more reluctant to follow a medical recommendation when the provider was automated than human. When human and automated providers were described as capable of tailoring care to each patient's unique characteristics, however, participants were as likely to follow the recommendation of AI as that of a human

provider. These results are consistent with the idea that consumers believe that AI providers are not able to account for one's unique symptoms, circumstances and characteristics as well as human providers. Personalized medicine appears to curb resistance to medical AI because it reassures consumers that care is tailored to their own unique characteristics, thus assuaging uniqueness neglect.

STUDY 8: RESISTANCE TO MEDICAL AI DOES NOT MANIFEST WHEN DECIDING FOR AVERAGE OTHERS

Study 8 further tested if uniqueness neglect drives resistance to medical AI, by manipulating the perceived uniqueness of the recipient of medical care. We reasoned that, if resistance to medical AI is driven by perceptions of the recipient being unique and different from others, it should not manifest when the recipient of care is not perceived to be unique. We tested this hypothesis by having participants make a decision either for themselves (i.e., high uniqueness), for an individuated other (i.e., high uniqueness), or for an average person (i.e., low uniqueness). We expected resistance to medical AI to emerge when making a decision for the self and for a unique other, but not for an average person.

Procedure

Four-hundred and one respondents ($M_{\text{age}} = 35.0$, $SD = 11.4$; 46.9% female) recruited from Amazon Mechanical Turk completed this study in exchange for money.

Participants were randomly assigned to one of the conditions in a 2 (provider: human versus automated) x 3 (recipient of care: self, average other, unique other) between-subjects design.

Participants read the same medical scenario as in studies 6 and 7. Between subjects, we manipulated whether a patient's health-related data would be analyzed by a human provider (i.e., a physician) or by an automated provider (i.e., a computer), specified to have the same accuracy. Also between-subjects, we manipulated whether participants made a decision for themselves, for an individuated other person, or an average other person. Specifically, participants read:

[Recipient of care: self] Please imagine the following scenario. The following scenario will be about you as a medical patient.

[Recipient of care: individuated other] Please imagine the following scenario. The following scenario will be about a person called Janice as a medical patient. These are some details about Janice. Janice lives in Kansas, is 38 years old, and is married with two kids. Janice loves watching football and ice skating. Janice has a mild peanut allergy. Janice has a family history of colon cancer.

[Recipient of care: average other person] Please imagine the following scenario. The following scenario will be about an average person as a medical patient.

These descriptions were calibrated to manipulate the perceived uniqueness of the recipient of care. Specifically, we expected that respondents would perceive both the self and Janice to be unique and different from an average patient. We validated this calibration by asking one-hundred and twelve respondents from Amazon Mechanical Turk to rate the perceived uniqueness, as medical patients, of the three profiles of the recipient of care in the scenario utilized in the main study (i.e., self, average person, unique other person; "As a medical patient in these circumstances, how unique do you think you are / Janice is / an average person is?" 1 = *not at all unique*, 7 = *very unique*). Indeed, the self and individuated other were rated as equally unique ($M_{Self} = 4.68$ $SD = 1.72$, $M_{Individuated Other} = 4.34$ $SD = 1.73$, $p = .40$) and more unique than an average other ($M_{Average other} = 3.29$, $SD = 1.71$, vs. self, $p = .001$; vs. individuated other, $p =$

.01). After reading the scenario, depending on the condition, participants indicated the extent to which *they* would follow the recommendation, or the extent to which they thought that *Janice* should follow the recommendation, or the extent to which they thought that *an average person* should follow the recommendation (1 = *definitely not follow*, 7 = *definitely follow*). Stimuli in web appendix J.

Results and Discussion

A 2x3 ANOVA revealed a significant main effect of provider ($F(1, 395) = 34.74, p < .001, \eta_p^2 = 0.08$), a significant main effect of recipient of care ($F(2, 395) = 10.94, p < .001, \eta_p^2 = 0.05$) and a significant interaction ($F(2, 395) = 5.54, p = .004, \eta_p^2 = 0.03$). Follow-up contrasts revealed that when the provider was human, participants' propensity to follow the recommendation was the same irrespective of whether the recipient of care was the self ($M_{Self} = 4.80, SD = 1.82$), an individuated other ($M_{Individuated\ other} = 4.67, SD = 1.77$), or an average other ($M_{Average\ other} = 5.02, SD = 1.68$; self vs. individuated other, $p = .64, p = .99$; self vs. average other, $p = .47$, individuated other vs. average other, $p = .23$). In contrast, when the provider was automated, participants were as reluctant to follow the automated provider's recommendation when choosing for themselves ($M_{Self} = 3.17, SD = 1.67$) as they were when choosing for an individuated other ($M_{Individuated\ other} = 3.55, SD = 1.62; p = .20$). However, when choosing for an average other person, participants were less reluctant to follow the automated provider's recommendation than both when choosing for themselves ($M_{Average\ other} = 4.76, SD = 1.62; p < .001$) and when choosing for an individuated other ($p < .001$). Similarly, when choosing for themselves and for an individuated other, participants were more reluctant to follow the

recommendation of an automated provider than the recommendation of a human provider ($p < .001$), whereas when choosing for an average other, participants' propensity to follow the recommendation was the same irrespective of the type of provider ($p = .38$; figure 6).

* Insert Figure 6 About Here *

These results provided converging evidence for uniqueness neglect as the psychological mechanism underlying resistance to medical AI. Participants exhibited resistance to medical AI when deciding for themselves, as well as when deciding for another person perceived to be unique. But, they did not exhibit resistance to medical AI when deciding for a person who was not perceived to be unique (i.e., an average person). Thus, these results supported the notion that resistance to medical AI was driven by the perceived inability of automated providers to satisfy unique needs of patients.

STUDY 9: MODERATION BY AI-ROLE (SUPPORT VS. REPLACEMENT)

Study 9 explored which stage of healthcare delivery evokes resistance to medical AI. We predicted that consumers would be averse to utilizing healthcare provided by medical AI when the automated provider replaces a physician as the ultimate decision maker. By contrast, we predicted that consumers would be receptive to medical AI if the automated provider simply supports a physician who retains the role of ultimate decision maker and thus can consider the unique needs of the patient. To benchmark consumers' responses to each of these two forms of healthcare delivery, we compared both conditions to one in which a physician is the sole decision maker.

Procedure

One-hundred and ninety-seven respondents ($M_{age} = 39.4$, $SD = 124$; 54.3% female) recruited from Amazon Mechanical Turk completed this study in exchange for money.

Participants read about a new skin cancer screening service.

Imagine that the American Academy of Dermatology is promoting a new type of preventive screening for skin cancer that you can take for a \$30 fee. This initiative is meant to foster detection of skin cancer in its earliest, most treatable stage, to reduce the incidence of the skin cancer, and raise awareness of effective skin cancer prevention techniques.

This is how this particular screening would work. You will go to see your primary care physician. Then, a nurse will take you to a private room, where she will take high-resolution photographs (scans) of your body. Screenings typically include an exam of scalp, face, mouth, hands, feet, trunk and extremities, eyes and eyelids, ears, fingers, toes and toenails, but you can ask for more (or fewer) areas to be checked.

Participants were then randomly assigned to one of three conditions that manipulated the role of medical AI as follows. In the *baseline* condition, the human provider was described as the sole decision maker:

Then, the scans of your body will be sent to a dermatologist trained to develop pattern recognition skills. This dermatologist is a physician trained to distinguish between cancerous and non-cancerous skin conditions. Training was made possible through learning the differences between cancerous and non-cancerous skin conditions using an extensive dataset of images. In the past, this dermatologist was able to achieve 87% sensitivity in picking up cancerous cells from body scans. This dermatologist will analyze the scans of your body, and send the results back to your doctor within one week from the screening. You will have no interaction with this dermatologist. Your doctor will call you to let you know these recommendations.

In the *AI as replacement* condition, the automated provider was described as the sole decision maker:

Then, the scans of your body will be sent to an artificial intelligence (AI) dermatologist trained to develop pattern recognition skills. This AI dermatologist is an algorithm trained to distinguish between cancerous and non-cancerous skin conditions. Training was made possible through learning the differences between cancerous and non-

cancerous skin conditions using an extensive dataset of images. In the past, this AI dermatologist was able to achieve 87% sensitivity in picking up cancerous cells from body scans. This AI dermatologist will analyze the scans of your body, and send the results back to your doctor within one week from the screening. You will have no interaction with this AI dermatologist. Your doctor will call you to let you know these recommendations.

In the *AI as support* condition, the automated provider was described as providing an input to support a doctor who would ultimately make the final decision:

Then, the scans of your body will be sent to an artificial intelligence (AI) dermatologist trained to develop pattern recognition skills. This dermatologist is an algorithm trained to distinguish between cancerous and non-cancerous skin conditions. Training was made possible through learning the differences between cancerous and non-cancerous skin conditions using an extensive dataset of images. In the past, this AI dermatologist was able to achieve 87% sensitivity in picking up cancerous cells from body scans. This AI dermatologist will analyze the scans of your body, and send the results back to your doctor within one week from the screening. You will have no interaction with this AI dermatologist. Even though the AI dermatologist will analyze the scans, your doctor will make the final call and use the AI dermatologist's recommendations simply as input and support to the final decision. Your doctor will call you to let you know these recommendations.

Participants indicated how likely they would be to utilize this service (1 = *very unlikely*, 7 = *very likely*). Participants then indicated their age and gender. Stimuli in web appendix K.

Results and Discussion

A one-way ANOVA on likelihood to utilize the service revealed a significant main effect of condition ($F(2, 194) = 5.19, p = .006, d = .11$). Follow-up contrasts revealed that participants were more resistant to utilize medical AI when it would replace the human provider ($M_{\text{AI replacement}} = 4.88, SD = 1.74$) than in either of the other two conditions: a human provider making the decision ($M_{\text{Human}} = 5.57, SD = 1.54; t(194) = 2.42, p = .017$), and medical AI providing input to support a human provider ($M_{\text{AI support}} = 5.75, SD = 1.62; t(194) = 3.06, p = .003$). The pattern of

results was the same when comparing the AI as replacement condition with the other two conditions (with weights, -2, +1, +1), $t(194) = 3.16, p = .002$. Likelihood to utilize the service was the same when a human provider would make the focal decision and when an automated provider supported a human provider who would make the final decision ($t(194) < 1, p = .53$).

These results showed that consumers were resistant to use medical AI when an automated provider replaced a human provider as the decision maker. By contrast, resistance to medical AI did not emerge when an automated provider only supported a human provider who made the decision. Indeed, participants were as likely to utilize the medical service in the latter condition as when no medical AI was involved in any stage of care delivery. Thus, the results of study 9 speak to the scope of resistance to medical AI, suggesting that consumers might be open to utilize health care when medical AI provides support to, but does not replace, a physician.

GENERAL DISCUSSION

Across a variety of medical decisions ranging from prevention, to diagnosis, to treatment, we document a robust resistance to use medical care delivered by AI providers rather than comparable human providers. In studies 1 and 2, participants were less likely to utilize a stress assessment, and would pay less for that assessment, when the medical care provider was automated than human, even though provider accuracy was the same. In studies 3A-3C, participants exhibited weaker preference for a provider that offered clearly superior performance when such provider was automated rather than human. The choice-based conjoint in study 4 showed a negative partworth utility for an automated provider, suggesting that, controlling for accuracy and cost, participants preferred a human to an automated provider.

We identify uniqueness neglect as a psychological driver of resistance to medical AI. Consumers believe that AI providers are unable to take into account the uniqueness of their case to the same extent as a human provider. In study 5, resistance to medical AI was moderated by the extent to which a consumer perceived herself to be unique. Study 6 showed that concern of uniqueness neglect mediated resistance to medical AI. Studies 7 and 8 provided further converging evidence for uniqueness neglect via moderation. Resistance to medical AI was reduced when uniqueness concerns were assuaged, as in the case of personalized medicine (study 7), or did not apply, as when making decisions for an average person (study 8). Finally, study 9 showed that resistance to medical AI was eliminated when the automated provider merely provided input to a physician's decision (i.e., supported rather than replaced a human provider).

Contributions

Our research makes three main contributions to the literatures on clinical versus actuarial judgment and medical decision making. First, we extend the literature on clinical versus statistical judgments. Since Meehl's seminal contribution (1954), substantial research has shown that people prefer to make decisions relying on their own intuition, or on the intuition of another human being, rather than on a statistical model, even though statistical models usually outperform human intuition (Dawes 1979). This work has primarily examined non-medical contexts, such as predictions of employee performance (Kuncel et al. 2013), market demand (Sanders and Manrodt 2003), fraud (Boatsman et al. 1997), crime (Kleinberg et al. 2017; Wormith and Goldstone 1984), and consumer preferences (Sinha and Swearingen 2001; Yeomans et al. 2017). We extend the scope of this literature to the medical domain. This is an

important contribution because medical decision making differs from decision making in other consumer contexts in important ways (Botti and Iyengar 2006; Kahn and Baron 1995; Kahn et al. 1997). Medical decisions are unfamiliar, rife with uncertainty (Epstein and Peters 2009), and have the potential for highly consequential and life-threatening outcomes (Botti et al. 2009). Thus, it remained to be tested whether findings from other domains would also apply to the medical domain.

Second, we identify a novel psychological driver of consumer reluctance to rely on statistical judgments: uniqueness neglect. Although our examination falls squarely in the medical domain, uniqueness neglect should drive resistance to AI in any domain in which consumers view themselves as unique, and their uniqueness is important when making decisions. This contribution is significant, because evidence of what drives reticence to rely on statistical judgment is still very limited. Prior research has in fact mainly focused on documenting the phenomenon across different domains, rather than identifying its psychological drivers. Several hypotheses have been advanced, including the possibility that people feel dehumanized when interacting with automated agents (Dawes 1979), that algorithms lack empathy and treat individuals as objects (Grove and Meehl 1996), that algorithms are unable to learn and improve through experience (Dawes 1979; Highhouse 2008) or process qualitative and unstructured information (Grove and Meehl, 1996), and a belief that it is unethical to entrust important decisions to algorithms (Dawes 1979). Yet, these hypotheses have for the most part remained speculative (for exceptions see Dietvorst et. al 2014; Castelo et. al 2018). By empirically documenting the role of uniqueness neglect as underlying resistance to medical AI, we add an important contribution toward explaining people's general reticence to rely on statistical judgments. In addition, our work provides initial evidentiary support, as well as insight into the

process, for the so-called broken leg hypothesis, which has long been used as a speculative explanation of people's resistance to rely on statistical judgment.

Third, our findings advance our understanding of medical decision making (Bolton et al. 2007; Kahn and Baron 1995; Kahn et al. 1997; Keller and Lehmann 2008). Prior research on medical decision making has explored how factors such as degree of desired autonomy (Botti et al. 2009), degree of involvement (Arora and McHorney 2000), and employment of decision aids (Ubel 2002) influence medical decisions. We identify a previously unexplored psychological mechanism that may drive decisions in the health domain—uniqueness neglect. Although our empirical investigation is limited to situations in which consumers face an automated medical care provider, we believe that uniqueness neglect is an important psychological mechanism that may influence medical decisions more broadly, and not just with respect to medical AI.

Practical Implications

The healthcare industry is experiencing tremendous interest in the use of AI to deliver efficacious and cost-effective care at scale (Das 2016). AI enables innovative solutions that have incredible potential across the spectrum of health stakeholders. Virtual nursing assistants available to interact with patients 24/7 could save the healthcare industry \$20 billion annually (Marr 2018). AI-based telemedicine could provide primary care support to remote areas without easy access to healthcare (Siwicki 2018). Our research points to a critical barrier that consumer-facing AI-based technological advancements will need to overcome to gain acceptance and diffusion: consumers might be reluctant to adopt medical AI because they believe it unable to

account for the unique facets of a person's case. Changing this belief will be fundamental to harness the full potential of medical AI to benefit our society in the future.

By identifying uniqueness neglect as a psychological driver of resistance to medical AI, our research yields insight into the kinds of interventions that might increase acceptance of medical AI. Study 7 indicated that resistance to medical AI is reduced when care is explicitly tailored to one's unique characteristics. This suggests that providers should actions aimed to increase the perceived personalization of medical care delivered by AI. This could be achieved by increasing the amount of information collected about each consumer, for instance, or by including cues suggesting personalization when providing patients with feedback and recommendations (e.g., "based on your unique profile"). The results of study 9 indicate that consumers are not resistant to medical AI when a physician remains in charge of the ultimate decision. This suggests that including the possibility of having a physician corroborate the recommendation of an AI provider should make consumers more receptive to AI providers.

Limitations and Future Research

Despite the robustness of the phenomenon documented and the converging evidence for the psychological account proposed, our research has limitations that offer several opportunities for future research. We focus our discussion on four directions that we believe offer particularly fruitful avenues to extend the current work.

First, our findings provide evidence for consumer resistance to medical AI across several medical domains, types of decisions, and preference elicitation procedures, but our research constitutes only the first step towards understanding this complex phenomenon. Several other

factors may sway resistance to medical AI. In a follow-up study ($N = 127$; $M_{\text{age}} = 38.1$, $SD = 12.9$; 46.5% female) we compared relative preference for a human vs. an AI provider as a function of the stakes of the medical decision. Specifically, we held the domain constant (i.e., a preventive service offering screenings of eye conditions) and manipulated the stakes via the severity of the medical condition (i.e., high stakes in case of diagnosis of macular degeneration vs. low stakes in case of diagnosis of dry eye; details in web appendix L). Participants in all cases preferred a human to an automated provider, but exhibited stronger reluctance to utilize an automated provider when the stakes were high ($M_{\text{High stakes}} = 2.42$, $SD = 1.53$) than when the stakes were low ($M_{\text{Low stakes}} = 3.19$, $SD = 1.74$, $t(125) = 2.64$, $p < .001$), suggesting that resistance to medical AI is less pronounced for less consequential medical decisions.

Second, more research is needed to map out the theoretical and practical boundaries of this complex phenomenon. In terms of mechanism, our research offers converging evidence for uniqueness neglect as a driver of resistance to medical AI. However, we believe that consumer resistance to medical AI is a complex phenomenon that is likely multiply determined (Inman et al. 2018; Kirmani 2015). As outlined above, several mechanisms may lead consumers to rely on clinical rather than statistical judgments including feelings of dehumanization and morality concerns. These mechanisms, along with unidentified other processes, may be similarly influential in the medical domain. Investigating the role that other psychological mechanisms play in influencing consumer preferences for medical AI is a fertile avenue for future research.

Third, just like these other drivers lead consumers to prefer clinical to statistical judgments in non-medical domains, so too may uniqueness neglect influence preferences in other domains in which resistance to statistical judgments has been observed. In the medical domain, uniqueness neglect manifests as a concern that an automated provider will neglect one's health-

related uniqueness—the constellation of a person’s unique biological, psychological, and behavioral characteristics. In other domains, uniqueness neglect might manifest as a concern that other characteristics that are relevant to one’s sense of uniqueness in that domain may be neglected.

We tested this conjecture in a follow-up study ($N = 92$; $M_{\text{age}} = 40.2$, $SD = 13.7$; 42.4% female). We selected eight domains where AI-based advice is currently available to consumers: fashion, financial investments, health, home décor, legal matters, movies, music, and restaurants. Participants read a brief description about each domain and the kind of advice that they would receive for that domain (order of domains randomized). Then, for each domain, participants made two judgments. They expressed a relative preference for receiving advice from an automated agent (i.e., a well-trained algorithm) versus a human agent (i.e., a well-trained human expert; 1 = *definitely the algorithm*, 7 = *definitely the human expert*), and then reported the extent to which they believed that their needs, preferences, and goals in that domain were unique to them (1 = *not at all*, 7 = *very much*). In all of the domains tested, there was a positive correlation between preference for human advice and self-uniqueness (correlations ranging from 0.25 to 0.46, all $ps < .016$; no difference between correlations, all $ps > .12$; details in web appendix M). These results offered a preliminary indication that uniqueness neglect may play a role in other domains in which resistance to statistical judgments has been documented. Further exploring whether uniqueness neglect applies to domains other than the medical one offers another fruitful avenue for future research.

Additionally, a particularly fruitful avenue for future research would be to explore ways to curb, consumer resistance to medical AI. Dietvorst, Simmons and Massey (2016) showed that people are more likely to rely on algorithms when they can even slightly modify them. This

intervention might overcome uniqueness neglect by giving people a sense of personalization, and thus might curb resistance to medical AI. Indeed, participants in study 9 were amenable to the inclusion of medical AI in their healthcare when a physician had the opportunity to check and correct the analysis of the AI.

Finally, we call on future research to explore when consumers may exhibit a preference for medical AI over human providers. For example, AI may be preferred to human providers when treatment requires the disclosure of stigmatized information (e.g., obesity or sexually transmitted diseases). Similarly, AI might be preferred to human providers when patients hold goals other than diagnosis or treatment accuracy, such as affordability or easy access to care. Furthermore, consumers might sometimes prefer AI to human providers when healthcare is framed as catering to unique needs. In this vein, it would be interesting to explore whether evidence that AI providers do deliver personal and individualized healthcare (e.g., patient-generated reviews, patient-generated word-of-mouth) can increase consumer receptivity toward medical AI. Given the speed at which AI-based healthcare technologies are being developed and deployed, these questions offer impactful and timely opportunities to better understand how to improve consumer wellbeing in the near and distant future.

Data Collection

Empirical illustration in the introduction: people from the Amazon Mechanical Turk panel participated in the study. All participants were from the US. Chiara Longoni collected and analyzed the data collected in this study and reported the results to the other two authors. Study 1: undergraduate students at Boston University and undergraduate students at New York University completed this study. The data from study 1 was collected by research assistants blind to the research hypotheses under the supervision of Chiara Longoni at Boston University and Andrea Bonezzi at New York University. Chiara Longoni analyzed the data collected in these studies and reported the results to the other two authors. Studies 2, 3A, 3B, 3C, 6, 7, 8, 9, 10, 11: people from the Amazon Mechanical Turk panel participated in these studies. All participants were from the US. Chiara Longoni collected and analyzed the data collected in these studies and reported the results to the other two authors. Studies 4-5: people from the Amazon Mechanical Turk panel participated in these studies. All participants were from the US. Andrea Bonezzi collected and analyzed the data and reported the results to the other two authors. Data was collected between October 2016 and February 2019.

REFERENCES

- Abràmoff, Michael D., Philip T. Lavin, Michele Birch, Nilay Shah and James C. Folk, (2018), “Pivotal Trial of an Autonomous AI-based Diagnostic System for Detection of Diabetic Retinopathy in Primary Care Offices,” *Nature Partners Journals Digital Medicine*, 1, 39.
- American College Health Association (2015). National College Health Assessment. Accessed February 19, 2018 http://www.acha-ncha.org/reports_ACHA-NCHAIic.html
- Arora, Neeraj K. and Colleen A. McHorney (2000), “Patient Preferences for Medical Decision-making: Who Really Wants to Participate?” *Medical Care*, 38(3), 335–41.
- Baumeister, Roy F. (1998), “The Self.” In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The Handbook of Social Psychology* (Vol. 3, pp. 680– 740). Boston: McGraw-Hill.
- Brown, Jonathon D. (1998), *The Self*, Boston: McGraw-Hill.
- Blanton, Hart, Danny Axsom, Kimberly P. McClive and Simani Price (2001), “Pessimistic Bias in Comparative Evaluations: A Case of Perceived Vulnerability to the Effects of Negative Life Events,” *Personality and Social Psychology Bulletin*, 27(12), 1627–36.
- Boatsman, James, Cindy Moeckel and Buck K.W. Pei (1997), “The Effects of Decision Consequences on Auditors' Reliance on Decision Aids in Audit Planning,” *Organizational Behavior and Human Decision Processes*, 71(2), 211-47.
- Böckenholt, Ulf and Elke Weber (1992), “Use of Formal Methods in Medical Decision Making: A Survey and Analysis,” *Clinical Use of Formal Decision Methods*, 12(4), 298-306.
- Bolton, Lisa E., Americus Reed II, Kevin G. Volpp, and Katrina Armstrong (2007), “How Does Drug and Supplement Marketing Affect a Healthy Lifestyle?” *Journal of Consumer Research*, 34(5), 713-26.
- Botti, Simona and Sheena S. Iyengar (2006), “The Dark Side of Choice: When Choice Impairs

- Social Welfare,” *Journal of Public Policy & Marketing*, 25(1), 24-38.
- Botti, Simona, Kristina Orfali and Sheena S. Iyengar (2009), “Tragic Choices: Autonomy and Emotional Responses to Medical Decisions,” *Journal of Consumer Research*, 36(10), 337-52.
- Bresnick, Jennifer (2017), “Artificial Intelligence in Healthcare Market to See 40% CAGR Surge,” available at <https://healthitanalytics.com>, July 24th
- Brewer, Marylann B. (1991), “The Social Self: On Being the Same and Different at the Same Time,” *Personality and Social Psychology Bulletin*, 17(5), 475-82.
- Cañas José, Quesada José F., Antolí A, Inmaculada Fajardo (2003), “Cognitive Flexibility and Adaptability to Environmental Changes in Dynamic Complex Problem Solving Tasks,” *Ergonomics*, 46, 482–501.
- Castelo, Noah, Marteen Bos and Don Lehman (2018), “Consumer adoption of algorithms that blur the line between human and machine,” working paper.
- Center for Collegiate Mental Health Annual Survey (2016). Accessed on February 19th, 2018 at https://sites.psu.edu/ccmh/files/2017/01/2016-Annual-Report-FINAL_2016_01_09-1gc2hj6.pdf
- Centers for Medicare and Medicaid Services. Accessed on February 23rd, 2019, at <https://www.cms.gov/Medicare/Quality-Initiatives-Patient-Assessment-Instruments/HospitalQualityInits/HospitalHCAHPS.html>
- Das, Reenita (2016), “Five Technologies That Will Disrupt Healthcare By 2020,” *Forbes*, March 30th
- Dawes, Robyn M. (1979), “The Robust Beauty of Improper Linear Models in Decision Making,” *American Psychologist*, 34(7), 571-82.

- Dawes, Robyn, M. David and Paul E. Meehl (1989), "Clinical Versus Actuarial Judgment," *Science*, 243, 1668-74.
- Dietvorst, Berkeley J. (2018), "People Reject (Superior) Algorithms Because They Compare Them to Counter-normative Reference Points," working paper
- 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*.
- 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*, 144(1), 114-26.
- Donnelly, Laura (2017), "Forget Your GP, Robots Will Soon Be Able to Diagnose More Accurately Than Almost Any Doctor," *The Telegraph*, March 7th.
- Dunning, David (1993), "Words to Live By: The Self and Definitions of Social Concepts and Categories," In J. M. Suls (Ed.), *Psychological Perspectives on the Self* (Vol. 4, pp. 99–126), Hillsdale, NJ: Erlbaum.
- Eastwood, Joseph, Brent Snook and Kirk Luther (2012), "What People Want From Their Professionals: Attitudes Toward Decision-Making Strategies," *Journal of Behavioral Decision Making*, 25, 458-68.
- Epstein, Seymour (1990). Cognitive-Experiential Self-Theory, In L. A. Pervin (Ed.), *Handbook of Personality: Theory and Research* (pp. 165–192). New York: Guilford Press.
- Epstein, Ronald M. and Ellen Peters (2009), "Beyond Information: Exploring Patients' Preferences," *Journal of the American Medical Association*, 302(2), 195-97.
- Esteva, Andre, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau

- and Sebastian Thrun (2017), “Dermatologist-level Classification of Skin Cancer With Deep Neural Networks, *Nature*, 542, 115-18.
- Gallagher, James (2017), “Artificial Intelligence As Good As Cancer Doctors,” *BBC News Website*, January 26th
- Granulo, Armin, Christoph Fuchs and Stefano Puntoni (2018), “Psychological Reactions to Human Versus Robotic Job Replacement,” working paper.
- Greenwald, Adam G. (1980), “The Totalitarian Ego: Fabrication and Revision of Personal History,” *American Psychologist*, 35(7), 603–18.
- Grove, William M., & Paul E. Meehl (1996), “Comparative Efficiency of Informal (Subjective, Impressionistic) and Formal (Mechanical, Algorithmic) Prediction Procedures: The Clinical-statistical Controversy,” *Psychology, Public Policy, and Law*, 2, 293–323.
- 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, 19–30.
- Gulshan, Varun, Lily Peng, and Marc Coram (2016), “Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs,” *Journal of American Medical Association*, 316, 2402–2410.
- Inman, Jeffrey J., Margaret C Campbell, Amna Kirmani and Linda L Price (2018), “Our Vision for the *Journal of Consumer Research*: It’s All about the Consumer,” *Journal of Consumer Research*, 44(5), 955–59.
- Israyelyan, Anna (2017), “Accelerating Precision Oncology with Artificial Intelligence,” *Cancer Genetics*, accessed at <http://www.cancergenetics.com>, June 20th
- Fromkin, Howard L. (1968), “Affective and Valuational Consequences of Self-perceived

- Uniqueness Deprivation,” unpublished doctoral dissertation, Ohio State University.
- Kahn, Barbara E and Jonathan Baron (1995), “An Exploratory Study of Choice Rules Favored for High-Stakes Decisions,” *Journal of Consumer Psychology*, 4(4), 305-28.
- Kahn, Barbara E., Eric Greenleaf, Julie R. Irwin, Alice M. Isen, Irwin P. Levin, Mary Frances Luce, Manuel C. F. Pontes, James Shanteau, Marc Vanhuele and Mark J. Young (1997), “Examining Medical Decision Making from a Marketing Perspective,” *Marketing Letters*, 8(3), 361-75.
- Keefe, Brian, Usha Subramanian, William M. Tierney, Edmunds Udris, Jim Willems, Mary McDonnell, and Stephan D. Fihn (2005), “Provider Response to Computer-Based Care Suggestions for Chronic Heart Failure,” *Medical Care*, 43(5), 461-65.
- Keller, Punam Anand and Donald R. Lehmann (2008), “Designing Effective Health Communications: A Meta-Analysis,” *Journal of Public Policy and Marketing*, 27(2), 117-30.
- Khosla, Vinod (2012), “Technology Will Replace 80% of What Doctors Do,” *Fortune*, December 4th
- Kirmani, Amna (2015), “Neatly Tied with a Bow,” *Journal of Consumer Psychology*, 25(2), 185–6.
- Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig and Sendhil Mullainathan (2017), “Human Decisions and Machine Predictions,” *The Quarterly Journal of Economics*, 133(1), 237-93.
- Kruger, Justin (1999), “Lake Wobegon Be Gone! The “Below-Average Effect” and the Egocentric Nature of Comparative Ability Judgments,” *Journal of Experimental Social Psychology*, 40(3), 332-40.

- Kuncel, Nathan R., David M. Klieger, Brian S. Connelly and Deniz S. Ones (2013), "Mechanical Versus Clinical Data Combination in Selection and Admissions Decisions: A Meta-Analysis," *Journal of Applied Psychology*, 98(6), 1060-72.
- Hayes, Andrew F. (2013), *Mediation, Moderation, and Conditional Process Analysis*, New York, NY: Guilford.
- Haenssle, H. A., Fink, C., Schneiderbauer, R., Toberer, F. T., Buhl, F.T., Blum, A., Kallou, A., A Ben Hadj Hassen, A., Thomas, L. Enk, A., Uhlmann, L.(2018), "Man Against Machine: Diagnostic Performance of a Deep Learning Convolutional Neural Network for Dermoscopic Melanoma Recognition in Comparison to 58 Dermatologists," *Annals of Oncology*, 29 (8), 1836–1842.
- Haslam, Nick (2006), "Dehumanization: An Integrative Review," *Personality and Social Psychology Review*, 10(3), 252-64.
- Haslam, Nick, Paul Bain, Lauren Douge, Max Lee and Brock Bastian (2005), "More human than you: Attributing Humanness to Self and Others," *Journal of Personality and Social Psychology*, 89, 937–50.
- Highhouse, Scott (2008) "Stubborn Reliance on Intuition and Subjectivity in Employee Selection," *Industrial and Organizational Psychology: Perspectives on Science and Practice*, 1, 333–42.
- Hsee, Christopher K., Sally Blount, Gerooge F. Loewenstein and Max H. Bazerman, (1999), "Preference Reversals between Joint and Separate Evaluations of Outcomes," *Psychological Bulletin*, 125, 576–90.
- Horney, Karen (1937), *The Neurotic Personality of Our Time*. New York: Norton.
- Hutson, Matthew (2017), "Self-taught Artificial Intelligence Beats Doctors at Predicting Heart

- Attacks,” *Science Magazine*, April 14th
- Larwood, Laurie and William Whittaker (1977), “Managerial Myopia: Self-serving Biases in Organizational Planning,” *Journal of Applied Psychology*, 62, 194–8.
- Leachman, Sancy A. and Glen Merlino (2017), “Medicine: The Final Frontier In Cancer Diagnosis,” *Nature*, 542, 36-8.
- Leyens, Jacques-Philippe, Armando Rodriguez-Perez, Ramon Rodrigues-Torres, Ruth Gaunt, Mariapaola Paladino, Jeroen Vaes and Stephanie Demoulin (2001), “Psychological Essentialism and the Differential Attribution of Uniquely Human Emotions to Ingroups and Outgroups,” *European Journal of Social Psychology*, 31, 395-411.
- Logg, Jennifer M. (2018), “Theory of Machine: When Do People Rely on Algorithms?” working paper.
- Lohr, Steve (2016), “IBM Is Counting on Its Bet on Watson, and Paying Big Money for It,” *The New York Times*, October 17th
- Loughnan, Stephen and Nick Haslam (2007), “Animals and Androids,” *Psychological Science*, 18(2), 116-21.
- Marr, Bernard (2018), “How Is AI Used In Healthcare: 5 Powerful Real-World Examples That Show The Latest Advances,” *Forbes*, July 27th
- Maslow, Abraham H. (1962), *Toward a Psychology of Being*, New York: Van Nostrand.
- Meehl, Paul (1954), *Clinical Versus Statistical Prediction: A theoretical Analysis and Review of the Literature*. Minneapolis, MN: University of Minnesota Press.
- Metha, Shivan J. (2015), “Patient Satisfaction Reporting and Its Implications for Patient Care,” *AMA Journal of Ethics*, 17, 616-621.
- Montagu, Ashley and Floyd W. Watson (1983), *The dehumanization of Man*, New York:

McGraw-Hill.

Nissenbaum, Helen and Decker Walker (1998), “Will Computers Dehumanize Education? A Grounded Approach to Values at Risk,” *Technology in Society*, 20(3), 237-73.

O’Hear, Steve (2017), “Babylon Heath Partners With UK’s NHS to Replace Telephone Helpline with AI-powered Chatbot,” *Techcrunch*, January 4th.

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,” *Journal of Behavioral Decision Making*, 22, 390–409.

Palmeira, Mauricio and Gerri Spassova (2013), “Consumer Reactions to Professionals Who Use Decision Aids,” *European Journal of Marketing*, 49(3), 302-26.

Promberger, Mariane and Jonathan Baron (2006), “Do Patients Trust Computers?” *Journal of Behavioral Decision Making*, 19, 455-68.

Sanders, Nada R. and Karl B. Manrodt (2003), “The Efficacy of Using Judgmental Versus Quantitative Forecasting Methods in Practice,” *The International Journal of Management Science*, 31, 511-22.

Sawyer, Jack (1966), “Measurement and Prediction, Clinical and Statistical,” *Psychological Bulletin*, 66, 178–200.

Scopelliti, Irene, Carey K. Morewedge, Erin McCormick, H. Lauren Min, Sophie Lebrecht and Karim S. Kassam (2015), “Bias Blind Spot: Structure, Measurement, and Consequences,” *Management Science*, 61(10), 2468-86.

Sensening and Jack W. Brehm (1968), “Attitude Change From an Implied Threat to Attitudinal Freedom,” *Journal of Personality and Social Psychology*, 8, 324-30.

Shaffer, Victoria A., Adam C. Probst, Edgar C. Merkle, Hal R. Arkes and Mitchell A. Medow

- (2013), "Why Do Patients Derogate Doctors Who Use a Computer-Based Diagnostic Support System?" *Medical Decision Making*, 1, 108-18.
- Simsek, Ömer Faruk and Berna Yalınçetin (2010), "I Feel Unique, Therefore I am: The Development and Preliminary Validation of the Personal Sense of Uniqueness (PSU) scale," *Personality and Individual Differences*, 49, 576–81.
- Sinha, Rashmi and Kirsten Swearingen (2001), "Comparing Recommendations Made by Online Systems and Friends," *Proceedings of the DELOS-NSF Workshop on Personalization and Recommender Systems in Digital Libraries*, available at www.onemweb.com
- Siwicki, Bill (2018), "Telemedicine Innovation: Rural Health System Keeping More People Alive at Lower Costs," *Healthcare News*, September 7th.
- Spiller, Stephen A, Gavan J. Fitzsimons, John G. Lynch Jr. and Gary H. McClelland (2013), "Spotlights, Floodlights, and the Magic Number Zero: Simple Effects Tests in Moderated Regression," *Journal of Marketing Research*, 50(2), 277-88.
- Snyder, Charles R. (1978), "The Illusion of Uniqueness," *Journal of Humanistic Psychology*, 18, 33-41.
- Snyder, Charles R., and Harold L. Fromkin (1980), *Uniqueness: The Human Pursuit of Difference*, New York: Plenum.
- Steele, Claude M. (1988), "The Psychology of Self-affirmation: Sustaining the Integrity of the Self," In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology* (Vol. 21, pp. 261–302). New York: Academic.
- Svenson Ola (1981), "Are We Less Risky and More Skillful Than Our Fellow Drivers?" *Acta Psychologica*, 47, 143-51.
- Taylor, Shelley E. and Jonathon D. Brown (1988), "Illusion and Well-being: A Social

- Psychological Perspective On Mental Health,” *Psychological Bulletin*, 103(2), 193–210.
- Windschitl, Paul D, Justin Krueger and Ericka Nus Simms (2003), “The Influence of Egocentrism and Focalism on People’s Optimism in Competitions: When What Affects Us Equally Affects Me More,” *Journal of Personality and Social Psychology*, 85(3), 389-408.
- Wormith, Stephen J. and Colin S. Goldstone (1984), “The Clinical and Statistical Prediction of Recidivism,” *Criminal Justice and Behavior*, 11(1), 3-29.
- Ubel, Peter A. (2002), “Is Information Always a Good Thing? Helping Patients Make ‘Good’ Decisions,” *Medical Care*, 40(9), V39–44.
- Yeomans, Michael, Sendhil Mullainathan, Anuj Shah and J Kleinberg (2018), “Making Sense of Recommendations,” working paper.
- Yu, Kun-Hsing, Andrew L. Beam, and Isaac S. Kohane, (2018) “Artificial Intelligence in Healthcare,” *Nature Biomedical Engineering*, 2(10), 719.

TABLES

TABLE 1 – DESIGN OF STUDIES 3A, 3B, 3C

	Care Delivered	Performance Framing	Degree of Agency	Format of Performance
Study 3A	prevention	positive/hits (accuracy rate)	high	numeric
Study 3B	diagnosis	positive/hits (accuracy rate)	high	numeric
Study 3C	treatment	negative/misses (complication rate)	low	numeric + visual
	Provider Performance			
	X = Y	X < Y	X << Y	
Study 3A	X=90% (accuracy rate) Y=90% (accuracy rate)	X=90% (accuracy rate) Y=93% (accuracy rate)	X=90% (accuracy rate) Y=96% (accuracy rate)	
Study 3B	X=90% (accuracy rate) Y=90% (accuracy rate)	X=90% (accuracy rate) Y=95% (accuracy rate)	X=85% (accuracy rate) Y=95% (accuracy rate)	
Study 3C	X=3.9% (complication rate) Y=3.9% (complication rate)	X=4.2% (complication rate) Y=3.6% (complication rate)	X=4.2% (complication rate) Y=3.3% (complication rate)	

TABLE 2 – PREFERENCE FOR PROVIDER IN STUDIES 3A, 3B, 3C

		Relative Performance of Medical Providers		
		X = Y	X < Y	X << Y
Study 3A	Human(X) - Human(Y)	4.00 (0.00)	6.06 (0.98)	6.62 (0.74)
	Human(X) - AI(Y)	3.11 (1.65)	4.64 (2.03)	5.97 (1.29)
		$F(1,199) = 8.6$ $p = .004$	$F(1,199) = 19.9$ $p = .000$	$F(1,199) = 4.29$ $p = .04$
Study 3B	Human(X) - Human(Y)	4.11 (0.53)	6.32 (1.16)	6.85 (0.36)
	Human(X) - AI (Y)	2.61 (1.47)	4.75 (2.01)	5.54 (1.76)
		$F(1,221) = 22.9$ $p = .000$	$F(1,221) = 24.9$ $p = .000$	$F(1,221) = 15.9$ $p = .000$
Study 3C	Human(X) - Human(Y)	4.00 (0.00)	6.08 (1.15)	6.45 (0.85)
	Human(X) - AI (Y)	2.82 (1.36)	5.25 (1.65)	5.48 (1.41)
		$F(1,229) = 18.6$ $p = .000$	$F(1,229) = 9.3$ $p = .003$	$F(1,229) = 13.2$ $p = .000$

Note: Scale endpoints marked, *Definitely Provider X* (1) and *Definitely Provider Y* (7); standard deviations are in parentheses.

TABLE 3 – UTILITIES FROM CHOICE-BASED CONJOINT

Attribute	Level	Average Utilities	Standard Deviation
Provider	Human, Robotic	-29.78	52.23
Accuracy	81%, 84%, 87%	26.54	10.94
Cost	\$20, \$40, \$60	-2.24	1.69

FIGURES

FIGURE 1 – OVERVIEW OF THE STUDIES

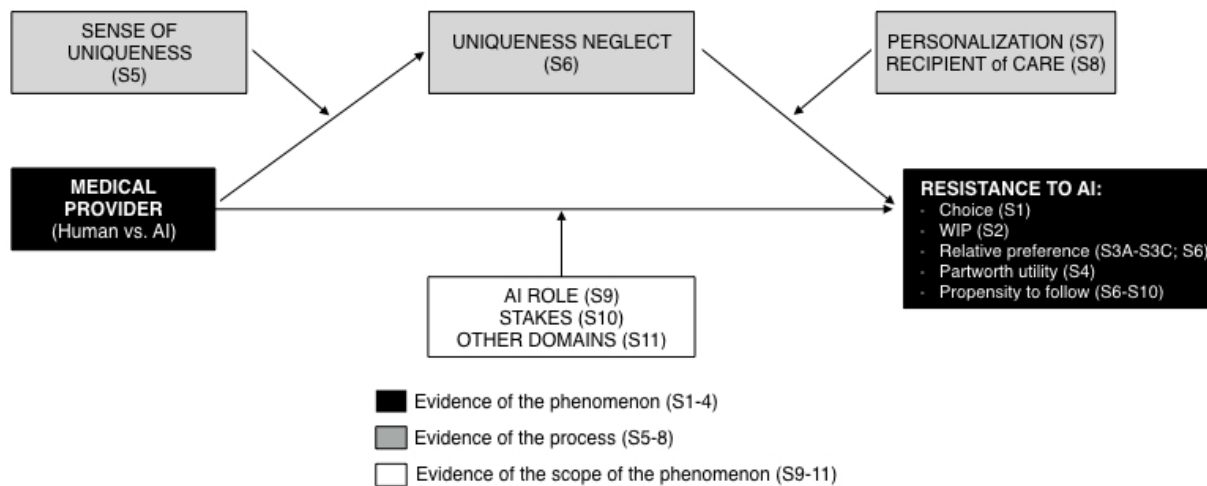
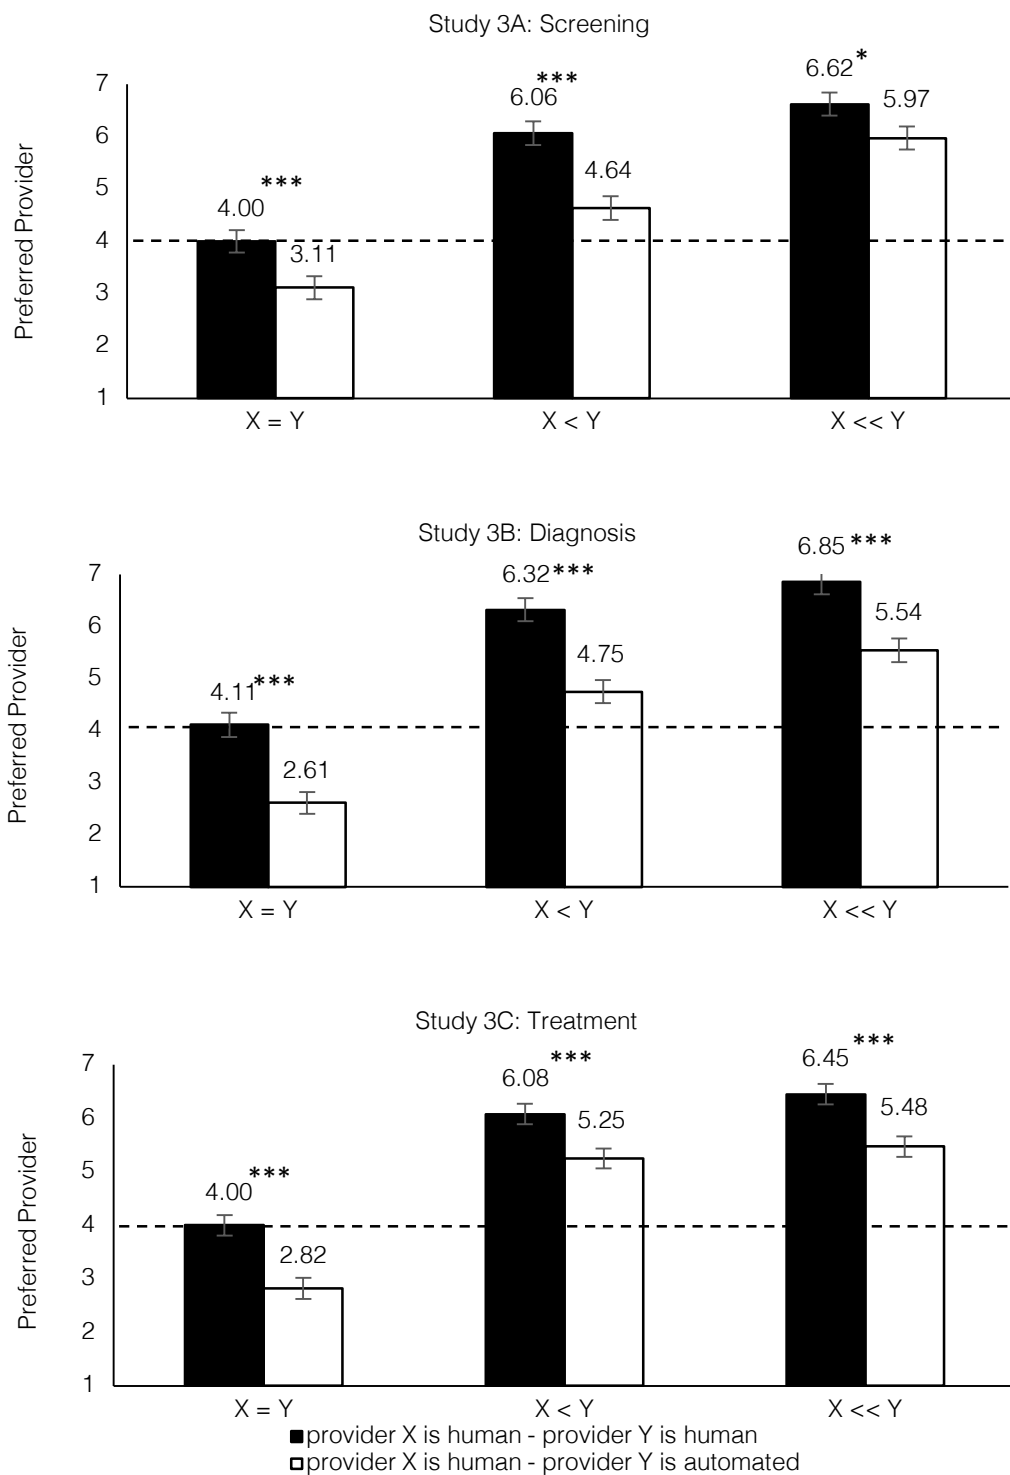


FIGURE 2 – PREFERENCE FOR AUTOMATED VERSUS HUMAN PROVIDER



Note: Scale endpoints marked, *Definitely Provider X* (1) and *Definitely Provider Y* (7)

* $p < .05$, ** $p < .01$, *** $p < .001$

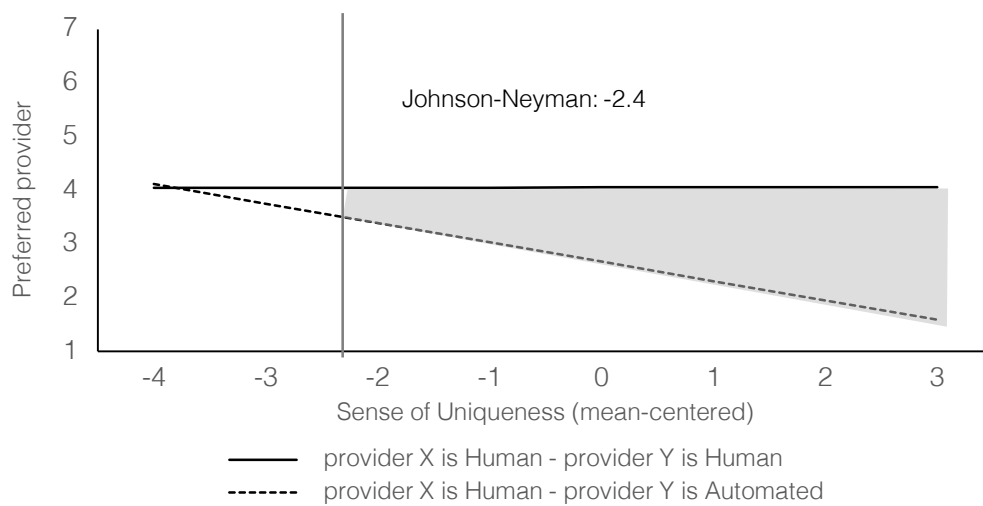
FIGURE 3 – SAMPLE CHOICE TASK IN STUDY 4

If these were your only options, which would you choose?

(1 of 7)


Provider	Human	Robotic	Human
Accuracy	81%	81%	87%
Cost to you	\$40	\$20	\$60
	<input type="button" value="Select"/>	<input type="button" value="Select"/>	<input type="button" value="Select"/>

FIGURE 4 - SENSE OF UNIQUENESS MODERATES PROVIDER PREFERENCE



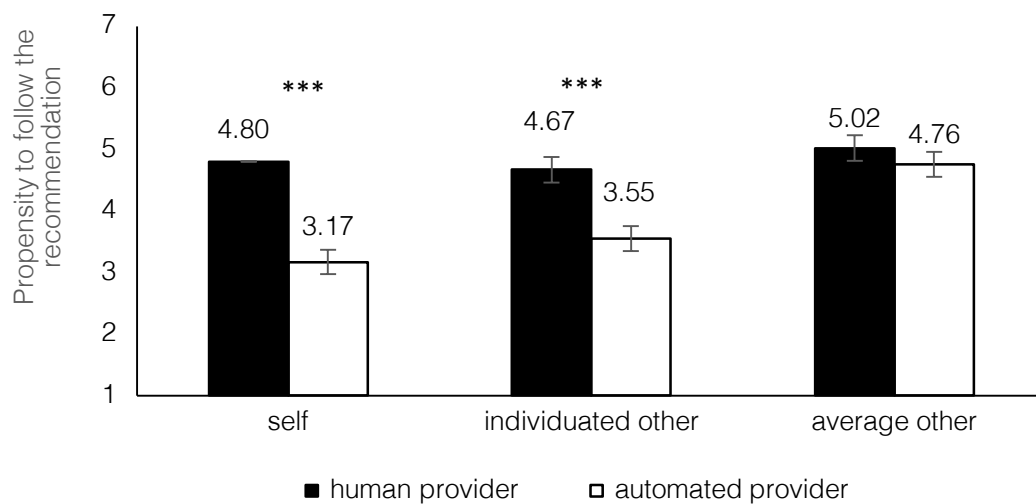
Note: Scale endpoints marked, *Definitely Provider X* (1) and *Definitely Provider Y* (7)

FIGURE 5 – PERSONALIZED CARE CURBS UNIQUENESS NEGLECT



Note: Scale endpoints marked, *Definitely not follow* (1) and *Definitely follow* (7); *** $p < .001$

FIGURE 6 –RESISTANCE TO MEDICAL AI IS MITIGATED FOR AVERAGE OTHERS



Note: Scale endpoints marked, *Definitely not follow* (1) and *Definitely follow* (7)
*** $p < .001$