Special Section: Open Forum

Engineering Medical Decisions

Computer Algorithms and the Manipulation of Choice

MEREDITH STARK and JOSEPH J. FINS

Algorithms and Health Decisions

Survey evidence continues to illustrate that U.S. adults rely rather substantially on the Internet as a resource for health information.[1] Yet other than those with advanced computer science training, few individuals likely understand, or perhaps even query, the processes behind the search results, advertising, and content that is displayed. Fewer still may be aware of the extent and impact of widespread data-collection practices and the sophisticated analytic algorithms in operation behind the scenes (and screens) to deliver targeted output to end users. A July 2011 article by computer science journalist Tom Spring, "Algorithms That Rule the Web," discussed how these newer "tools help Google, Facebook, Amazon, and others stay one step ahead of us online. . . . It's hard to say whether the computer algorithms that these services use to anticipate our needs and wants are turning us into puppets or geniuses. But algorithms have a huge impact on our tastes, buying habits, and decisions."[2] Whereas toying with choice in the realm of books, movies, and vacation destinations may yield an array of assessments ranging from appreciation to irritation, in the realm of medical decisions, manipulating patient choices, and perhaps ultimately treatment decisions, arguably treads on shakier ethical ground, terrain in need of further scrutiny.

In considering the particular challenges posed by the application of algorithmic online advertising processes to health decisions, it may be instructive to consider, at the outset, a hypothetical example of how influence might be exerted on a health decision. Imagine an individual searches for information on a painful arthritic knee. Having typed keywords into a search engine, browsed various websites, and joined the arthritis group of a social networking site, the individual begins to see, on the margins of various web pages, advertisements for a specific brand of pain reliever. Perhaps the ads even tout the suitability of that brand for arthritic pain. Soon ads are displayed for a nearby orthopedic practice and associated surgical center, the leading regional provider of knee replacements. As the consumer moves between the websites of a shoe store, a travel site, a pet-supply center, and a news page, the same ads for the surgical center are featured time and again. Likewise, the surgical center website now tops the list of her search engine results. To learn more about the condition and the surgery and to inform her decision, the patient visits the website.

When thinking about how to treat this intractable knee pain, the patient now keeps returning to the idea of a knee replacement. (Coincidence? We tend to think not.)

Cambridge Quarterly of Healthcare Ethics (2013), **22**, 373–381.

doi:10.1017/S0963180113000224

From information on the surgical center website, the individual has learned that the surgery would provide a permanent solution to the pain; she wouldn't need to go to physical therapy or have to try various medications. The patient learns that she would not need to stay overnight at a hospital, and a gallery of photographs shows luxurious and state-of-the art operating and recovery rooms. The most prominent of the patient testimonials features the narrative of a woman who was able to resume jogging with her dog after the surgery. Our patient, a dog owner, feels an immediate kinship and considers this a sign. From other ads now displayed, the patient even knows the brand name of a leading type of artificial joint and why that brand is best. Conversations ensue online and offline, and these exchanges all confirm her gut feeling; she is now certain she wants to have the surgery, and she already knows just the place to go, and just the brand of artificial joint to request. In discussions with the surgeon at the center, the patient conveys her choice, and the physician incorporates these strong patient preferences into his treatment recommendation, as this is clearly an informed patient who has researched and thought comprehensively about all of the various options. Or has she?

Behind the Pixels

Drawing on the dramaturgic metaphors of sociologist Erving Goffman,[3] in order to understand how individuals interact with health information in front of the monitor, on the putative stage, it may be necessary to journey behind the scenes to examine the background data collection and analytic algorithms in play. Originally termed online profiling,[4,5] in current industry parlance the use of technology to target advertising to individuals is referred to as online behavioral advertising (OBA) or behavioral targeting.[6] These strategies involve "the practice of collecting data about an individual's online activities for use in selecting which advertisement to display."[7] A 2009 Federal Trade Commission (FTC) press release, announcing a new staff report, notes the durability of these challenges: "Over the last decade, the FTC has periodically examined the consumer privacy issues raised by online behavioral advertising—which is the practice of tracking an individual's online activities in order to deliver advertising tailored to his or her interests."[8]

Notably, the extent and manner of such intrusion, and the associated impact, may be escalating with the growing sophistication and proliferation of highly advanced computational processes and algorithms. A June 2012 article posted on cnn.com—the web is also a source of scholarship in itself—discussed the development of new capabilities that will allow Facebook to display ads, not predicated solely on activity within the social network but based on integration of information across online venues.[9] "A new Facebook system will use your activity on other websites to send you what Facebook thinks are ads about your current interests. Advertisers will, in effect, be bidding to get their ads in front of you."[10] Worryingly, an executive of a technology firm was quoted as heralding the subtlety of the new algorithms: "'It's not going to be discernible to most consumers,' he said. 'Most people won't notice any difference.'"[11]

Industry-wide, online behavioral advertising is big business, with advertisers spending $8 billion annually, not including spending on search engine ads.[12]

A July 2012 *New York Times* article discussed how Facebook, in a critical effort to shore up its value in the financial sector, is in the midst of developing new strategies and algorithms to capitalize on "the pile of personal data it collects from 900 million users . . . to serve up effective, profitable advertisements."[13] Efforts by Facebook include the acquisition of a start-up company that builds "artificial-intelligence-based algorithms to analyze online data"; however, no other details surrounding these new initiatives were released.[14] Regrettably, apart from accounts in the press, there is scant information on the inner workings of these mostly proprietary strategies and algorithms.

In assessing the qualitative nature of online health advertising and the possible impact on health decisions, consideration of pharmaceutical direct-to-consumer advertising (DTCA) may be informative. First, empirical research on DTCA highlights a critically important pathway through which influencing patient preferences might translate into impact on care; research on the efficacy of DTCA consistently illustrates that patient requests for pharmaceuticals following direct-to-consumer ads remain "a powerful driver" of physician prescribing behavior.[15] Second, attention to the novel strategies developed for advertising pharmaceuticals online[16] may be instructive for those concerned with behaviorally targeted advertising of an array of health services. Of substantial concern is that industry sponsorship of online content may be increasingly difficult for patients to detect, with strategies including the use of third-party testimonials and the establishment of "unbranded" websites to obscure corporate sponsorship.[17]

In terms of oversight, although online DTCA practices remain underregulated,[18] the Food and Drug Administration (FDA) has been called on to intervene with new guidelines and limits.[19] Notably, in the case at hand, the diffusion of targeted online behavioral advertising of health products and services across a broader swath of providers of healthcare products and services, coupled with the individualized and nonuniform impact on consumers based on discrete search, browsing, and social media histories, may further hamper oversight of this emerging form of health advertising.

The computational power of these algorithms and the application to shaping medical choice carry distinct challenges, marking a departure from the literature to date. Scholars have comprehensively examined online health information;[20] the issues addressed range from health information technology and electronic medical records[21] to the inherent privacy and sensitivity issues of health data[22,23] and from direct-to-consumer advertising of pharmaceuticals on the Internet[24,25] to informational access as potentially undermining the physician-patient relationship.[26,27] Furthermore, scholars from the disciplines of law[28,29] and information technology[30,31] have addressed issues arising from behavioral targeting, with a focus on consumer privacy and data protection. Still unexamined is the confluence of these issues, the use of increasingly complex algorithms to make use of unprecedented quantities of data, so as to meaningfully target individuals for the purposes of expressly interfering with health decisions. Although targeted advertising practices are not especially novel, the development of increasingly sophisticated capabilities to make use of ever-larger caches of data, coupled with the application to health decisions, may sufficiently alter the ethical calculus so as to compel careful reconsideration of the challenges posed.

Hazards of Commercial Targeting of Patient Medical Choice

Vulnerabilities in Decisionmaking

Industry efforts to individually target individuals with health advertising, making use of complex technologies and unprecedented quantities of personal data to skew choice, effectively takes aim at free and informed medical decisionmaking. Methods of informational distortion—tinkering with the choices presented, promoting and touting some options in searches for health information while obscuring others, utilizing social networks to craft the appearance of peer recommendations and endorsements where there may be none, and delivering content that appears neutral when it may be commercially sponsored—all conducted in the name of influencing health decisions to benefit not the well-being of the patient but the fiscal health of a third party, severely challenge ideals of respect for persons and decisional autonomy. Moreover, as Thaler and Sunstein write, "the emerging science of choice, consisting of careful research by social scientists over the past four decades" has, in turn, "raised serious questions about the rationality of many judgments and decisions that people make."[32] To the extent that the cognitive resources of patients may be outmatched by industry investment and outpaced by computational power, and to the extent that some populations may be rendered more vulnerable to this interference, considerations of fairness and justice also rise to the fore.

Compounding the potential for interference is the unfortunate notion that, even absent heavy interference, individuals consistently exhibit flawed decisional strategies marked by an array of cognitive heuristics and biases.[33] For example, individuals may rely on the *availability heuristic*, whereby an option is credited or favored according to the ease with which instances are brought to mind.[34] Also relevant is *anchoring*, the tendency to set on an initial position, even a faulty value, with only fine adjustments thereafter.[35] Further salient may be *confirmation bias*, the tendency to search for and credit information that supports an initial hypothesis.[36] A bias toward *overconfidence* and a "universal tendency to believe we know more than we do"[37] may likewise shape the decisional dynamic, perhaps blunting further inquiry and conveying confidence in the strength of the process and the accuracy of the selection.

Returning to our example of the arthritis patient, the application of behavioral targeting resulted in repeated exposures to advertisements for the leading regional knee replacement surgical center. That repetition may directly play into the availability heuristic, as ready retrieval from memory is fostered, conferring advantage to the surgical option. Application of the anchoring heuristic suggests that the patient may dig in and adhere to an initial position, which is particularly salient if the surgical option is positively presented early in the decision process. The confirmation bias may serve to constrict latitude in information gathering and thinking, thereby ensuring that subsequent conversations and thinking on the matter are filtered and directed so as to support the initial choice. In light of inherent biases toward overconfidence, the patient may summarily conclude her fact finding, certain that she has arrived on the right choice; she assuredly conveys her preferences to the surgeon, who is influenced by her strong convictions and the apparent comprehensiveness of her decisional process.

Voluntariness and Consent

The assembly of extensive caches of personalized data, thereafter analyzed and applied to shape information delivery, calls into question the extent to which these practices are unfolding with the understanding and voluntary participation of those targeted. Practices relying on the collection and imperceptible algorithmic analysis of information, from web searches, browsing history, purchasing behavior, and participation in social media, absent the knowledge and consent of individuals, challenge the boundaries of voluntariness. A 2009 empirical study conducted at Carnegie Mellon, for example, found that individuals "have a poor understanding of how Internet advertising works, . . . believe that their actions online are completely anonymous unless they are logged into a website, and believe that there are legal protections that prohibit companies from sharing information they collect online. We found that participants have substantial confusion about the results of the actions they take within their browsers."[38] Subsequent studies have likewise found considerable misinformation among end users as to the extent and nature of these practices.[39] Lengthy legal disclosures, if available to consumers, may do little to remedy issues of understanding and consent.

The vast and steadily growing quantities of individual information now available online have further muddied the picture. Individuals may exercise a voluntary sharing of information in an online community for the purposes of promoting the health of themselves or others, or for building and extending social bonds. Likewise, individuals may browse websites or utilize a search engine to learn more about a health condition. Although these are voluntary actions, what takes place thereafter may very well not be. Contributing extensive personal information for commercial use, providing a valuable warehouse of input to be translated by algorithmic functions into later attempts to impact their own health decisions, all in the name of augmenting industry profit, may extend far beyond the scope of original intent. "At the heart of this industry," write marketing and IT professors Goldfarb and Tucker, "is the detailed collection, parsing, and analysis of consumer data, often without consumers' consent or knowledge."[40]

Worryingly, the impact of targeted ads may be exacerbated by the centrality of patient preferences in treatment decisions. In the last decades, bioethics has successfully championed a shift away from paternalism in medical decisions toward autonomy and a broadening of the space accorded patient preference in healthcare choices.[41] Parallel efforts in the health services and policy communities have advanced the construct of shared decisionmaking, advocating for a collaborative process that blends physician guidance and patient preferences.[42] In the burgeoning realm of preference-sensitive care, in which patient input into decisions ranging from joint replacements to cholesterol lowering and to bariatric surgery is increasingly taken into account,[43] the demands on patients to be informed consumers of healthcare have grown more urgent, as has the critical import of safeguarding medical choice.

Intrusion in and appropriation of the space afforded patients to be more active participants in healthcare decisions are perhaps all the more paradoxical, as the availability of these online tools ostensibly should function to increase autonomy, allowing individuals more access to information and more input into health decisions. Instead, owing to intricate means of behavioral targeting, these strategies may be lulling individuals into a sense of informed decisionmaking while actually

directing them toward specified outcomes, in a manner that thwarts voluntariness, engineers choice, and surreptitiously exploits decisional vulnerabilities. Whereas issues of differential access to online health information, a "digital divide," customarily yield concerns of disadvantage for those without adequate digital access,[44] in this case, the traditional model may be upended, with vulnerability conferred on those who craft health decisions with an unwary reliance on the information delivered to their screens.

Transparency and Conflicts of Interest

Wholly intertwined with issues of voluntariness are concerns of transparency, as lack of transparency stands as dispositive to informed voluntary participation. Transparency concerns are salient across the process, from the extent of data aggregation, to the complex means of analysis, and to the informational distortion and the subtle impact on decisionmaking. As with lack of voluntariness, transparency is particularly at issue in regard to data collection and analysis, as the extent to which data is harvested from buying behavior, web browsing, and social media use, and is thereafter aggregated and utilized to shape delivery of health information, remains substantially opaque.[45] Further troubling is the latency of the decisional interference, the manner by which some options may be prioritized and advertised while others are obscured from view, yielding a faint distortion likely to escape consumer detection and, thus, remedy. Critically, for our arthritis patient, she likely remains unaware of both the targeted advertising practices and, more so, any compromise of or intrusion into her decisionmaking.

Furthermore, these practices may engender conflicts of interest and may precipitate concerns of beneficence and nonmaleficence. As healthcare providers engage in online advertising and behavioral targeting to market services, mediated by third-party vendors, conflicts may be realized, pitting the interests and informed decisions of patients against the financial aims of providers. These practices may invoke associated concerns of dual agency, as vectors of provider financial interest and patient best interest may be in opposition. By engaging in these targeting practices to manipulate demand for their services, providers may be complicit, even unknowingly, in preying on the decisional vulnerabilities of those they are charged with protecting. In the case of our patient with arthritis, the same surgeon whose group practice paid for marketing and advertising to increase visibility, patient demand, and revenue may also collaborate in the patient's decision to have the surgery and may yet perform the operation.

Meanwhile, the relevance and import of these assembled challenges will likely only deepen with time. As the Internet, particularly social media sites, grows as a resource for health information,[46,47] as access expands via smart phones,[48] and as individuals report greater trust in the Internet as a resource for health information,[49] the critical import of safeguarding health decisions will intensify accordingly.

Conclusion and Next Steps

Although discussions of online targeting algorithms, including concerns regarding privacy and surveillance, are not particularly new in the scholarly communities of computer science, information technology, and Internet law,[50,51,52] what is new is

the application of these technologies to influencing medical choice; this subtle encroachment in the realm of health decisions yields far deeper concerns and may also require dedicated remedies. In recent years, the FTC has generated much-needed attention to the issue of behavioral targeting, although the relevant reports have largely confined concerns for health to the potential threats to privacy in using sensitive health information as input into these algorithms.[53,54] Unexamined are the distinct challenges that occur when the *output* is expressly trained on health decisions. Whether specific regulations are warranted remains an open question, although at minimum, the current situation calls for further inquiry. In contemplating regulatory efforts, we must take care to balance our closely held and hallmark protections of free speech, in all forms, with concomitant mandates to protect patients and preserve free and informed medical decisions.

In the meantime, there is another balance to be realized with respect to the potential for new technologies to both expand and constrict thinking and choice. Tom Spring writes,

> Back in the 20th century—the primordial age of algorithms—life was simpler and harder at the same time. We never knew what else we might want to buy at Amazon; we didn't know what the most "important" news stories of the day were; and before the Netflix movie recommendation engine, we had no mechanized assistance in determining which DVD to rent next. When we're looking for something online, Google's algorithm frees us from having to sort and search through multitudes of only not-very-relevant results. On the other hand, algorithms might trap us in a world where advertisers and government agencies couple behavioral data with computer formulas to predict and manipulate what we do or buy next.[55]

Whether these algorithms are rendering us geniuses or puppets in our health decisions remains unknown. Assured is the ongoing need to venture backstage, behind the curtain and beyond the pixels, in order to ask these important questions, answer them, and, not long after, ask yet again.

Notes

1. Fox S. Pew internet: Health 2012. *Pew Internet*; 2012 Mar 1; available at http://www.pewinternet.org/Commentary/2011/November/Pew-Internet-Health.aspx (last accessed 19 June 2012).
2. Spring T. Algorithms that rule the web. *PCWorld*; 2011 July 24; available at http://www.pcworld.com/printable/article/id,236226/printable.html (last accessed 18 June 2012).
3. Goffman E. *The Presentation of Self in Everyday Life*. New York: Penguin Books; 1959.
4. Federal Trade Commission. Online profiling: A report to Congress. *Federal Trade Commission*; 2000 June; available at http://www.ftc.gov/os/2000/06/index.shtm#13 (last accessed 8 Aug 2012).
5. Cantos L, Fine L, Singer R. FTC releases online profiling report. *Intellectual Property & Technology Law Journal* 2000;12(10):22–3.
6. Federal Trade Commission. FTC staff revises online behavioral advertising principles. *Federal Trade Commission*; 2009 Feb 12; available at http://www.ftc.gov/opa/2009/02/behavad.shtm (last accessed 25 July 2012).
7. McDonald AM, Cranor LF. An empirical study of how people perceive online behavioral advertising. *Carnegie Mellon University Cylab*; 2009 Nov 10; available at http://www.cylab.cmu.edu/research/techreports/2009/tr_cylab09015.html (last accessed 1 Aug 2012).
8. See note 6, Federal Trade Commission 2009.
9. Gross D. Facebook to show you ads based on your Web browsing. *CNN*; 2012 June 18; available at http://www.cnn.com/2012/06/18/tech/social-media/facebook-ads-web-browsing/index.html (last accessed 19 June 2012).

10. See note 9, Gross 2012.
11. See note 9, Gross 2012.
12. Goldfarb A, Tucker CE. Economic and business dimensions: Online advertising, behavioral targeting, and privacy. *Communications of the ACM* 2011;54(5):25–7.
13. Sengupta S. Facebook efforts on advertising face a day of judgment. *New York Times* 2012 July 22; available at http://www.nytimes.com/2012/07/23/technology/facebook-advertising-efforts-face-a-day-of-judgment.html (last accessed 25 July 2012).
14. See note 13, Sengupta 2012.
15. Mintzes B, Barer ML, Kravitz RL, Kazanjian A, Bassett K, Lexchin J, et al. Influence of direct to consumer pharmaceutical advertising and patients' requests on prescribing decisions: Two site cross sectional survey. *British Medical Journal* 2002;324:278–9.
16. Liang BA, Mackey T. Direct-to-consumer advertising with interactive internet media: Global regulation and public health issues. *JAMA* 2011;305(8):824–5.
17. See note 16, Liang, Mackey 2011, at 824.
18. See note 16, Liang, Mackey 2011.
19. Kuehn BM. FDA weighs limits for online ads. *JAMA* 2010;303(4):311–3.
20. Nuffield Council on Bioethics. Chapter 5 online health information: Personalised healthcare. *Nuffieldbioethics.org*; 2010 Oct; available at http://www.nuffieldbioethics.org/personalised-healthcare/personalised-healthcare-online-health-information (last accessed 19 June 2012).
21. Goldzweig CL, Towfigh A, Maglione M, Shekelle PG. Costs and benefits of health information technology: New trends from the literature. *Health Affairs* 2009;28(2):w282–w293.
22. See note 20, Nuffield Council on Bioethics 2010.
23. Love V. Privacy ethics in health care. *Journal of Health Care Compliance* 2011 July/Aug:15–24.
24. Lurie P. DTC advertising harms patients and should be tightly regulated. *Journal of Law, Medicine & Ethics* 2009 Fall:444–50.
25. See note 16, Liang, Mackey 2011.
26. Mathieu E. The internet and medical decision making: Can it replace the role of health care providers? *Medical Decision Making* 2010 Sep–Oct:14S–16S.
27. Rosenberg L. Does direct-to-consumer marketing of medical technologies undermine the physician-patient relationship? *American Journal of Bioethics* 2009;9(4):22–3.
28. Cox JT, Cline KM. Parsing the demographic: The challenge of balancing online behavioral advertising and consumer privacy considerations. *Journal of Internet Law* 2012 Mar:3–12.
29. Bennet K. Privacy and online behavioral advertising. *Journal of Internet Law* 2007 Dec:24–6.
30. Backes M, Kate A, Maffei M, Pecina K. ObliviAd: Provably secure and practical online behavioral advertising. *IEEE Symposium on Security and Privacy* 2012 Spring; available at www.ieee-security.org/TC/SP2012/papers/4681a257.pdf (last accessed 3 Aug 2012).
31. See note 12, Goldfarb, Tucker 2011.
32. Thaler RH, Sunstein CR. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. New Haven, CT: Yale University Press; 2008, at 7.
33. Tversky A, Kahneman D. Judgment under uncertainty: Heuristics and biases. *Science* 1974;185:1124–31.
34. See note 33, Tversky, Kahneman 1974.
35. See note 33, Tversky, Kahneman 1974.
36. Croskerry P. The importance of cognitive errors in diagnosis and strategies to minimize them. *Academic Medicine* 2003;78(8):775–80.
37. See note 36, Croskerry 2003, at 778.
38. See note 7, McDonald, Cranor 2009, at 1.
39. Ur B, Leon PG, Cranor LF, Shay R, Wang Y. Smart, useful, scary, creepy: Perceptions of online behavioral advertising. Presented at the Symposium on Usable Privacy and Security (SOUPS); 2012 July 11–13; available at cups.cs.cmu.edu/soups/2012/proceedings/a4_Ur (last accessed 3 Aug 2012).
40. See note 12, Goldfarb, Tucker 2011, at 26.
41. Wolpe PR. The triumph of autonomy in American bioethics: A sociological view. In: DeVries R, Subedi J, eds. *Bioethics and Society: Sociological Investigations of the Enterprise of Bioethics*. New York: Prentice-Hall; 1998:38–59.
42. Moulton B, King JS. Aligning ethics with medical decision making: The quest for informed patient choice. *Journal of Law, Medicine & Ethics* 2010;38(1):85–97.
43. Dartmouth Atlas Project. Topic brief: Preference-sensitive care. *Dartmouth Atlas of Health Care*; 2007; available at www.dartmouthatlas.org/downloads/reports/preference_sensitive.pdf (last accessed 19 June 2012).

44. See note 20, Nuffield Council on Bioethics 2010, at 86.
45. See note 7, McDonald, Cranor 2009.
46. See note 1, Fox 2012.
47. Fox S. The social life of health information. *Pew Internet*; 2011 May 12; available at http://pewinternet.org/Reports/2011/Social-Life-of-Health-Info/Summary-of-Findings.aspx (last accessed 19 June 2012).
48. See note 1, Fox 2012.
49. Wolters Kluwer Health. Survey: Consumers show high degree of trust in online health information, report success in self-diagnosis. *Wolters Kluwer Health*; 2012 May 16; available at http://www.wolterskluwerhealth.com/News/Pages/Survey-Consumers-Show-High-Degree-of-Trust-in-Online-Health-Information,-Report-Success-in-Self-Diagnosis--.aspx (last accessed 9 July 2012).
50. See note 39, Ur et al. 2012.
51. See note 12, Goldfarb, Tucker 2011.
52. See note 28, Cox, Cline 2012.
53. Federal Trade Commission Staff Report. Self-regulatory principles for online behavioral advertising: Behavioral advertising tracking, targeting & technology. *Federal Trade Commission*; 2009 Feb; available at www.ftc.gov/os/2009/02/P085400behavadreport.pdf (last accessed 1 Aug 2012).
54. Federal Trade Commission Report. Protecting consumer privacy in an era of rapid change: Recommendations for businesses and policymakers. *Federal Trade Commission*; 2012 Mar; available at www.ftc.gov/os/2012/03/120326privacyreport.pdf (last accessed 8 Aug 2012).
55. See note 2, Spring 2011.