Deloitte Review

ISSUE 16 | 2015

Complimentary article reprint

THE LAST-MILE PROBLEM

How data science and behavioral science can work together

BY JAMES GUSZCZA
> ILLUSTRATION BY JOSIE PORTILLO


FINISH

THE LAST-MILE PROBLEM

How data science and behavioral science can work together

BY JAMES GUSZCZA > ILLUSTRATION BY JOSIE PORTILLO

"Many programs and services are designed not for the brains of humans but of Vulcans. Thanks in large part to Kahneman and his many collaborators, pupils, and acolytes, this can and will change."
— Rory Sutherland, Ogilvy & Mather[1]

TWO OVERDUE SCIENCES

Near the end of *Thinking, Fast and Slow*, Daniel Kahneman wrote, "Whatever else it produces, an organization is a factory that produces judgments and decisions."[2] Judgments and decisions are at the core of two of the most significant intellectual trends of our time, and the one-word titles of their most successful popularizations have become their taglines. "Moneyball" is shorthand for applying data analytics to make more economically efficient decisions in business, health care, the public sector, and beyond. "Nudge" connotes the application of findings from psychology and behavioral economics to prompt people to make decisions that are consistent with their long-term goals.

Other than the connection with decisions, the two domains might seem to have little in common. After all, analytics is typically discussed in terms of computer technology, machine learning algorithms, and (of course) big data. Behavioral nudges, on the other hand, concern human psychology. What do they have in common?

Quite a bit, as it turns out. Business analytics and the science of behavioral nudges can each be viewed as different types of responses to the increasingly commonplace observation that people are predictably irrational. Predictive analytics is most often about providing tools that correct for mental biases, analogous to eyeglasses correcting for myopic vision. In the private sector, this enables analytically sophisticated competitors to grow profitably in the inefficient markets that exist thanks to abiding cultures of biased decision making in business. In government and health care, predictive models enable professionals to serve the public more economically and effectively.

When the ultimate goal is behavior change, predictive analytics and the science of behavioral nudges can serve as two parts of a greater, more effective whole.

The "behavioral insights" movement is based on a complementary idea: Rather than try to equip people to be more "rational," we can look for opportunities to design their *choice environments* in ways that comport with, rather than confound, the actual psychology of decision making.[3] For example, since people tend to dislike making changes, set the default option to be the one that people would choose if they had more time, information, and mental energy. (For example, save paper by setting the office printer to the default of double-sided printing. Similarly, retirement savings and organ donation programs are more effective when the default is set to "opt in" rather than "opt out.") Since people are influenced by what others are doing, make use of peer comparisons and "social proof" (for example, asking, "Did you know that you use more energy than 90 percent of your neighbors?"). Because people tend to ignore letters written in bureaucratese and fail to complete buggy computer forms, simplify the language and user interface. And since people tend to engage in "mental accounting," allow people to maintain separate bank accounts for food money, holiday money, and so on.[4]

Richard Thaler and Cass Sunstein call this type of design thinking "choice architecture." The idea is to design forms, programs, and policies that go with, rather than against, the grain of human psychology. Doing so does not restrict choices; rather, options are arranged and presented in ways that help people make day-to-day choices that are consistent with their long-term goals. In contrast with the hard incentives of classical economics, behavioral nudges are "soft" techniques for prompting desired behavior change.

Proponents of this approach, such as the Behavioural Insights Team and ideas42, argue that behavioral nudges should be part of policymakers' toolkits.[5] This article goes further and argues that the science of behavioral nudges should be part of the toolkit of mainstream predictive analytics as well. The story of a recent political campaign illustrates the idea.

YES, THEY DID

The 2012 US presidential campaign has been called "the first big data election."[6] Both the Romney and Obama campaigns employed sophisticated teams of data scientists charged with (among other things) building predictive models to optimize the efforts of volunteer campaign workers. The Obama campaign's strategy, related in Sasha Issenberg's book *The Victory Lab*, is instructive: The team's data scientists built, and continually updated, models prioritizing voters in terms of their estimated likelihood of being persuaded to vote for Obama. The strategy was judicious: One might naively design a model to simply identify likely Obama voters. But doing so would waste resources and potentially annoy many supporters already intending to vote for Obama. At the opposite extreme, directing voters to the doors of hard-core Romney supporters would be counterproductive. The smart strategy was to identify those voters most likely to *change their behavior* if visited by a campaign worker.[7]

So far this is a "*Moneyball* for political campaigns" story. A predictive model can weigh more factors—and do so more consistently, accurately, and economically—than the unaided judgment of overstretched campaign workers. Executing an analytics-based strategy enabled the campaign to derive considerably more benefit from its volunteers' time.

But the story does not end here. The Obama campaign was distinctive in combining the use of predictive analytics with outreach tactics motivated by behavioral science. Consider three examples: First, campaign workers would ask voters to fill out and sign "commitment cards" adorned with a photograph of Barack Obama. This tactic was motivated by psychological research indicating that people are more likely to follow through on actions that they have committed to. Second, volunteers would also ask people to articulate a specific plan to vote, even down to the specific time of day they would go to the polls. This tactic reflected psychological research suggesting that forming even a simple plan increases the likelihood that people will follow through. Third, campaign workers invoked social norms, informing would-be voters of their neighbors' intentions to vote.[8]

The Obama campaign's combined use of predictive models and behavioral nudge tactics suggests a general way to enhance the power of business analytics applications in a variety of domains.[9] It is an inescapable fact that no model will provide benefits unless appropriately acted upon. Regardless of application, the implementation must be successful in two distinct senses: First, the model must be converted into a functioning piece of software that gathers and combines data elements and produces a useful prediction or indication with suitably short turn-around time.[10] Second, end users must be trained to understand, accept, and appropriately act upon the indication.

In many cases, determining the appropriate action is, at least in principle, relatively straightforward. For example, if an analysis singles out a highly talented yet underpaid baseball player, scout him. If an actuarial model indicates that a policy-holder is a risky driver, set his or her rates accordingly. If an emergency room triage model indicates a high risk of heart attack, send the patient to intensive care. But in many other situations, exemplified by the challenge of getting out the vote, a predictive model can at best point the end user in the right direction. It cannot suggest how to prompt the desired behavior change.

I call this challenge "the last-mile problem." The suggestion is that just as data analytics brings scientific rigor to the process of estimating an unknown quantity or making a prediction, employing behavioral nudge tactics can bring scientific rigor to the (largely judgment-driven) process of deciding how to prompt behavior change in the individual identified by a model. When the ultimate goal is behavior change, predictive analytics and the science of behavioral nudges can serve as two parts of a greater, more effective whole.

PUSH THE WORST, NUDGE THE REST

Once one starts thinking along these lines, other promising applications come to mind in a variety of domains, including public sector services, behavioral health, insurance, risk management, and fraud detection.

Supporting child support

Several US states either have implemented or intend to implement predictive models designed to help child support enforcement officers identify noncustodial parents at risk of falling behind on their child support payments.[11] The goal is to enable child support officers to focus more on prevention and hopefully less on enforcement. The application is considered a success, but perhaps more could be done to achieve the ultimate goal of ensuring timely child support payments.[12] The Obama campaign's use of commitment cards suggests a similar approach here. For example, noncustodial parents identified as at-risk by a model could be encouraged to fill out commitment cards, perhaps adorned with their children's photographs. In addition, behavioral nudge principles could be used to design financial coaching efforts.[13]

Taking the logic a step further, the model could also be used to identify more moderate risks, perhaps not in immediate need of live visits, who might benefit from outreach letters. Various nudge tactics could be used in the design of such letters. For example the letters could address the parent by name, be written in colloquial and forthright language, and perhaps include details specific to the parent's situation. Evidence from behavioral nudge field experiments in other applications even suggests that printing such letters on colored paper increases the likelihood that they will be read and acted upon. There is no way of knowing in advance which (if any) combination of tactics would prove effective.[14] But randomized control trials (RCTs) could be used to field-test such letters on treatment and control groups.

Several US states either have implemented or intend to implement predictive models designed to help child support enforcement officers identify noncustodial parents at risk of falling behind on their child support payments. The goal is to enable child support officers to focus more on prevention and hopefully less on enforcement.

The general logic common to the child support and many similar applications is to use models to deploy one's limited workforce to visit and hopefully ameliorate the highest-risk cases. Nudge tactics could help the case worker most effectively prompt the desired behavior change. Other nudge tactics could be investigated as a low-cost, low-touch way of prompting some medium- to high-risk cases to "self-cure." The two complementary approaches refine the intuitive process case workers go through each day when deciding whom to contact and with what message. Essentially the same combined predictive model/behavioral nudge strategy could similarly be explored in workplace safety inspections, patient safety, child welfare outreach, and other environments.

Let's keep ourselves honest

Similar ideas can motivate next-generation statistical fraud detection efforts. Fraud detection is among the most difficult data analytics applications because (among other reasons) it is often the case that not all instances of fraud have been flagged as such in historical databases. Furthermore, fraud itself can be an inherently ambiguous concept. For example, much automobile insurance fraud takes the form of opportunistic embellishment or exaggeration rather than premeditated schemes. Such fraud is often referred to as "soft fraud." Fraud "suspicion score" models inevitably produce a large proportion of ambiguous indications and false-positives.

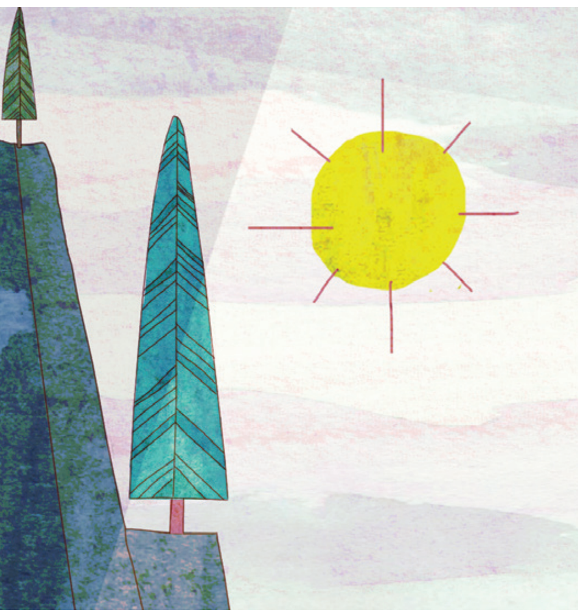

Acting upon a fraud suspicion score can therefore be a subtler task than acting on, for example, child welfare or safety inspection predictive model indications.

Behavioral nudge tactics offer a "soft touch" approach that is well suited to the ambiguous nature of much fraud detection work. Judiciously worded letters could be crafted to achieve a variety of fraud mitigation effects. First, letters that include specific details about the claim and also remind the claimant of the company's fraud detection policies could achieve a "sentinel effect" that helps ward off further exaggeration or embellishment. If appropriate, letters could inform individuals of a "lottery" approach to random fraud investigation. Such an approach is consistent with two well-established psychological facts: People are averse to loss, and they tend to exaggerate small probabilities, particularly when the size of the associated gain or loss is large and comes easily to mind.[15]

A second relevant finding of behavioral economics is that people's actions are influenced not only by external (classically economic) punishments and rewards but by an *internal* reward system as well. For example, social norms exert material influences on our behaviors: People tend to use less energy and fewer hotel towels when informed that people like them tend to conserve. Physician report cards have been found to promote patient safety because they prompt physicians to compare their professional conduct to that of their peers and trigger such internal, noneconomic rewards as professional pride and the pleasure of helping others.[16] Similarly, activating the internal rewards system of a potential (soft) fraudster might provide an effective way of acting upon ambiguous signals from a statistical fraud detection algorithm. Behavioral economic research suggests that a small amount of cheating flies beneath the radar of people's internal reward mechanism. Dan Ariely calls this a "personal fudge factor." In addition, people simply deceive themselves when committing mildly dishonest acts.[17] Such nudge tactics as priming people to think about ethical codes of honor and contrasting their actual behavior with their (honest) self-images are noneconomic levers for promoting honest behavior.

Physical and financial health

Keeping health care utilization under control is a major goal of both predictive analytics and behavioral economics. In the United States alone, health expenditures topped $2.8 trillion in 2013, accounting for more than 18 percent of the country's GDP.[18] Furthermore, chronic diseases, exacerbated by unhealthy lifestyles, account for a startling proportion of this spending.[19] All of this makes health spending an ideal domain for applying predictive analytics: The issue is one of utmost urgency; unhealthy behaviors can be predicted using behavioral and lifestyle data; and models are capable of singling out small fractions of individuals accounting for a disproportionate share of utilization. For example, in his *New Yorker* article "The Hot Spotters," Atul Gawande described a health utilization study in which the highest health utilizers, 5 percent of the total, accounted for approximately 60 percent of the total spending of the population.[20]

Of course, identifying current—and future—high utilizers is crucial, but it's not the ultimate goal. Models can point us toward those who eat poorly, don't exercise enough, or are unlikely to stick to their medical treatments, but they do not instruct us about which interventions prompt the needed behavior change. Once again, behavioral economics is the natural framework to scientifically attack the last-mile problem of going from predictive model indication to the desired action.

In contrast with the building inspection and fraud detection examples, it is unlikely that purely economic incentives are sufficient to change the behavior of the very worst risks. In this context, the worst risks are those with multiple chronic diseases. A promising behavioral strategy, described in Gawande's article, is assigning health coaches to high-utilizing patients who need personalized help to manage their health.

Such coaches are selected based more on attitudes and cultural affinities with the patient than on medical training. Indeed, this suggests a further data science application: Use the sort of analytical approaches employed by online matchmaking services to hire and match would-be health coaches to patients based on personality and cultural characteristics. Gawande recounts the story of a diabetic, obese woman who had suffered three heart attacks. After working with a health coach, she managed to leave her wheelchair and began attending yoga classes. Asked why she would listen to her health coach and not her husband, the patient replied, "Because she talks like my mother."[21] In behavioral economics, "messenger effects" refer to the tendency to be influenced by a message's source, not just its content.

Closely analogous ideas can be pursued in such arenas as financial health and back-to-work programs. For example, a recent pilot study of a "financial health check" program reported a 21 percent higher savings rate compared with a control group. The health check was a one-hour coaching session designed to address such behavioral bottlenecks as forgetfulness, burdensome paperwork, lack of self-control, and letting short-term pleasures trump long-term goals.[22] Predictive models (similar to credit-scoring models) could be built to identify in advance those who most need such services. As with the child support case study, this would enable prevention as well as remediation. But behavioral science is needed to identify the most effective ways to provide the coaching.

> Viewed through a traditional lens, the telematics black-box devices increasingly used by insurance companies are the ultimate actuarial segmentation machine. They can capture fine-grained details about how much or abruptly we drive, accelerate, change lanes, turn corners, and so on.

IN 3D: DATA MEETS DIGITAL MEETS DESIGN

The discussion so far has focused on behavioral nudge tactics as a way of bridging the gap between predictive model *indications* and the appropriate actions. But nudge tactics are more commonly viewed in the context of a more everyday last-mile problem: the gap between people's long-term *intentions* and their everyday actions. An overarching theme of behavioral economics is that people's short-term actions are often at odds with their own long-term goals. We suffer inertia and information overload when presented with too many choices; we regularly put off until "tomorrow" following up on important diet, exercise, and financial resolutions; and we repeatedly fool ourselves into believing that it's okay to indulge in something risky "just this once."[23]

Just as behavioral science can help overcome the last-mile problem of predictive analytics, perhaps data science can assist with the last-mile problem of behavioral economics: In certain contexts, useful nudges can take the form of digitally delivered, analytically constructed "data products."

Consider that much of what we call "big data" is in fact *behavioral* data. These are the "digital breadcrumbs" that we leave behind as we go about our daily activities in an increasingly digitally mediated world.[24] They are captured by our computers, televisions, e-readers, smartphones, self-tracking devices, and the data-capturing "black boxes" installed in cars for insurance rating. If we refocus our attention from data capture to data delivery, we can envision "data products," delivered via apps on digital devices, designed to help us follow through on our intentions. Behavioral economics supplies some of the design thinking needed for such innovations. The following examples illustrate this idea.

Driving behavior change

Viewed through a traditional lens, the telematics black-box devices increasingly used by insurance companies are the ultimate actuarial segmentation machine. They can capture fine-grained details about how much or abruptly we drive, accelerate, change lanes, turn corners, and so on.

But in our digitally mediated world of big data, it is possible to consider new business models and product offerings. After all, if the data are useful in helping an insurance company understand a driver's risk, they could also be used to help the driver to better understand—and control—his or her own risk. Telematics data could inform periodic "report cards" that contain both judiciously selected statistics and contextualized composite risk scores informing people of their performance. Such reports could help risky drivers better understand (and hopefully improve) their behavior, help beginner drivers learn and improve, and help older drivers safely remain behind the wheel longer.

Doing so would be a way of countering the "overconfidence bias" that famously clouds drivers' judgment. Most drivers rate themselves as better than average, though a moment's reflection indicates this couldn't possibly be true.[25] Peer effects could be another effective nudge tactic. For example, being informed that one's driving is riskier than a large majority of one's peers would likely prompt safer driving.

It is worth noting that employing nudge tactics does not preclude using classical economic incentives to promote safer driving. For example, the UK insurer ingenie uses black-box data to calculate risk scores, which in turn are reported via mobile apps back to drivers so that they can better understand and improve their driving behavior. In addition, the driver's insurance premium is periodically recalculated to reflect the driver's improved (or worsened) performance.[26]

Resisting the siren song

Self-tracking devices are the health and wellness equivalent of telematics black boxes. Here again, it is plausible that appropriately contextualized periodic feedback reports ("you are in the bottom quartile of your peer group for burning calories") could nudge people for the better.

Perhaps we can take the behavioral design thinking a step further. In behavioral economics, "present bias" refers to short-term desires preventing us from achieving long-term goals. Figuratively speaking, our "present selves" manifest different preferences than our "future selves" do. An example from classical mythology is recounted in Homer's *Odyssey*: Before visiting the Sirens' island, Odysseus instructed his crew to tie him to a mast. Odysseus knew that in the heat of the moment, his present self would give in to the Sirens' tempting song.[27]

It is reasonable to anticipate better results when the two approaches are treated as complementary and applied in tandem. Behavioral science principles should be part of the data scientist's toolkit, and vice versa.

Behavioral economics offers a somewhat less dramatic device: the commitment contract. For example, suppose Jim wants to visit the swimming pool at least 15 times over the coming month but also knows that on any given day he will give in to short-term laziness and put it off until "tomorrow." At the beginning of the month, Jim can sign a contract committing him to donate, say, a thousand dollars to charity—or even to a political organization he loathes—in the event that he fails to keep his commitment to his future self.

Commitment contracts digitally implemented on mobile app devices offer an intriguing application for self-tracking devices that goes beyond real-time monitoring or feedback reports. Data from such devices could automatically flow into apps preprogramed with commitment contracts that reflect diet, exercise, or savings goals. Oscar Wilde famously said that "the only way to get rid of temptation is to yield to it." Perhaps. But Wilde wrote long before behavioral economics.[28]

Washing our hands of safety noncompliance

I yield to the temptation to give one final example of data-fueled, digitally implemented, and behaviorally designed innovation. A striking finding of evidence-based medicine is that nearly 100,000 people die each year in the United States alone from preventable hospital infections. A large number of lives could therefore be saved by prompting health care workers to wash their hands for the prescribed length of time. Medical checklists of the sort popularized in Atul Gawande's *The Checklist Manifesto* are one behavioral strategy for improving health care workers' compliance with protocol.[29] An innovative digital soap dispenser suggests the possibilities offered by complementary, data-enabled strategies.

The DebMed Group Monitoring System is an electronic soap dispenser equipped with a computer chip that records how often the members of different hospital wards wash their hands.[30] Each ward's actual hand washing is compared with statistical estimates of appropriate hand-washing that are based on World Health Organization guidelines as well as data specific to the ward. Periodic feedback messages report not specific individuals' compliance but the compliance of the entire group. This design element exemplifies what MIT Media Lab's Sandy Pentland calls "social physics": Social influences help shape individual behavior.[31] In such cases, few individuals will want to be the one who ruins the performance of their team. Administered with the appropriate behavioral design thinking, an ounce of data can be worth a pound of cure.

ACT NOW

In predictive analytics applications, the last-mile problem of prompting behavior change tends to be left to the professional judgment of the model's end user (child support enforcement officer, safety inspector, fraud investigator, and so on). On the other hand, behavioral nudge applications are often one-size-fits-all affairs applied to entire populations rather than analytically identified sub-segments. It is reasonable to anticipate better results when the two approaches are treated as complementary and applied in tandem. Behavioral science principles should be part of the data scientist's toolkit, and vice versa.

Behavioral considerations are also not far from the surface of discussions of "big data." As we have noted, much of big data is, after all, *behavioral* data continually gathered by digital devices as we go about our daily activities. Such data are controversial for reasons not limited to a basic concern for privacy. Behavioral data generated in one context can be repurposed for use in other contexts to infer preferences, attitudes, and psychological traits with accuracy that many find unsettling at best, Orwellian at worst.

In addition, many object to the idea of using psychology to nudge people's behavior on the grounds that it is manipulative or a form of social engineering. These concerns are crucial and not to be swept under the carpet. At the same time, it is possible to view both behavioral data and behavioral nudge science as tools that can be used in either socially useful or socially useless ways. Hopefully the examples discussed here illustrate the former sort of applications. There is no unique bright line separating the usefully personalized from the creepily personal, or the inspired use of social physics ideas from hubristic social engineering. Principles of social responsibility are therefore best considered not merely topics for public and regulatory debate but inherent to the process of planning and doing both data science and behavioral science.

Behavioral design thinking suggests one path to "doing well by doing good" in the era of big data and cloud computing.[32] The idea is for data-driven decision making to be more of a two-way street. Traditionally, large companies and governments have gathered data about individuals in order to more effectively target market and actuarially segment, treat, or investigate them, as their business models demand. In a world of high-velocity data, cloud computing, and digital devices, it is increasingly practical to "give back" by offering data products that enable individuals to better understand their own preferences, risk profiles, health needs, financial status, and so on. The enlightened use of choice architecture principles in the design of such products will result in devices to help our present selves make the choices and take the actions that our future selves will be happy with. DR

James Guszcza *is the US chief data scientist for Deloitte Consulting LLP.*

Endnotes

1. See the edge.org conversation “On Kahneman,” March 27, 2014, <http://edge.org/conversation/on-kahneman>.
2. See “Conclusions” chapter of Daniel Kahneman, *Thinking, Fast and Slow* (Farrar, Straus and Giroux), 2011.
3. I use the terms “behavioral insights” and “nudge tactics” interchangeably to denote the application of the “choice architecture” principles outlined in Richard Thaler and Cass Sunstein’s book *Nudge: Improving Decisions about Health, Wealth, and Happiness* (New Haven, CT: Yale University Press, 2008). Many of these ideas were in turn anticipated in the fall 2003 *University of Chicago Law Review* article “Libertarian paternalism is not an oxymoron,” by Sunstein and Thaler <http://faculty.chicagobooth.edu/Richard.Thaler/research/pdf/LIbpatLaw.pdf>. To the best of my knowledge, the term “behavioral insights” originates from the UK Cabinet Office’s Behavioural Insights Team <http://www.behaviouralinsights.co.uk/>, led by David Halpern and advised by Richard Thaler.
4. Richard Thaler’s *Nudge* blog reports such a mental accounting scheme by ING Direct; see Thaler and Sunstein, “Nudge in the classroom part II,” *Nudge*, June 10, 2010, <http://nudges.org/tag/mental-accounting/page/2/>. A pioneer of using behavioral science to reduce power consumption is Opower <http://www.thermostat.opower.com/>. See the discussion of Opower in William D. Eggers and Paul Macmillan, “A billion to one,” *Deloitte Review* 16, January 2015. *EAST: Four simple ways to apply behavioural insights* (April 11, 2014, <http://behaviouralinsights.co.uk/publications/east-four-simple-ways-apply-behavioural-insights>) by the Behavioural Insights Team reports that six months after employees in a sample of large firms were automatically enrolled in pension schemes (that is, the default was changed from “opt in” to “opt out”), participation rates rose from 61 to 83 percent.
5. For information about the Behavioural Insights Team and ideas42 see the organizations’ respective websites: <http://www.behaviouralinsights.co.uk/about-us> and <http://www.ideas42.org/about/background/>.
6. Eric Hellweg, “2012: The first big data election,” *HBR Blog Network*, November 13, 2012, <http://blogs.hbr.org/2012/11/2012-the-first-big-data-electi/>.
7. The analytical technique is known as “uplift modeling,” and is also used in marketing applications. See, for example, Eric Siegel, “The real story behind Obama’s campaign victory,” *Fiscal Times*, January 21, 2013, <http://www.thefiscaltimes.com/Articles/2013/01/21/The-Real-Story-Behind-Obamas-Election-Victory>.
8. Benedict Carey, “Academic ‘dream team’ helped Obama’s effort,” *The New York Times*, November 12, 2012. Note that I am highlighting the Obama campaign’s behavioral nudge tactics that can be viewed as complementary to predictive model-guided voter outreach. The article discusses other behavioral science-motivated tactics unconnected with the use of predictive models. For example, the campaign’s Consortium of Behavioral Scientists suggested ideas on how to counter such false rumors as the notion that Barack Obama is a Muslim: Don’t deny the charge, since doing so would not remove the association that has been planted. Instead, affirm a competing notion.
9. I do not suggest that the Obama campaign deliberately set out to pursue a unified predictive modeling/behavioral nudge strategy. It is possible that different parts of the campaign independently pursued different strategies that proved complementary. This narrative weaves together threads from different news articles, each highlighting a different aspect of the campaign’s scientific strategies. However, I do suggest that in applications where behavior change is the desired result, the story of the Obama campaign helps motivate viewing behavioral nudge methods as part of the business analytics toolkit.
10. This is often easier said than done. Alex Madrigal’s November 2012 *Atlantic* article “When the nerds go marching in” <http://www.theatlantic.com/technology/archive/2012/11/when-the-nerds-go-marching-in/265325/> recounts the dramatic story of the team that managed and continually stress-tested the Obama campaign’s IT infrastructure. The team’s efforts were crucial to the success of the campaign’s predictive modeling strategy. This story is consistent with our (admittedly more mundane) experience that data cleansing, model implementation, and various model risk mitigation activities often account for the bulk of the work in major predictive modeling projects. This is a basic “no free lunch” principle of business analytics. Interested readers might also wish to consult the personal website of Harper Reed <www.harperreed.com>, the campaign’s chief technology officer who was profiled in Madrigal’s article.
11. Daniel Richard et al., “Leveraging advanced analytics: The evolution of Pennsylvania’s child support enforcement program,” *Policy & Practice*, February 1, 2014, <http://www.readperiodicals.com/201402/3280496061.html>.
12. For example in 2012 the application was recognized by a “bright idea award” by the Harvard Kennedy School’s Ash Center for Democratic Governance and Innovation. <http://www.ash.harvard.edu/ash/Home/News-Events/Press-Releases/Innovations/Harvard-s-Ash-Center-Announces-111-Bright-Ideas-in-Government/2012-Bright-Ideas>.
13. For example, the Behavioral Interventions to Advance Self-Sufficiency (BIAS) program, funded by the US Department of Health and Human Services, used behavioral design thinking to eliminate hassle factors and devise coaching programs to increase the rate of incarcerated noncustodial parents applying for modifications of their child support orders. See Mary Farrell et al., *Taking the first step: Using behavioral economics to help incarcerated parents apply for child support order modifications*, Office of Planning, Research, and Evaluation, August 2014, <http://www.acf.hhs.gov/sites/default/files/opre/bias_texas_executive_summary_2014.pdf>.
14. For example, the above-mentioned BIAS child support program used blue paper in letters sent to prisoners. Peer effects were famously used by the UK Behavioural Insights team to increase tax collections. Using a randomized control trial, the team demonstrated that adding a line informing that most people pay their taxes on time led to a 5 percent increase in tax collection. Adding a (behavioral science-motivated) line to a letter therefore yielded millions of pounds of additional tax collections. See *EAST: Four simple ways to apply behavioural insights* by the Behavioural Insights Team. < http://behaviouralinsights.co.uk/publications/east-four-simple-ways-apply-behavioural-insights>.

15. Random inspections and activating a "sentinel effect" are certainly not new ideas in fraud prevention, but behavioral science helps explain their usefulness. The Behavioural Insights Team's research highlights the idea of including very specific information in communications: When letters sent to nonpayers of a road tax included photographs of the offending vehicle, payments rose from 40 to 49 percent. See Behavioural Insights Team, EAST.

16. "How physician report cards can improve health care," Knowledge@Wharton, August 28, 2014, <http://knowledge.wharton.upenn.edu/article/how-physician-report-cards-can-improve-health-care/>. The blog discusses findings of the health economist Jonathan Kolstad.

17. Nina Mazar and Dan Ariely, "Dishonesty in everyday life and its policy implications," *Journal of Public Policy & Marketing*, 2006, <http://people.duke.edu/~dandan/Papers/PI/dishonesty.pdf>; Dan Ariely, "Our buggy moral code," filmed February 2009. TED Talks video, 16:17, <http://www.ted.com/talks/dan_ariely_on_our_buggy_moral_code?language=en>. Here, Ariely calls people's tendency to cheat a little bit without changing their honest self-images a "personal fudge factor." For specific examples of randomized control trials of behavioral nudges for prompting honesty, using social norms, and so on, see Behavioural Insights Team, *Applying behavioural insights to reduce fraud, error, and debt*, UK Cabinet Office, February 2012, <https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/60539/BIT_FraudErrorDebt_accessible.pdf>.

18. Ian Duncan, "Predictive modeling of healthcare costs using hierarchical regression," University of California at Santa Barbara, 2014.

19. For example, the Centers for Disease Control estimated that chronic diseases accounted for 75 percent of 2005 health care spending in the United States. See Johnson & Johnson, "Prevention and wellness," <http://www.jnj.com/sites/default/files/pdf/prevention-and-wellness.pdf>, accessed October 29, 2014.

20. Atul Gawande, "The hot spotters," *The New Yorker*, January 24, 2011, <http://www.newyorker.com/magazine/2011/01/24/the-hot-spotters>. A more recent presentation reports that the 5 percent of the highest utilizers account for 50 percent of utilization; see Michael Breslow, "Corporate and executive spotlight: Michael Breslow, Philips Healthcare," filmed at 2012 mHealth Summit; *Healthcare IT News* video, 15:58; <https://www.youtube.com/watch?v=_zY88SfsmGg>, posted December 3, 2013.

21. Ibid.

22. Antoinette Schoar and Piyush Tantia, "The financial health check: A behavioral approach to financial coaching," *ideas42*, March 2014, <http://www.ideas42.org/publication/view/the-financial-health-check-a-behavioral-approach-to-financial-coaching/>. The above-referenced BIAS child support study can be viewed along similar lines.

23. "Report: Texting while driving okay if you look up every couple seconds," *The Onion*, May 27, 2013, <http://www.theonion.com/articles/report-texting-while-driving-okay-if-you-look-up-e,32583/>.

24. James Guszcza and Bryan Richardson, "Two dogmas of big data: Understanding the power of analytics for predicting human behavior," *Deloitte Review* Issue 15, July 18, 2014, <http://dupress.com/articles/behavioral-data-driven-decision-making/>.

25. This phenomenon from behavioral economics, the tendency to overestimate one's abilities, is also known as the "Lake Wobegon Effect" in homage to Garrison Keillor's radio show *The Prairie Home Companion*. The show is set in the fictional town of Lake Wobegon, "where all the women are strong, all the men are good looking, and all the children are above average."

26. ingenie, "How it works" (video), <http://www.ingenie.com/how-it-works>, accessed October 29, 2014. Our previous *Deloitte Review,* Issue 14, article "The personalized and the personal: Socially responsible innovation through big data," published January 2014, <http://dupress.com/articles/dr14-personalized-and-personal/> explored the idea of behavioral science-motivated and digitally delivered driving feedback reports.

27. For the record, the first draft of this article is being written several days after the deadline. In his preface to *Behavioral Economics and Policy Design: Lessons from Singapore* (World Scientific, October 2011), Donald Low brilliantly invokes another modern illustration, from the comedian Jerry Seinfeld: "I never get enough sleep. I stay up late at night because I'm Night Guy. Night Guy wants to stay up late. 'What about getting up after five hours' sleep?' Oh, that's Morning Guy's problem. That's not my problem. I'm Night Guy—I stay up as late as I want. So you get up in the morning . . . you're exhausted, you're groggy. 'Oh I hate that Night Guy!' Night Guy always screws Morning Guy. There's nothing Morning Guy can do. The only thing Morning Guy can do is try to oversleep often enough so that Day Guy loses his job and Night Guy has no money to go out anymore."

28. stickK (http://www.stickk.com/) is a service founded by the Yale professors Dean Karlan (a prominent researcher of commitment devices) and Ian Ayres (well known as the author of *Super Crunchers*). Interestingly, adding a social element dramatically increased the effectiveness of the commitment contracts administered by stickK. When users name a referee to keep themselves honest, the success rate is 80 percent. Without a referee, it is only 29 percent. See Michael Blanding, *The business of behavioral economics*, Harvard Business School, August 11, 2014, <http://hbswk.hbs.edu/item/7588.html>. Beeminder <https://www.beeminder.com/> is an app-based commitment contract service. Parmy Olson's *Forbes* article "Wearable tech is plugging into health insurance" (June 19, 2014; <http://www.forbes.com/sites/parmyolson/2014/06/19/wearable-tech-health-insurance/>) reports that BP, which self-insures for employee health, offers its 14,000 employees free self-tracking devices. Employees that take over 1 million steps qualify for lower insurance premiums. One can envision such data feeding commitment contracts as well.

29. Although medical checklists are not always characterized in behavioral nudge terms, they can be viewed as an example of what behavioral economists call "heuristics": simple rules of thumb that enable an individual to make the right decision without the need to process large amounts of information. The daily diet "food pyramid" is another familiar example. See Antoinette Schoar and Saugato Datta, "The power of heuristics," *ideas42*, January 2014, <http://www.ideas42.org/publication/view/the-power-of-heuristics/>.

30. "First, wash your hands," the *Economist*, July 17, 2013, <http://www.economist.com/news/technology-quarterly/21584441-biomedicine-smart-antiseptic-dispensers-promise-save-lives-subtly-encouraging>.

31. See Alex Pentland, *Social Physics: How Good Ideas Spread—The Lessons from a New Science*, (Penguin, 2014).

32. James Guszcza, David Schweidel, and Shantanu Dutta, "The personalized and the personal," *Deloitte Review*, Issue 14, January 17, 2014, <http://dupress.com/articles/dr14-personalized-and-personal/>; Eggers and Macmillan, "A billion to one," *Deloitte Review,* Issue 16, January 19, 2015.