


How AI Affects Our Sense of Self

by Gizem Yalcin and Stefano Puntoni

From the Magazine (September–October 2023)




Anna Parini

Summary. The increasing frequency of interactions we have with AI and automated technologies means it is vital to understand how those things make people feel about themselves. Why? Because how people feel about themselves affects a wide range of... [more](#)

If you ever took a marketing course, you may remember the famous case from the 1950s about General Mills' launch of Betty Crocker cake mixes, which called for simply adding water, mixing, and baking. Despite the product's excellent performance, sales were initially disappointing. That was puzzling until managers figured out the problem: The mix made baking too easy, and buyers felt they were somehow cheating when they used it. On the basis of that insight, the company

removed egg powder from the ingredients and asked customers to crack an egg and beat it into the mix. That small change made those bakers feel better about themselves and so boosted sales. Today, 70 years later, most cake mixes still require users to add an egg.



To hear more, download [Noa app](#) for iPhone or Android.

We can take a lesson from that story today. As companies increasingly embrace automated products and services, they need to understand how those things make their customers feel about themselves. To date, however, managers and academics have usually focused on something quite different: understanding what customers *think* about those things. Researchers have been studying, for example, whether people prefer artificial intelligence over humans (they don't), how moral or fair AI is perceived to be (not very), and the tasks for which people are likely to resist the adoption of automation (those that are less quantifiable and more open to interpretation).

All that is important to consider. But now that people are starting to interact frequently and meaningfully with AI and automated technologies, both at and outside work, it's time to focus on the emotions those technologies evoke. That subject is psychological terra incognita, and exploring it will be critical for businesses, because it affects a wide range of success factors, including sales, customer loyalty, word-of-mouth referrals, employee satisfaction, and work performance.

We have been studying people's reactions to autonomous technology and the psychological barriers to adopting it for more than seven years. In this article, drawing on recent research from our lab and reviewing real-life examples, we look at the

psychological effects we've observed in three areas that have important ramifications for managerial decision-making: (1) services and business-process design, (2) product design, and (3) communication. After surveying the research and examples, we offer some practical guidance for how best to use AI-driven and automated technologies to serve customers, support employees, and advance the interests of organizations.

Services and Business-Process Design

Today AI and automated technologies are embedded in a wide range of services and business processes that directly or indirectly affect consumers and employees. Upstart, for example, uses AI to decide which applicants to lend to, and Monster and Unilever use it to assess job candidates' potential. GEICO's DriveEasy program uses it to evaluate customers' driving skills and determine car-insurance premiums, while IBM and Lattice help businesses adopt AI-based performance-feedback processes, which have an impact on promotion and layoff decisions.

Given this trend, we need to ask: How do people react to decisions and feedback from AI and automated technologies? And how can businesses best incorporate them into their services and business processes to maximize customer and employee satisfaction?

Let's start with the first question. Together with Sarah Lim of the University of Illinois Urbana-Champaign and Stijn M.J. van Osselaer of Cornell University, we've recently examined situations in which the applications that people made to companies (perhaps for a loan or some benefits) were either accepted or rejected. In 10 studies, which involved a total of more than 5,000 participants, we found that in the case of acceptance, they reacted differently to decisions made by AI than to those made by humans.

Their reactions were psychologically revealing: Study participants whose requests were granted by a person felt more joy than did those whose requests were granted by AI, even though the

outcome was identical. Why? Because the latter felt reduced to a number and thought they couldn't take as much credit for their success. When their requests were turned down, however, participants felt the same way whether the rejection was by a person or by AI. In both cases, and to the same degree, they tended to blame the decision-maker for their failure rather than themselves.

The takeaway here is clear: People's feelings about themselves may differ depending on who or what evaluates them, and that has important consequences for business.

Study participants whose requests were granted by a person felt more joy than did those whose requests were granted by AI, even though the outcome was identical.

Consider the results of one of our studies, in which we asked people to imagine applying for a bank loan. Half the participants were told that a loan algorithm would evaluate their applications and the other half that a loan officer would evaluate them. Later half the participants in each group were told that their application had been approved and the other half that it had been denied.


Participants whose applications had been approved by an algorithm gave lower ratings to the bank and were less likely to recommend it to others than were people whose applications had been approved by a loan officer. But all the participants whose applications had been denied rated the bank similarly and felt the same degree of interest in recommending it to others.

We've observed this pattern in real-world contexts as well. For example, we asked workers who were part of an online labor platform to apply for membership on a select panel formed by a

research company. Half were told that AI would evaluate their applications, and the other half that a human employee would do so. Those who won admission to the panel through AI evaluated the research company less positively than those who won admission through an employee, but everybody who was rejected felt the same way about the company.

In short, when delivering good news about decisions and evaluations, companies can generate more-positive reactions among customers and employees if they rely on humans rather than on AI—but that effect disappears when they deliver bad news.

Most of the experienced leaders and managers we've interviewed in our research seemed unaware of these effects. In a survey we found that almost none of them could foresee the actual results. Executives will need to understand people's probable reactions if they hope to effectively engage customers and employees with new AI and automated technologies.



Anna Parini

Let's now turn to our second question: How can businesses integrate AI into their services and business processes to maximize customer and employee satisfaction? Our experimental findings offer some suggestions.

First, when AI or automated technologies are adopted for the purposes of evaluation and feedback, we recommend having some active human involvement in those processes and making that involvement clear to customers or employees. In one of our studies, we assessed how people rate a company when a human is only passively involved in evaluations (perhaps just monitoring algorithmic decisions). We compared that condition with one in

which a human is in charge of the evaluation process and one in which just an algorithm is, and we found that participants reacted positively only when human involvement was active.

Second, we recommend that managers be selective about the degree to which they rely on their (expensive) human workforce for decision-making. Because people tend to react the same way to negative news, whether it comes from a person or from AI, companies may not need the “human touch” to deliver it—even though that contradicts traditional managerial thinking. They should, however, consider using humans as often as possible to deliver good news.

Another research project also throws light on when humans can most effectively be deployed in business processes. Stefano worked with Armin Granulo of the Technical University of Munich and Christoph Fuchs of the University of Vienna to study symbolic products and services, which offer consumers more than just instrumental functionality. Such products and services embody abstract concepts that convey something about personality, beliefs, social-group membership, class status, or other intangibles. A few examples are tattoos, fashion jewelry, and varsity jackets. It’s important to remember, though, that a single product may have both physical and symbolic uses. Eyeglasses, for example, consist of lenses, which allow consumers to see (a physical use), and frames, which both hold the lenses in place (a physical use) and serve as a fashion accessory that may be central to self-expression (a symbolic use).

Because people tend to react the same way to negative news, whether it comes from a person or from AI, companies may not need the “human touch” to deliver it.

For that project—which consisted of four experiments using different product categories and involving more than 1,000 respondents—the authors compared consumers’ attitudes toward symbolic products that had been made by either automated technologies or humans. What they consistently found was that human labor adds distinctive value to symbolic products. In one of the experiments participants revealed that they preferred eyeglass lenses made by automated technology—presumably for their machine-based precision—but frames made by humans. In another study participants were more likely to purchase a poster designed by a human than one designed by AI.

These findings lead us to a third recommendation, which is that companies should carefully consider why customers are likely to buy their offerings—and whether they might add distinct value to the product by maintaining at least some human involvement in the production process, even if they intend to automate most of it.

Product Design

AI technologies and advanced automated features are integrated in many products and are transforming how we accomplish a variety of tasks in our personal lives: iRobot’s Roomba cleans your floors, Tesla’s Autopilot lets you enjoy the ride, Jura’s fully automatic coffee machine prepares your coffee from bean to cup and even cleans itself. Increasingly, too, people are working with AI-driven applications on the job. IBM’s Watson teams up with employees at many companies on a wide range of business tasks, including financial estimates and the management of marketing communication strategies; Adobe’s AI empowers designers and enhances their creative expression with Photoshop and other applications; and workers at Toyota operate highly automated tools and machinery. The recent advent of large language models and generative AI, such as OpenAI’s DALL-E and ChatGPT, is likely to accelerate these trends. How will our interactions with all

these automated technologies influence our sense of identity and accomplishment? And how will that influence the demand for products?

Our lab has explored how people react to automated products in the context of identity-based consumption, which helps people define who they are. Stefano worked on that project with Eugina Leung of Tulane University and Gabriele Paolacci of Erasmus University Rotterdam. In six studies and across various product categories, they found that people who identify with a particular activity, such as fishing, cooking, or driving, may experience automation as a threat to their identity, leading to reduced product adoption and lower product approval.

To learn more about this phenomenon, the authors conducted a study with Dutch participants and focused on cycling, an activity that is central to many Dutch people's sense of self. To temporarily make them identify even more strongly with cycling, half the participants were asked to write a short essay about the Dutch national passion for it, and the remaining half were asked to write an essay about the Dutch passion for flowers (the control condition). After that task they took part in an ostensibly unrelated study. The authors told them about a special offer from a bike shop and asked about their interest in adding a free automated feature to their own bikes: a rechargeable battery to assist with pedaling. Participants who had written about cycling were 20% less likely to accept the feature, even though it was free.

In another project, with the same team and Maria Cristina Cito of Bocconi University, the researchers examined a complementary issue: how people who are motivated by identity-relevant goals respond to companies' digitalization efforts. Across three main studies and five follow-up experiments, they found that symbolic products are adopted less often in digital form than they are in physical form. People simply can't express who they are as easily with digital products. Seeing the collected works of Shakespeare

on your Kindle is not nearly as powerful a way of validating your literary identity as seeing that same collection on your living room bookshelf.

Findings from these two projects indicate that when people identify with a certain product category, or when products help them express their beliefs and personalities, they sometimes resist any technological enhancement of those products. When that's the case, what should businesses do?

Many people's sense of self is rooted in their professional identity, and AI and automation can be perceived as undermining it.

First, we recommend that companies refrain from targeting identity-motivated consumers with fully automated products, and that when they do target such consumers, they focus on features or tasks that allow users to feel proud and involved. Consider the case of a bicycle-component manufacturer we worked with. Sometime earlier the company had introduced an expensive automatic gear-shifting device in the European market and had targeted cycling enthusiasts, who are more willing to pay for mechanical gadgets. But those consumers showed little interest in the device, because they felt that it would eliminate a central part of the cycling experience for them. If the company had marketed to commuters or casual bikers or had designed the feature in a way that gave riders a feeling of more control, it might have had greater success.

Second, we recommend that companies conduct market research to assess the extent to which automation risks triggering an identity threat.

Communication

With the adoption of AI and automated technologies, as with so much else, communication matters. In our research we've discovered three important ways that companies can optimize their communication strategies to minimize the risk of resistance or backlash.

First, companies that use AI interfaces to communicate with customers or employees should consider humanizing those interfaces. This is particularly important, we've found, in business processes that involve evaluation and decision-making. In one of our studies we tested whether adding humanlike features to AI would lead people to internalize positive news and rate the company more favorably. When we gave the AI a name (Sam), added an avatar, and made its interaction with people more conversational, they responded much as they would to a human employee. For companies that cannot employ humans for various reasons—such as a high volume of requests, limitations on time, or computational restrictions—this finding suggests that simply humanizing their AI might mitigate less-positive reactions to feedback or news from it.

Consider the case of a fintech company we worked with, which relies on AI technology to evaluate users' financial health. In its interactive and fully automated process, users fill out a questionnaire, the AI evaluates their answers, and the system produces an assessment of their financial health. At that point users are encouraged to click on a link for information about the company's services. In an attempt to boost consumer interest in those services, the company, working on behalf of a major global bank, created a chat format in which the AI engaged users with emotionally expressive cues such as emojis. When users received positive feedback about their financial health from the humanized AI versus the standard display format, they were more likely to click on the link and seek more information.

Second, we recommend that businesses modify how they communicate with customers and employees about their automated products. As noted, when people identify with a certain domain or activity, they sometimes resist automation if they feel that they can't attribute outcomes to their own skill or effort. But what if companies describe automated features not as replacing people but as complementing their skills?



Subscribe to our Monthly Newsletter

Leadership

The qualities of the most effective leaders are always changing. Read our latest.

[Sign Up](#)

Part of Stefano's project with Leung and Paolacci tested whether people's reaction to an automated product can be changed if it's framed in those terms. The authors created two advertisements in which they described an automated cooking machine in different ways: One ad read that the appliance would handle all the cooking steps "at the touch of a button," and the other that it would guide the cooking process and prepare the meal with the help of the user. Participants were randomly given one of them. Although the ads were for the same product, the results revealed that framing does indeed matter: When the appliance was described as allowing people to at least partly use their skills, identity-motivated consumers had more-positive attitudes toward it.

Although our studies were conducted primarily in the context of consumption activities, identity-related motivation is often important in the workplace as well. Many people's sense of self is rooted in their professional identity, and AI and automation can be perceived as undermining that identity if they threaten to

devalue skills, expertise, or status. Internal communication about their complementary potential will be crucial if companies hope to deploy them at scale.

...

Automated technologies are changing not only product and labor markets but also how the people using those technologies feel about themselves. Increasingly, companies will need to overcome psychological barriers by strategically designing their business processes and products to take human feelings into account and by employing well-thought-out communication strategies. In some cases automation may introduce the risk of reduced employee commitment or customer satisfaction, and companies will need to weigh its benefits against that risk. In such situations the appropriate question when considering a move to AI and automation is not “Can we?” but “Should we?”

A version of this article appeared in the September–October 2023 issue of *Harvard Business Review*.

GY

Gizem Yalcin is an assistant professor of marketing at the McCombs School of Business, University of Texas at Austin. She focuses on how AI and automation are changing the way consumers behave and feel about themselves, companies, and others.

SP

Stefano Puntoni is the Sebastian S. Kresge Professor of Marketing at The Wharton School and a codirector of the Wharton Impact of Technology Initiative. He investigates how AI and automation are changing consumption and society.

