

BEHAVIORAL ECONOMICS 2025 GUIDE



Introduction
Ernst Fehr

Guest Editorial
Suzanne Shu
Geoffrey Fisher
David Just
Stephen Shu

Editor
Alain Samson



The Behavioral Economics Guide 2025

ISSN 2398-2020

Author information:

Alain Samson (Editor)

Ernst Fehr (Introduction)

Suzanne B. Shu, Geoffrey Fisher, David R. Just and Stephen Shu (Guest Editorial)

Lucas Amherd (FehrAdvice), Olympia Arango Castro (Neovantas), Rasha Attar (Behavioral Science Group), Jörn Basel (Lucerne School of Business), Emma Bernardo Sampedro (Neovantas), Chiara Cappellini (Behavioral Science Group), Soraya Catoja (Allianz), Ann-Kathrin Crede (Allianz), Juande Rus Gutiérrez (Neovantas), Gerhard Fehr (FehrAdvice), Ailidh Finlayson (Behavioral Science Group), Emilio Guerra (Discovery Health), Ariam Hailemariam (ThinkPlace Ethiopia), Antonio Hanna-Amodio (Behavioral Science Group), Minh Hua, Sandra Lehmann (Vitality, South Africa), Nanxi Li (Vitality, South Africa), Jerry Luukkonen (Dectech), Mosima Mabunda (Vitality, South Africa), Miriam Makin (Dectech), Melaku Muleta (ThinkPlace Ethiopia), Christelle T. Ngnoumen (Voya Financial), Kim H. Nguyen (Voya Financial), Akinyi Odera (ThinkPlace Ethiopia), Deepak Patel (Vitality, South Africa), Carla Pool (Vitality, South Africa), Miguel Antonio Rodríguez Ramírez (Neovantas), Brett Rowland (Discovery Health), Natasha Silverman (Vitality, South Africa), Henry Stott (Dectech), Christian Weibel (Lucerne School of Business) and Malin Wenzke (Lucerne School of Business) (Contributing Authors)

Romneya Robertson (Graphic Design)

Povilas Daknys, SEMIHUMAN (Cover Design)

Published on www.behavioraleconomics.com by Behavioral Science Solutions Ltd

Copyright © by the authors

All rights reserved

Requests for permission to reproduce materials from this work should be sent to info@behavioraleconomics.com or directly to contributing authors

Suggested citation:

Samson, A. (Ed.)(2025). *The Behavioral Economics Guide 2025*.
<https://www.behavioraleconomics.com/be-guide/>.

Suggested citation for individual chapters/articles:

[Author(s)] (2025). [Chapter/Article Title]. In A. Samson (Ed.), *The Behavioral Economics Guide 2025* (pp. nnn-nnn). <https://www.behavioraleconomics.com/be-guide/>.

Contributions By



Cornell University



Universität
Zürich^{UZH}



Further Support By



CITY
ST GEORGE'S
UNIVERSITY OF LONDON



THE LONDON SCHOOL
OF ECONOMICS AND
POLITICAL SCIENCE



Reichman University
Raphael Recanati
International School



University College Dublin



University
of Exeter



University
of Glasgow



Penn
Master of Behavioral
and Decision Sciences



UNIVERSITY OF
SAN FRANCISCO
CHANGE THE WORLD FROM HERE



WARWICK
THE UNIVERSITY OF WARWICK

In Association With



Table of Contents

INTRODUCTION	V
Social Preferences, Cooperation, and Corporate Public Goods.....	VI
PART I – EDITORIAL	1
Advancing Behavioral Economics Research for Health and Finance.....	2
PART II – APPLICATIONS	15
A New Tool for Behavioral Scientists: Exploring the Potential of Synthetic Participants to Accelerate How We Study Attitudes and Behaviors.....	16
Symbolic Decision Support: Enhancing Financial Help-Seeking and Investment Decisions Using GenAI Imagery.....	26
One Nudge Is Not Enough: Exploring the Limits and Potentials of Behavioral Approaches for Sustainable Tourism Mobility.....	37
The Behaviorally-Informed Insurance Company: Insights From the Allianz Global Center for Behavioral Economics.....	48
Preventive Screening and Habit Formation in the Discovery-Vitality Shared-Value Health Promotion and Disease Prevention Program.....	55
From Points to Profits: Designing the Optimal Loyalty Programme.....	69
Can a Simple Message Make Customers More Honest? Insights From the Insurance Industry.....	83
What Lies Behind Our Behavior? An Approach to Understanding Human Decision-Making.....	92
The Integrated Ideation Model for Data-Use Culture (IIMDC): Tackling Behavioral Challenges Within Health Information Systems.....	101
CEO Briefing: Behavioral Science in the Workplace.....	110
PART III – RESOURCES	117
Postgraduate Programs.....	118
City St George’s, University of London.....	119
London School of Economics and Political Science.....	122
Reichman University.....	125
University College Dublin.....	128
University of Exeter.....	130
University of Glasgow.....	133
University of Pennsylvania.....	135
University of San Francisco.....	138
University of Warwick.....	140
Other Resources.....	143
APPENDIX	144
Author Profiles.....	145
Contributing Organizations.....	147

Acknowledgements

The editor would like to thank Ernst Fehr for writing the introduction to this edition, as well as Suzanne Shu, Geoffrey Fisher, David Just, and Stephen Shu for their guest editorial.

We are grateful for the support received from Allianz, Behavioral Science Group, Dectech, Discovery Vitality, FehrAdvice, Lucerne School of Business, Neovantas, ThinkPlace, Voya Financial, as well as City St George's, University of London, London School of Economics and Political Science, Reichman University, University College Dublin, University of Exeter, University of Glasgow, University of Pennsylvania, University of San Francisco, and University of Warwick.

Many thanks also go to Povilas Daknys for the cover design, as well as Rojhat Avsar, Charlotte Beckett, Raphaela Bruckdorfer, Vivek Choudhary, Darcie Dixon, Marcelo Fiszner, Tobias Heldt, Nika Kabiri, Annika Kleinsteuber, Myles Leighton, Xinxin Lyu, Lee McCauley, Martin Morgenstern, Sarah Murray, Heather Rupp, Rom Severin, Jargalan Tsogt, and Alexandros Tzikas for their valuable feedback

INTRODUCTION

Social Preferences, Cooperation, and Corporate Public Goods

ERNST FEHR

University of Zurich

Why Are Social Preferences Economically Relevant?

Up until the 1990s, traditional economic theory long assumed that individuals act primarily in their self-interest, striving to maximize their material payoffs without considering the well-being of others. This perspective has shaped models across markets, contracts, and organizational behavior. However, a vast body of research in behavioral economics has demonstrated that human behavior is driven not solely by material incentives, but also by social preferences such as concerns for fairness, reciprocity, and others' well-being (Rabin, 1993; Fehr & Schmidt, 1999; Andreoni & Miller, 2002; Fehr & Charness, forthcoming).

These preferences play a critical role in fostering cooperation in everyday economic and social life (Fehr & Charness, forthcoming), and they help explain why individuals often support colleagues, collaborate with team members, or resist actions they perceive as unfair, even when such behavior comes at a personal cost. In addition, in settings that involve shared goals, mutual dependence, or uncertain outcomes—such as group projects, workplace collaborations, or neighborhood initiatives—social preferences can help maintain trust and reduce the need for constant monitoring or formal enforcement. This is particularly important when contracts are incomplete, when outcomes depend on discretionary effort, or when trust is required to navigate complexity and ambiguity. In such contexts, the presence of individuals who value fairness and reciprocity can encourage cooperative behavior more broadly, in turn helping groups avoid breakdowns in collaboration.

In labor markets and organizations, social preferences shape how people respond to wage differences, managerial decisions, and workplace interactions. Fair treatment and mutual respect often lead to greater motivation and discretionary effort, while

perceived injustices—such as unexplained pay gaps or one-sided decision-making—can undermine morale and performance. When organizations acknowledge and account for these preferences, for instance by promoting transparent practices, equitable pay structures, or norms of mutual support, they are more likely to foster productive and stable work environments. However, social preferences are not a universal remedy: their effects depend on context and on how well they are understood and integrated into institutional design. Firms that overlook them may miss opportunities to strengthen internal cooperation, while those that account for them thoughtfully can often benefit from enhanced trust, lower turnover, and more resilient performance.

The purpose of this introduction is to translate insights from research on social preferences and their role in cooperation into practical implications for business organizations. Through real-world cases, I will illustrate how firms actively incorporating social preferences into how they design their organization and financial incentives can generate public goods, strengthen corporate culture, and enhance their overall performance. Furthermore, I will introduce the Willingness-Awareness Matrix as a practical framework for companies to systematically leverage the presence of social preferences for their strategies and improve their positioning in an increasingly complex and interconnected economic environment.

Social Preferences and the Creation of Corporate Public Goods

Social preferences are essential not only for fostering cooperation within organizations, but also for creating what can be described as “corporate public goods.” In economic terms, a public good is characterized by its non-excludability and non-rivalry, which means that once a public good is produced, it is available to all members of the group, without

diminishing its availability to others. In a corporate context, public goods include assets such as organizational culture, the company's or a department's reputation, shared knowledge, innovation climate, and team spirit, all of which are critical for a firm's long-term success and resilience.

Yet, despite their value, these goods are often fragile, in that they depend on voluntary contributions from employees who may not directly benefit from their efforts. In other words, while everyone gains from a strong culture or a high-trust environment, individuals may find it easier to withhold effort—leading to under-provision. This is precisely where social preferences—such as a preference for fairness and reciprocity—become essential, because they motivate individuals to contribute to the collective good, even in the absence of direct rewards.

Table 1 outlines several examples of such public goods and describes what makes them valuable, how they depend on social preferences, and why they can be difficult to sustain in practice. It also highlights the positive outcomes when these goods are present—and the risks when they are absent. This structured overview is meant to help practitioners recognize the invisible infrastructure that underpins long-term cooperation and corporate performance.

Firms that successfully integrate social preferences into their organizational structure can actively contribute to the generation of corporate public goods. To understand better what is at stake, consider two of the most fundamental ones: a firm's corporate culture and reputation, both of which are central to long-term performance, yet difficult to build and maintain.

Corporate culture refers to the shared values and behaviors that shape workplace interactions. It is a corporate public good because it is non-excludable—everyone in the organization is affected by it—and non-rivalrous—benefiting from a strong culture does not diminish others' ability to do the same. Further, culture creates value by aligning individual actions with collective goals, thereby enhancing collaboration and promoting a sense of belonging. However, it is difficult to sustain, since contributing to a strong culture requires ongoing effort, and the benefits are often diffuse, delayed, and shared. This consequently creates incentives to withhold effort—especially when others can carry

the burden. Social preferences, such as a preference for fairness or a concern for reciprocity, help resolve this problem, in that they motivate individuals to contribute to culture, even when doing so yields little personal reward—simply because it feels right or because others are doing the same. When too few employees act on such preferences, culture weakens and cooperation breaks down.

An example in this regard is Migros, one of Switzerland's largest retailers, whose corporate culture is deeply shaped by its cooperative ownership structure. Unlike shareholder-driven companies, Migros belongs to its customers, embedding values of fairness, mutual responsibility, and long-term societal benefit into everyday operations. This culture is non-excludable—every employee and customer is affected by it—and non-rivalrous, as one person's benefit does not diminish others' ability to benefit. Maintaining this culture requires individuals across the organization to voluntarily align their behavior with shared values, even when personal incentives to do so are weak or diffuse.

Migros actively reinforces this culture through practices like reinvesting profits into social and cultural projects (the Migros-Kulturprozent), promoting fair labor conditions, and advancing sustainability without immediate financial gain. These actions rely on employee and managers' willingness to contribute to collective goals, motivated by fairness and reciprocity rather than direct rewards.

Another interesting real-world example of corporate culture as a public good is provided by the Styria Media Group. After years of acquisitions, Styria operated through dozens of decentralized profit centers, each optimizing its own website traffic rather than contributing to collective goals like group-wide reach. Thus, if a client approached a particular profit center, this center had an active interest in keeping that client, even if another profit center in the company could provide greater value to the client, which clearly did not maximize the company's overall profit opportunities. In this setting, a culture of cooperation would have required that profit centers link their websites, such that the clients could easily find out which unit of the company provided the highest value, but the self-interest of the profit centers prevented the emergence of a cooperative culture whereby clients could transfer to a higher-value option. One

important reason for this was that the benefits of linking to other centers' websites were diffuse, while the immediate costs were center-specific so that individual units had strong incentives to defect. Without widespread voluntary cooperation based on fairness and reciprocity preferences, the shared cultural fabric of the group was unstable.

To address this issue, Styria Media implemented deliberate cultural interventions. They raised collective awareness of the group's shared interests, formulated explicit social norms favoring cooperation, introduced peer feedback mechanisms to create social accountability, and partially linked manager bonuses to overall group success. These measures strengthened social preferences and hence encouraged individuals to contribute to the collective culture, even when it involved personal effort. The result was a revitalized corporate culture in which cooperation became the norm rather than the exception.

Reputation is also a corporate public good because it is non-excludable—it reflects the organization as a whole, not just individual contributors—and non-rivalrous—benefiting from a good reputation does not diminish others' ability to do the same. Moreover, a strong reputation builds trust with clients, partners, and employees, and it helps attract opportunities across the organization. However, it is fragile, in that individuals often benefit from a good reputation without actively contributing to it, and maintaining it may involve personal restraint or effort, such as rejecting opportunistic gains. Social preferences, for example fairness concerns, motivate individuals to act in the interest of the company's reputation, even when personal incentives are weak. Without such preferences, inconsistent or self-serving behavior can erode collective credibility and damage a reputation in ways that are difficult to repair.

A negative illustration of this logic can be seen in the case of McKinsey & Company. Historically, McKinsey consultants benefited from the firm's prestigious reputation as a trusted advisor. However, the maintenance of this public good required voluntary adherence to fairness norms and social responsibility, even when short-term incentives suggested otherwise. In key cases, such as advising Purdue Pharma on strategies to boost opioid sales during the U.S. opioid crisis, and consulting for authoritarian regimes

despite ethical concerns, McKinsey's actions prioritized immediate client fees over societal well-being. As a result, individual- and team-level opportunism eroded the collective credibility of the firm.

The consequences were severe. McKinsey's reputation suffered significant damage across industries and governments, leading to legal settlements, public backlash, and long-term trust deficits. This erosion exemplifies how inconsistent behavior, driven by the absence of fairness-motivated restraint, can destroy a corporate public good—a reputation painstakingly built over decades. McKinsey's case highlights that without social preferences guiding individual and organizational behavior, corporate reputation becomes highly vulnerable to collective failure.

These examples illustrate how corporate public goods both rely on and reinforce social preferences. When firms actively support these goods, they generate positive externalities that extend beyond immediate stakeholders. Conversely, when social preferences are absent or ignored, individuals are more likely to withhold effort, act opportunistically, or focus narrowly on personal gain. In this sense, social preferences are not just desirable traits, they are essential foundations for sustaining the public goods on which modern organizations depend.

The Willingness-Awareness Matrix: A Practical Framework for Companies

One of the most persistent challenges in organizations is the under-provision of corporate public goods. Since contributing to these goods is often personally costly and benefits are shared collectively, firms frequently face free-rider problems on the employee level. Understanding and addressing these challenges requires a structured diagnostic approach.

The Willingness-Awareness Matrix (WAM) provides such a tool and is based on insights from the literature on social preferences and voluntary cooperation in public goods. I have described it the first time in Fehr (2018). The WAM, and its application to solving practical problems, is based on the following insights and assumptions:

Firstly, once individuals understand that a problem is generated by the free-rider incentives and externalities that are inherent in a public goods problem, there automatically emerge normative demands on individuals' behavior to avoid negative externalities

Table 1: Key Corporate Public Goods: Their Strategic Value, Dependence on Social Preferences, and Consequences of Their Presence or Absence

Public Good	Why Is It a Public Good?	Why Is It Valuable?	Why Social Preferences Matter	Why It's Hard to Achieve	Positive Consequences (if Present)	Negative Consequences (if Absent)	Example
Corporate Culture	Shared values and behaviors that shape workplace interactions; once established, they benefit all, regardless of who contributed to them.	Boosts productivity, loyalty, and collaboration.	Preferences for fairness and reciprocity motivate individuals to uphold shared values, even when cultural contributions go unrewarded.	Requires ongoing effort to maintain, with benefits spread across the entire organization.	High motivation, creativity, effective teamwork.	Toxic culture, low motivation, high turnover.	Migros' cooperative identity (see above)
Reputation	A company's collective image affects all stakeholders; everyone benefits from a good reputation, but no one alone can secure it.	Builds trust with customers, partners, and employees.	Fairness concerns drive reputational behavior despite weak personal incentives to protect collective image.	Individuals rarely benefit directly from improving the company's reputation.	Strong partnerships, customer loyalty, attractive brand.	Loss of clients, talent, and credibility.	LEGO's reputation for safety and quality
Infrastructure	Shared systems and resources that enable operations; use by one does not reduce availability for others.	Supports productivity and smooth operations.	Social preferences facilitate the cooperation that is needed to maintain resources and avoid overuse.	Maintaining resources requires effort, but the benefits are shared by everyone.	Improved productivity, efficient operations.	Lack of resources, reduced productivity.	Firms' usage of Microsoft's Office platform (e.g., SharePoint)
Data & Knowledge-Sharing	Information and know-how that, once shared, are accessible to all without being depleted by use.	Supports innovation, decision-making, and learning.	Preferences for reciprocity encourage sharing instead of hoarding valuable knowledge.	Individuals may hoard knowledge instead of sharing it for the collective good.	Better innovation, faster learning, smart decisions.	Knowledge silos, duplication of effort, reduced efficiency.	NASA's open technical knowledge archives
Rules & Policies	Shared guidelines that apply to all; once norms are in place, everyone benefits from predictable standards.	Provides structure and clarity for decisions and behavior.	Perceived fairness is essential for acceptance and compliance.	Following rules may feel restrictive, with unclear personal benefits.	Clear processes, reduced conflicts, efficient workflows.	Ambiguity, poor compliance, chaotic environment.	Toyota's standardized lean processes
Innovation Climate	A supportive environment for new ideas that, once in place, benefits everyone—even those who don't contribute.	Drives problem-solving, competitiveness, and growth.	A concern for group success encourages idea-sharing and risk-taking, even when own rewards are uncertain.	Risk-taking and creativity often require personal effort without immediate rewards.	Increased creativity, faster adaptation, problem-solving.	Stagnation, reduced competitiveness, missed opportunities.	Amazon's two-pizza teams

Public Good	Why Is It a Public Good?	Why Is It Valuable?	Why Social Preferences Matter	Why It's Hard to Achieve	Positive Consequences (if Present)	Negative Consequences (if Absent)	Example
Expertise & Best Practices	Once shared, effective methods improve work across teams, regardless of who developed them.	Improves efficiency, quality, and consistency.	Reciprocity preferences support sharing of know-how, although sharing is costly for the contributor.	Sharing expertise requires effort without direct personal gain.	Better productivity, reduced errors, higher quality.	Knowledge gaps, inefficiencies, wasted effort.	Airbus's standardized engineering protocols
	Norms of constructive feedback improve learning and coordination across the organization, even for those who don't actively provide input.	Enhances learning, improves performance, strengthens communication, and fosters adaptive behavior.	Fairness and reciprocity preferences motivate people to provide feedback, even when it involves personal cost.	Providing feedback is effortful and may not benefit the giver directly.	Higher engagement, stronger alignment, continuous improvement, and psychological safety.	Mistrust, stagnation, misalignment, low morale, and unresolved conflicts.	Firm-wide anonymous feedback app for employees
Leadership	Provides direction and coordination that benefits the entire group, but the effort and responsibility fall on the leaders only.	Aligns people, creates purpose, and supports coordinated action.	Preferences for fairness, responsibility, and reciprocity motivate individuals to lead despite high effort and shared rewards.	Leadership requires ongoing effort and responsibility without immediate rewards.	High trust, motivation, and shared vision.	Distrust, poor guidance, fragmented goals.	See case study section discussing leadership principles
Compliance Systems	Once accepted, shared rules promote predictable behavior for all—even if not everyone contributes to upholding or enforcing them.	Ensures fairness, reduces misconduct, and protects reputation.	A concern for fairness can sustain rule-following, even when enforcement is weak or costs are individual.	Following rules can be demanding and feel restrictive without direct rewards.	Consistent behavior, ethical culture, reduced risks.	Misconduct, confusion, reputational damage.	Volkswagen's compliance reforms after "Dieselgate"
Team Spirit & Trust	Mutual trust and team spirit strengthen cooperation for everyone but require ongoing contributions that many can opt out of.	Encourages collaboration, motivation, and loyalty.	Preferences for reciprocity motivate individuals to invest in team relationships, even when others may free-ride on their effort.	Building trust takes effort and time, but the benefits are collective.	Better cooperation, high morale, efficient teamwork.	Conflicts, mistrust, low productivity, high turnover.	Spotify's squad-based team model

and engage in behaviors that generate positive externalities. One of the most powerful examples of the normative forces that are unleashed by the explicit knowledge of negative externalities is provided by the emergence of the normative demands involved in the anti-smoking norm and the subsequent anti-smoking policies that emerged over the last few decades after it became clear that the passive inhalation of cigarette smoke is detrimental to the health of non-smoking individuals.

Secondly, however, in many instances, people do not automatically perceive a public goods problem as what it is, namely, a problem that is generated by free-rider incentives, and that everybody could, in principle, be made better-off if the problem were solved. Thus, often neither the free-rider incentives nor the collective benefits are salient enough. Nevertheless, when they can be made salient, and when the collective benefits of cooperation are fairly distributed, the existence of social preferences implies that there will automatically be many individuals who are willing to contribute voluntarily to the relevant public good. This is one of the key lessons that follows from the myriad laboratory public goods experiments that have been conducted over the previous three decades (Fehr & Schurtenberger 2018). In these experiments, the experimenter ensures that the subjects fully understand the public goods problem. As a consequence, we typically observe that many subjects are willing to contribute even in one-shot public goods experiments. They do so despite the fact that they would be individually better off if they engaged in free-riding, which is an indication of their social preferences.

It is simultaneously a strength and a weakness of the laboratory literature on public goods experiments that the experimenters carefully explain the economic incentives associated in public goods situations to the subjects. It has the advantage that bounded rationality and bounded awareness are unlikely to play a role. In reality, however, they do play a role and are intermingled with social preferences whose impact is diminished if individuals are not aware of the externalities involved in public goods. Therefore, in practical applications, such as in the case of corporate public goods, it is of paramount importance to create an awareness of the public goods nature of the relevant problems.

The two dimensions of the WAM are therefore “willingness” and “awareness”. In an organizational or corporate context,

- **willingness** refers to the extent to which employees are prepared to contribute voluntarily to corporate public goods despite the personal cost and the possibility of free-riding, while
- **awareness** describes whether individuals understand that their behavior creates externalities, meaning (positive or negative) effects for others and the organization as a whole, and that increasing the public good (by avoiding actions that create negative externalities or engaging in behaviors that create positive externalities) can make everybody better off.

By measuring employees’ willingness and awareness to contribute to important corporate public goods that are under-provided, a company acquires powerful knowledge about how it can solve the involved public goods problem. Specifically, depending on where a firm’s employees are located in the WAM, the solution to the public goods problem requires different intervention strategies.

High Willingness, High Awareness

In this case, the employees understand the negative consequences of noncompliance on the company, including its impact on other employees, and they are willing to contribute voluntarily to shared goals and norms. Still, cooperative behavior in this quadrant can be undermined by distractions, competing priorities, or lack of reinforcement. Organizations should reinforce attention and stabilize contributions through behavioral prompts and low-cost nudges—such as reminders, peer cues, or small public commitments—to sustain voluntary cooperation over time.

High Willingness, Low Awareness

Individuals are willing to contribute but don’t fully understand how their actions impact others or the organization. This often occurs in siloed teams and fast-changing environments, or with new employees. Once employees gain clarity about the broader effects of their behavior, cooperation can improve substantially. Organizations should increase awareness through targeted communication, transparent feedback, and tools that help make interdependencies and collective impact more visible.

Low Willingness, High Awareness

Employees understand that failing to follow cooperative norms has negative consequences for the company but choose not to contribute—often due to mistrust, poor incentives, or prior negative experiences. Here, the problem is motivational, not informational. Organizations should improve incentive structures, strengthen leadership credibility through consistent behavior, and implement systems to recognize and reward contributions to shared goals.

Low Willingness, Low Awareness

Employees neither realize that their actions have negative consequences for the company and their colleagues nor feel motivated to contribute to shared goals. This quadrant signals deep cooperation failure and typically correlates with dysfunctional norms or disengagement. Shifting behavior here requires foundational work along both dimensions. Organizations should begin by building basic awareness and establishing clear behavioral expectations associated with strong formal or informal incentives to cooperate, reinforced by leadership actions and, where necessary, external accountability or pressure.

By locating specific cooperation challenges within this matrix, firms can better understand the behavioral roots of their culture problems and design targeted interventions. For example, low willingness often calls for revised incentive systems and strong leadership by example, whilst low awareness may require storytelling, transparency, or nudges. In either case, the matrix reframes culture as a behavioral system—not a static value statement—and offers a map for shaping it.

Understanding where an organization stands in the WAM is not about labeling companies but about identifying where and why cooperation for public good fails—and what to do about it.

Case Study: Building a Shared Leadership Culture in a Financial Services Firm

A Swiss-based financial services firm launched an organization-wide initiative to embed seven shared leadership principles¹ across all levels. These principles were intended to strengthen leadership behavior, improve employee development and retention, and foster alignment between individual actions and the firm's long-term goals. Rather than treating leadership as a formal role, the initiative emphasized that any employee could act as a leader by living these principles in everyday decisions.

To inform the approach, the company conducted a behavioral survey involving over 470 employees from across its regional offices. The goal was not to introduce new values but to make the existing principles—already seen as relevant by 83% of employees—concretely actionable and sustainably practiced throughout the organization. This focus was essential, as 38% of employees reported that the principles were not helpful in their day-to-day work—highlighting the need to move beyond broad agreement and toward real behavioral application. The gap underscored a familiar cultural challenge: while all employees stood to benefit from a strong leadership culture, its maintenance depended on voluntary, ongoing contributions that often go unrewarded. This created a typical corporate public goods dilemma.

Prior to this initiative, the company had already introduced structural measures to support decentralized decision-making and empowerment—such as

1 *Leadership Principle 1 – Maximum Customer Orientation:* We consistently align our work to provide our customers with the best possible service. All other activities are subordinated to this mission.

Leadership Principle 2 – Continuous Improvement: We are convinced: Methodical innovation lies in continuous improvement through small steps. Progress starts with each of us. Leaders never stop learning.

Leadership Principle 3 – Long-Term Thinking: Standing up for one's convictions can require a lot of energy. We accept that we may be misunderstood in the short term when we pursue our long-term goals.

Leadership Principle 4 – Evidence-Based Decision Making: We do not rely solely on our gut feeling. When making important decisions, we get to the bottom of things and base our actions on numbers, data, and facts.

Leadership Principle 5 – Taking Ownership: Leaders never say, 'That's not possible'. They don't manage problems; they tackle them and find solutions. If something goes wrong, true leaders always look in the mirror first.

Leadership Principle 6 – Focus and Simplify: We do less, but better. Our customers prefer simple solutions over complex ones. Our entire organization is aligned toward the same goals.

Leadership Principle 7 – Creating Value: Everything we do must deliver the greatest possible value. Our work is successful when it leads to a measurable improvement for our customers. We are only satisfied with the best services, offerings, and processes.

flexible work models and delegated authority for team leads. However, these efforts did not translate into reliable behavioral change. Internal assessments revealed that attrition was linked not to dissatisfaction with the company but with individual supervisors. Furthermore, formal flexibility on working from home existed but was applied inconsistently, often out of fear of misuse, and although employees were formally empowered, decisions frequently reverted to senior management. Leadership roles were filled, but often without systematic development or coaching. In short, while the structural tools existed, the behavioral foundation was weak, and as a result, voluntary contributions to a shared leadership culture were neither expected nor reinforced.

To address this gap, the company applied the WAM as a diagnostic tool. Each leadership principle was assessed based on two dimensions: employees' willingness to act voluntarily in line with the principle, and their awareness of the consequences if the principle was not upheld. The survey revealed a consistent pattern: willingness was high—with on average 80–90% of employees expressing readiness to contribute to the principles—while awareness lagged behind, with only 50–70% recognizing the negative consequences of noncompliance. In particular, many employees underestimated how their own inaction—such as not speaking up or offering feedback to other leaders—contributed to cultural erosion.

In response, the company convened a leadership retreat with regional managers to co-create a portfolio of targeted interventions, in an attempt to promote voluntary cooperation by taking into account social preferences. The roadmap focused on five key areas based on behavioral levers that were identified in the survey. First, the company aimed to strengthen feedback and recognition by creating regular formats where positive contributions would be acknowledged and peer feedback normalized. Second, it addressed orientation and clarity by illustrating how each principle could be applied in specific work situations. Third, supportive structures were introduced by embedding the leadership principles into coaching practices, 1:1 conversations, and employee development formats. Fourth, to promote long-term engagement, the organization linked principle adherence explicitly to career development. Finally, practical initiatives

were launched to increase autonomy and trust—such as giving teams more decision-making authority on customer-facing issues like pricing. These measures were not framed as compliance tasks but as behavioral enablers for norm-driven cooperation.

As a part of the first set of interventions, the initiative focused on strengthening the feedback culture. The core idea was to ensure that giving and receiving feedback would not be perceived as an extra burden but as something embedded in the flow of work. To achieve this aim, feedback practices were deliberately attached to ongoing strategic projects, making them part of regular routines rather than isolated activities. Clear feedback rules were defined—not only regarding how to give feedback, but also when and in what situations it should occur, such as during project debriefs, team check-ins, or regular 1:1 conversations. The implementation process was tracked over time and regularly reflected upon. Two concrete measures made this shift visible. First, an upward feedback process was introduced, enabling employees to provide feedback to their leaders in a systematic and psychologically safe way. Second, a “coin system” was piloted to foster peer-to-peer recognition. Each employee could distribute symbolic coins (e.g., 100 per year) to colleagues as a gesture of appreciation, thereby anchoring feedback in positive reinforcement.

Together with all other interventions, these measures did more than normalize feedback—they established concrete routines and formats that made contributions to the leadership principles as a corporate public good. By reinforcing the perception that living these principles leads to recognition and appreciation, the interventions increased awareness of how one's actions (or inactions) affect the organization. At the same time, they activated social preferences such as the concern for fairness and reciprocity, consequently increasing employees' willingness to contribute voluntarily. In doing so, the company laid the behavioral foundation for a self-sustaining leadership culture built on cooperation rather than formal enforcement.

This case shows how building a leadership culture involves addressing a classic corporate public goods problem. Although a healthy leadership environment benefits all, individual contributions often require effort without immediate reward—and they can

be easily withheld. Defining values is not enough; organizations must identify and address the behavioral barriers that prevent everyday enactment of those values. The WAM helped the firm understand the specific dynamics at play, namely, high motivation but low awareness of their (in-)actions and consequences. This insight guided the design of interventions that targeted behavior, not just belief. The initiative's success depended on leveraging social preferences—particularly preferences for fairness, reciprocity, and social recognition. By reinforcing the belief that others were also contributing, and by making cooperative behavior visible and valued, the company created conditions under which voluntary cooperation became more likely. Ultimately, rather than enforcing culture top-down, it built a system that invited employees to cooperate voluntarily.

Conclusion: Embedding Social Preferences for Sustainable Success

Social preferences are not ancillary to business strategy—they are a key behavioral/motivational characteristic of many employees that must be taken into account. In addition, companies can take applicants' social preferences into account when hiring them, and in so doing they can shape the composition of social preferences among their employees. Social preferences are not simply moral ideals; they are productive forces that can help generate core corporate assets such as trust, culture, and reputation. These corporate public goods, in turn, improve collaboration, reduce friction, and build long-term resilience. Recognizing their strategic value requires moving beyond ad hoc initiatives and embedding these preferences into how companies design leadership, incentives, and decision-making processes.

The Willingness-Awareness Matrix provides a structured way for companies to assess how they can increase the provision of relevant corporate public goods. It offers diagnosis and guides action. Companies that take social preferences seriously as a source of competitive advantage have a higher chance of being better positioned to navigate uncertainty, attract talent, and sustain cooperation over time. The success of companies, in this view, also rests on how they leverage social preferences to generate behavioral foundations for the provision of corporate public goods.

THE AUTHOR

Ernst Fehr is Professor of Microeconomics and Experimental Economics at the University of Zurich and Director of the UBS Center for Economics in Society. His research examines the role of social preferences and social norms for incentives, market behavior, and organizational outcomes. He has also contributed to the fields of neuroeconomics and behavioral finance. Ernst has published in leading academic journals across economics and interdisciplinary science, and his work is widely cited. He is a member of the American Academy of Arts and Sciences, the John Kenneth Galbraith Fellow of the American Academy of Political and Social Sciences, and a Foreign Honorary Member of the American Economic Association. He also served as a president of the Economic Science Association and the European Economic Association. He co-founded FehrAdvice & Partners AG, a consultancy systematically applying behavioral economics to business and policy.

REFERENCES

- Andreoni, J., & Miller, J. (2002). Giving according to GARP: An experimental test of the consistency of preferences for altruism. *Econometrica*, 70(2), 737–753. <https://www.jstor.org/stable/2692289>
- Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3), 817–868. <https://doi.org/10.1162/003355399556151>
- Fehr, E., & Charness, G. [Forthcoming]. Social preferences: Fundamental characteristics and economic consequences. *Journal of Economic Literature*.
- Fehr, E. (2018). Behavioral foundations of corporate culture. *University of Zurich, UBS International Center of Economics in Society, Public Paper*, (7). <https://dx.doi.org/10.2139/ssrn.3283728>
- Fehr, E., & Schurtenberger, I. (2018). Normative foundations of human cooperation. *Nature Human Behaviour*, 2(7), 458–468. <https://doi.org/10.1038/s41562-018-0385-5>
- Rabin, M. (1993). Incorporating fairness into game theory and economics. *American Economic Review*, 83(5), 1281–1302. <https://www.jstor.org/stable/2117561>



EDITORIAL

Advancing Behavioral Economics Research for Health and Finance

SUZANNE B. SHU¹, GEOFFREY FISHER, DAVID R. JUST AND STEPHEN SHU

Cornell University

Behavioral economics research has thrived in recent years both through descriptive research, which explores how the ways in which people make real decisions can differ from normative theory, and through prescriptive research, which designs interventions (nudges) to improve decision-making. In this editorial, we look forward to new research opportunities in both of these broad areas. On the side of descriptive research, we consider how individuals' choices can reflect misalignment between stated and revealed preferences, and what this implies for theory. We also consider how newer methods in process-tracing, including eye tracking, can be used to unpack the non-choice elements of the decision process. On the side of prescriptive research, we discuss how work on choice architecture can continue to evolve to reflect more layers of the choice environment, including situational context and decision-maker heterogeneity. We also discuss opportunities and challenges for bringing behavioral interventions into the field including partnering with industry.

Introduction

The overall field of behavioral economics is broad in terms of the types of research undertaken by both academics and practitioners. On the theory side, we investigate how actual choices and behavior may differ from normative theories of utility. Research has documented the ways in which individuals regularly employ heuristics in their judgments and decisions that lead to predictable biases. Within the domain of methods, we employ a wide range of empirical approaches, including laboratory studies for hypothetical and/or incentive-compatible choices, and field experiments that document actual behavior in natural environments. And on the application side, we have both researchers who focus on designing interventions that are optimized for contexts and individuals, and practitioners who focus on bringing those interventions successfully into field environments.

In this chapter, we look at the current issues being explored by behavioral economists in the health and financial domains, with special attention given to four main topics in which this author team has particular expertise: misaligned goals as a driver

for consumption problems; process-tracing and other new testing methods; new opportunities for behavioral economics research; and challenges in implementing our learnings in applied contexts.

In our first section, we use the context of food choices to explore how consumption preferences can be misaligned with stated goals. We note that testing these behaviors can be challenging when done solely through choice data, and so we explore new directions in process-tracing methods in our second section. Our third section looks forward to new directions in behavioral economics research, which will in turn require new empirical methods that can account better for the influence of situational contexts and individual differences on behavior. Finally, we dive more deeply into applied problems, with a special focus on the context of financial choices, and an investigation of implementation challenges for bringing behavioral thinking to industry. We hope that this end-to-end perspective, from consumer preferences and goals through testing and application, can provide guidance and inspiration to new generations of behavioral economics researchers and practitioners.

¹ Corresponding author: sbs78@cornell.edu

Consumers Plagued by Misaligned Goals

We begin by looking carefully at the individual decision-maker and our ability to understand their preferences. Consumers increasingly seek goods and services with credence qualities that are not clearly observed at the time of purchase. These qualities include desired positive impacts on health and fitness, sustainable environmental impacts, and adherence to ethical production processes. However, in many cases, the goods people purchase to address specific credence qualities are demonstrably misaligned with the goals of the consumer. Thus, we may purchase fattening foods that have a reputation for being natural in the interest of health (Rahman et al., 2020), or wear t-shirts with positive social messages sourced from sweatshop labor (e.g., Ellery, 2014), or buy goods produced locally with a high impact on the environment in the interest of lowering our carbon footprint (Stein & Santini, 2022).

From a behavioral perspective, the rise of misalignment between stated and revealed preferences, due to a lack of direct observability of the desired attributes, is both important and understudied. To the extent that systematic biases create these misalignments, such are of interest to both behavioral decision researchers and policymakers.

In their annual survey of food consumers, the International Food Information Council reports that more than half of Americans follow a specific diet or eating pattern—a number that has been growing for several years (IFIC, 2024). The stated goals behind these efforts vary, but primarily they center around feeling better, losing weight, or meeting other health-related goals. Despite these stated goals, however, the types of diets engaged vary dramatically (both ketogenic and vegetarian diets feature prominently), with many of the cited eating patterns having little to no connection to the underlying goal (e.g., gluten-free diets are often associated with weight loss goals or general nutrition, rather than gluten sensitivity) (Prada et al., 2019; Priven et al., 2015).

Products are often marketed with this behavior in mind. Consumer fads enable marketers to use vague terms or specific signals to imply virtuous properties without actually possessing those properties. Thus, these marketers can fuel prominent trends in consumption around concepts like “clean label” or “natural” that are disconnected enough from specific

health claims to avoid more stringent regulation. Even government-imposed nutrition labels have been shown to mislead consumers (Barahona et al., 2023).

Prior to the popularity of GLP-1 Agonists, commercial weight loss programs were perhaps the most blatant example of such misalignment and the strange relationship that arises between producer promises, consumer perceptions, and revealed preference. The weight loss industry takes in a combined \$70 billion every year in the US (Harvard, 2022). A variety of commercial weight loss programs are marketed, including many with significant research supporting their program.

For example, WeightWatchers and Nutrisystem are both based upon research literature and have been the subject of effectiveness studies (Ahern et al., 2011; Habowski et al., 2017). These programs involve adherence to substantive guidelines or rules of thumb governing what and how much one eats. Not surprisingly, they are typically marketed to emphasize the amount of weight some have lost on the program and the ease of adherence. Reading the ads for programs with less grounding in evidence becomes a comical exercise, with each touting more and more unbelievable claims of weight loss results accompanied by claims of eating what you want and not requiring any exercise (see Lim et al., 2020; Sopher, 2004). While such claims are at face value unbelievable, the firms making them are commercially successful. Consumers are either not capable of seeing through the claims at the point of decision or they are willing to suspend disbelief (perhaps out of desperation).

Several behavioral theories present alternative explanations for this conduct. It has been common to think of indulgence and procrastination in dieting and exercise being the result of present biased preferences (Cheung et al., 2022; Hunter et al., 2018). One may simply act on this information, recognizing their inability to adhere to a more effective program. Furthermore, acting on unbelievable claims may also present a low effort resolution to the natural cognitive dissonance that arises around diet—a desire to be more fit, but resistance to change in diet or exercise (Ong et al., 2017).

Understanding the origins and results of magical thinking and misaligned preferences and goals in consumer transactions is key to finding effective policy to manage misinformation and predatory marketing behaviors within a variety of markets.

Behavioral work in psychology has produced a number of promising directions that help elucidate such apparent divergence in goals and behavior. For example, Barahona et al. (2023) find evidence that nutrition warning labels cause consumers to make incorrect inferences about products that escape labeling requirements.

One promising explanation applies fuzzy trace theory (Blalock & Reyna, 2016) to consumer behavior. Fuzzy trace theory is a dual process model, supposing that consumers economize on cognitive effort by summarizing and simplifying information into the gist. Thus, instead of making sense of the entire list of nutritional facts panel on a food product, an individual may make a general judgment of what that panel and any other labeling is trying to communicate. Similarly, an individual may consider an attribute like “organic,” which has complicated and nuanced requirements, and take it as a simple indicator of nutritional content as a means of a decision shortcut. Fuzzy trace theory has been useful in describing risky decisions in a variety of contexts such as drug use and unprotected sexual activity. Other work has also helped to highlight how perceptions can become so misaligned, for example through the misperception of causation (Spiegler, 2020). To the extent that such models accurately represent decision processes, they indicate an important wedge between revealed preference and the underlying goals of the consumer.

In the context of behavioral economics, such theories could be extremely useful in considering social welfare implications as well as potential policy responses. The wedge driven between revealed preference and underlying goals creates perverse incentives for those marketing goods to play on misperceptions of the performance and attributes of their products. Moreover, given these misperceptions of the usefulness and meaning of attributes, marketers are able to do this without running afoul of laws regarding truth in advertising and necessary documentation of explicit health claims (Hastak & Mazis, 2011; Vladeck, 2000). To understand the full implications requires rigorous and credible modeling of both the consumer behavior and profit incentives created for marketers.

In differentiated—yet somewhat competitive—markets, such as those for food or weight loss programs, these incentives could potentially cause substantial harm to unwitting consumers. Designing

effective policy to combat such predatory marketing behavior must engage the consumer decision-making process and specifically the way in which the consumer makes sense of available information. We would need substantial understanding of how consumers read (or ignore) labels, claims, and third party certification badges.

To date, most studies in this area have focused very specifically on which types of labels produce specific behaviors. For example, Ares et al. (2023), in their review of the literature, find that consumers tend to choose more nutritious foods in response to information presented as warning labels. Similarly, consumers respond appropriately to simple traffic lights (Trudel et al., 2015). However, they tend not to respond in such straightforward ways to more complicated or nuanced information (Balasubramanian & Cole, 2002; Grunert & Wills, 2007). These results are helpful in designing government-required labels but do little to shape policy on the information landscape as a whole. In order to address the consumer response to the broad landscape of information in a free society, we must have a greater understanding of the information acquisition and decision process.

Newer Methods for Tracing Decisions

The above examples of how greater health goals can be misaligned with actual product qualities demonstrate many of the pervasive problems that trouble consumers. Although economists are traditionally concerned with exclusively analyzing choices, recent work has made use of non-choice data and allows for the potential to add value to standard economic analysis by providing new perspectives about the decision process. This broader class of data includes physiological measures (e.g., heartrate or galvanic skin response), neural measures (e.g., fMRI or EEG), and digital metadata (e.g., clickstream or browsing behavior), each offering insights into internal states, attention, and preferences. Among these tools, response times and eye-tracking have emerged as particularly influential in economics, as they center around capturing attentional allocation and processing effort to enrich our understanding better in terms of how consumers arrive at their choices.

Of these two process-tracing tools, response time has historically received more attention in the

economics literature. Response time, which captures the total duration it takes for a decision-maker to make a choice, has a rich history of use across various research domains beyond economics, notably in psychology. It plays a central role in dual-process theories, which distinguish between automatic, intuitive responses and deliberate reasoning, and it is often employed to uncover implicit attitudes (Kahneman & Frederick, 2002; Shiffrin & Schneider, 1977). The key finding underlying much of this work is that decisions that are more difficult yield longer response times.

Measuring response time in economic contexts is often straightforward, as it is endogenously generated in the decision process, and most computerized platforms used in experimental data collection do not require specialized software to capture it. Hence, most prior work that has utilized response time has occurred in laboratory settings where such measurements are simple, and this work has encompassed a wide range of economic scenarios, including risky decision-making (Alós-Ferrer & Garagnani, 2024), information cascades (Alós-Ferrer et al., 2021; Frydman & Krajbich, 2022), consumer product decisions (Clithero, 2018), and coordination games (Schotter & Trevino, 2021). These findings largely suggest that making use of response times improves choice prediction, and that participants intuitively recognize the link between processing duration and difficulty.

One potential concern about utilizing response times is that they may be of use only in laboratory studies where the decision environment is carefully controlled and decisions are relatively fast (i.e., in the order of just a few seconds). However, several recent studies suggest that the analysis of response times still provides useful data in field settings. For example, Cotet et al. (2025) find that response time predicts eBay buyer and seller behavior so that it takes individuals more time to reject good offers and more time to accept bad offers. In this case, response time is often in the order of hours or days. Chiong et al. (2024) use user-generated response times from mobile advertisements to estimate a decision-making model, establishing that utility estimates from their model are correlated with external measures of ad engagement. Finally, Card et al. (2024) investigate response time at top economics journals and find that the probability that authors receive a positive decision

increases in line with decision time. Together, this work suggests that the analysis of response time is generalizable to certain real-world settings.

A second process-tracing tool that has received attention in the literature is eye-tracking. Traditionally, it has used specialized high-frequency cameras that sit below a computer monitor and record where on a computer screen the participant is fixating (i.e., gazing) at each time point. High-quality versions typically allow for measurement of pupil size, in addition to gaze patterns, which can be further analyzed. The analysis of eye-tracking data can be relatively straightforward using specialized software that preprocesses the raw data and can generate simplistic visualizations. Although academic research typically employs eye-tracking in highly controlled laboratory settings to ensure accuracy, recent technological advances have expanded its applicability, with computer and mobile webcams being increasingly used despite potentially lower accuracy compared to traditional setups (Sammelmann & Weigelt, 2018; Yang & Krajbich, 2021). In commercial marketing contexts, eye-tracking is often used to assess engagement with advertisements, product designs, and online user interfaces. Both open source software and commercial services (e.g., Lumen Research) allow for a broader adoption of eye-tracking across academic and industry settings.

Within economics, a number of papers have highlighted the utility that eye-tracking data provides by testing whether fixation patterns are consistent with economic models in contexts such as sender-receiver games (Wang et al., 2010), consumption decisions (Reutskaja et al., 2011), and risky choice (Arieli et al., 2011). Yet, other work outside of economics is also relevant for considering how eye-tracking may be used in the future. One stream of literature from psychology and neuroscience has sought to add fixation data explicitly into models. Here, these models not only make predictions about what individuals choose, but also determine how those choices are correlated with response time and fixation data (e.g., Krajbich et al., 2010). A second stream has addressed whether certain fixation patterns are associated with decisions. For example, Reeck et al. (2017) note that when participants compare intertemporal options along their attribute values, rather than integrating information about each option, they are more likely to choose a larger-later

option. Better understanding of the relationship between time spent fixating to information, and how the information is processed, is an interesting open question to consider for future work.

Response time and eye-tracking are some of the most common process measures used in behavioral economics, but it is worth asking what additional measures can be integrated with choice data. Psychology and neuroscience have utilized a wide range of tools, including functional magnetic resonance imaging (fMRI) and electroencephalography (EEG), but some of these are challenging in relation to collecting and analyzing the data they produce, without specialized training or collaborating with researchers with specific expertise. Therefore, these techniques may have received less attention in the literature compared to the measures discussed above.

Looking towards the future, one potentially underutilized tool measures the path that individuals take to physically select a choice option, often through finger trajectories or mouse cursor trajectories (Dotan et al., 2019). Although this data is not typically saved as a default in most experimental platforms, it can be captured with custom JavaScript code or by using many of the freely available programs that do so (e.g., MouseTracker).

Trajectory tracking has a rich history in psychology and has been utilized for several different types of analysis. In the first type, reaching trajectories have been shown to provide additional information about decision conflict, which may improve predictions of choices above and beyond utilizing solely response time (Stillman et al., 2020). In the second type, prior work has used reaching trajectories to estimate the time at which an attribute is first utilized in the decision-making process, revealing that attributes are not always assessed at the same time point. Moreover, differences in attribute timing are associated with differences in choices (Sullivan et al., 2015). Overall, process-tracing offers a number of techniques to unpack the decision-making processes that have yet to be explored fully in economic contexts.

Expanding Frontiers in Behavioral Economics Research

The sections above focus on two important areas of current behavioral economics research that are

receiving well-deserved recent attention. New developments in these two research areas are crucial to developing a better understanding of both how consumers' choices may or may not reflect actual preferences and how we may better track the thought processes behind those choices. Nonetheless, there are topics in behavioral economics in even earlier stages of research development that are also deserving of greater effort if we are to move the field forward.

One of the great success stories for bringing theory into practice in the behavioral economics landscape is the growth of choice architecture: the theory-driven, systematic development of behavioral interventions (nudges) that can influence decisions in positive ways while preserving individuals' freedom of choice (Thaler & Sunstein, 2009). Traditional interventions in this literature stream have included defaults (Madrian & Shea, 2001; Johnson et al., 1993), option partitioning (Langer & Fox, 2005; Bardolet et al., 2009), labeling (Peters et al., 2009; Riis & Ratner, 2014), limited time windows (Shu & Gneezy, 2010; O'Donoghue & Rabin, 1999), and cues and feedback (Milkman et al., 2014; Giuntella et al., 2024).

While these interventions are often highly cost-effective at changing behavior (Benartzi et al., 2017), they are typically tightly designed to address a decision bias in a single type of choice context with a single, clear outcome goal. Consider, for example, interventions designed around the concept of psychological ownership in the context of public spaces, as studied by Peck et al. (2020). In that context, the outcome goal is to overcome the free-rider thinking associated with the tragedy of the commons, and the intervention's success is measured through stewardship behaviors toward public spaces. Psychological ownership interventions have also been tested in contexts such as government benefit acceptance (De La Rosa et al., 2022; Shu, 2018), vaccination uptake (Milkman et al., 2022), and cancer screening (Bakr et al., 2020). In each case, the context and desired decision outcome are clearly defined, and so a one-size-fits-all intervention to achieve the goal within that specific context can be developed.

However, humans rarely make decisions in such clearly defined situations, where a single type of intervention can work for all individuals. First, the situations themselves can differ in significant ways

that make a particular intervention designed for similar choices inappropriate for the new environment.

In the health domain, designing interventions for chronic disease patients first requires a recognition that not all chronic diseases operate the same way, in which case they can require different approaches (Mogler et al., 2013). For example, hypertension is a disease that is mostly asymptomatic for long periods, and patients do not have natural tangible feedback that may guide behavior change. In contrast, a disease such as diabetes has immediate feedback when uncontrolled, and thus patients may be more motivated to be consistent with their treatment. In behavioral intervention terms, these context differences may imply that an intervention designed to overcome present bias, by highlighting long-term outcomes for hypertension (Shapiro et al., 2020), is less impactful for a disease like diabetes, where reminder-based interventions could be most effective (Quinn et al., 2011).

In the financial decision-making domain, interventions that have been found to work in one context can even have opposite effects in other contexts, such as the impact of an informational display that can increase future-biased decisions in an annuity choice context (Shu et al., 2016) but decreases them in a social security-claiming choice context (Shu & Payne, 2025). In the health and environmental domains, a social comparison intervention that successfully changes antibiotic prescribing behavior among doctors (Meeker et al., 2016) has also been documented to have a boomerang effect on already well-performing homeowners in an energy consumption intervention (Schulz et al., 2007). More broadly, choice architects who hope to depend on behavioral economics findings need more research focusing on defining the characteristics of contexts that determine which interventions will work for changing behavior in different types of situations. For practitioners, this means that it is essential to test previously effective interventions for any new context to which it is being applied.

In addition to the characteristics of the situational contexts, future successful behavioral economics intervention research will need to be pay more attention to the heterogeneity of individual decision-makers. Just as personalized medicine recognizes that genetic differences between patients may require different

drugs per individual to treat a condition, behavioral economists will need to consider how to design nudges that are specific to the personal traits of the decision-maker (Sunstein, 2013; Goldstein et al., 2008; Banerjee & Galizzi, 2024; Munshi & Ramani, 2023; Delaney & Bucher, 2022; Heal et al., 2022; Mills, 2022).

Returning to the health domain, designing interventions for chronic disease patients also requires understanding each individual's knowledge and beliefs, costs and barriers, attention levels, and self-efficacy (Mogler et al., 2013). One patient with an uncontrolled chronic disease, for instance, might hold biased beliefs that could be influenced through a social norms intervention, while another patient with that same disease may hold accurate beliefs but be forgetful about treatment schedules and thus responsive to appointment defaults.

In the financial decision-making domain, efforts have been made to test the interaction of various interventions with individual decision-maker differences in the context of the social security-claiming decision (Greenberg et al., 2023). Customized interventions, such as smart defaults, have also been explored in the areas of passwords (Peer et al., 2019), tax payment reminders (Halprin, 2016), and electricity use feedback (Costa & Kahn, 2013).

However, the development of personalized interventions faces several methodological challenges for an empirical researcher. The greatest of these is the very large sample sizes required; not only is the researcher attempting to run a horse race among multiple interventions, but they are also measuring multiple individual psychological measures and looking for the significant interactions between these factors. To avoid the cherry-picking of data and/or reporting of spurious effects, good data science practices such as preregistration are critical.

Progress has been made with the increase of so-called “megastudies,” which simultaneously test multiple interventions in the field on large numbers of participants (Milkman et al., 2022), but unfortunately, it is rare for these field studies to collect individual measures, and so the connection between the successful interventions and the population segments they work for is not made. New techniques in working with large datasets, such as recursive partitioning, provide opportunities for analyzing them in this way, such as the work done by Shah et

al. (2024) to evaluate the effects of savings nudges on population segments in Mexico. Beyond simply customizing nudges for different subpopulations, new approaches like dynamic nudging (Dalecke & Karlsen, 2020; Nixon & Gilbert, 2024) could allow interventions to evolve with additional user input. New empirical techniques that leverage machine learning, AI, or large language models may also provide future research breakthroughs on this challenging topic.

Bringing Behavioral Economics Findings to Industry

Having now considered concerns about misalignments in consumer preferences, new methods for tracing decision processes, and the importance of considering factors of heterogeneity and situational influences, we turn to the question of how external audiences can best apply the findings of our field. Consistent with the research focus on heterogeneity, one next frontier for behavioral economics applications within industry lies at its intersection with hyper-personalization. This ultimately means addressing heterogeneity, such as how people differ in terms of their demographics, worldviews, individual behavioral differences, situations, and decision-making processes. However, there are several areas that need to be addressed to pursue this frontier: 1) improving research on heterogeneity while also paving the way for high-potential use cases in industry; 2) changing the organizational DNA within industry to implement behavioral insights better; and 3) facilitating the development of foundational behavioral platforms.

Historically, the culture of scientific publication has placed immense value on mass behavioral interventions and average treatment effects. While good for the publication industry, this has potentially had the impact of diminishing the importance of heterogeneity findings and the opportunity to do good in the applied world. For example, anecdotally, when relaying studies that focus on the benefits of reframing emergency savings as \$5 per day versus \$150 per month (Hershfield et al., 2020), or retirement savings as X% of my salary versus X pennies for every dollar I earn (Shu et al., 2022), it has not been unusual for industry professionals to recall average treatment effects yet overlook heterogeneity details,

such as whether those with the lower incomes were helped more relative to others.

The lack of saliency in the mind of professionals potentially results in missed opportunities in the transition from scientific findings to industry application, where people might be helped through heterogeneous treatments and personalized interventions. The opportunity to improve research on heterogeneity has been noted by prior researchers (Bryan et al., 2021). Key notes include proactively addressing heterogeneity upfront as opposed to making it a nuisance to explain after the fact, measuring more holistically-related characteristics of study participants and contexts, using newer statistical techniques for exploratory analyses, and potentially developing shared infrastructure to make data collection and related costs more accessible to researchers. However, as a behavioral community, we could also go further in encouraging the building of bridges from science to the applied world and industry, where heterogeneity is vast.

Moreover, we need more focus on and investment in bridges to stronger generalization, experimentation, and application for high-potential use cases. Perhaps this can be achieved through more academic-industry collaborations and funding to address areas ripe with heterogeneity, like personal finance, debt management, insurance, savings, and retirement income.

A second area that should be addressed involves changing the organizational DNA within industry to implement behavioral insights better. One starting point is that organizations should widely incorporate behavioral audits as part of their business processes to identify behavioral obstacles hampering company, end-user, and constituent goals, including segmenting groups based on either hypothesized or known individual behavioral differences.

The supply of people who know how to apply behavioral economics insights, while still limited, has expanded substantially over the past 15 years since the book *Nudge* first came out (Thaler & Sunstein, 2009). As such, conducting an assessment of a portion of the business, using the lenses of behavioral economics, should be readily achievable by a large fraction of companies (and relatively inexpensive in the grand scheme of things). After identifying behavioral obstacles through behavioral audits, organizations then need to work on their management

ability to ideate, design, develop, validate, measure (e.g., A/B test), and deploy behavioral solutions. Part of the process should include adapting traditional management principles from product development but applying them to the management of behavioral solutions. For example, behavioral solution glidepaths could involve “platform thinking,” which covers the process of moving behavioral interventions and tests from lab to field, marketing to customer experience, and single product to complete product line. The management discipline of behavioral economics should also involve staying fresh relative to new technologies, such as those mentioned earlier for eye-tracking and emerging areas such as LLM-based synthetic sampling (Horton, 2023) and agentic AI simulations for business and marketing, which can help reduce cycle times and go-to-market costs.

Finally, from a business policy perspective, companies should also encourage professionals to weigh ethical considerations systematically (such as using principles from moral psychology, recognizing dark patterns, or, at minimum, getting them to reflect on the degree of goal alignment between parties for behavioral interventions). When it comes to the implementation of behavioral economics, ethical lenses need to be deliberate and not accidental.

Thirdly, organizations will need to create behavioral architecture foundations to support a hyper-personalized world. The potential “behavioral models” to either motivate or underpin a foundation are plentiful, and they can include behavior-centric models, like COM-B (Michie et al., 2011), design and choice architecture approaches (Thaler & Sunstein, 2009), psychology-centric models (e.g., involving perception, memory, social versus self, motivation theory), and others.

However, they will also have to integrate behavioral model data with traditional, contextual business data and frameworks. The data collection mindset of companies for data collection might also need to shift from passively collecting data to actively getting end-users to want to express their preferences and goals, such as if they had their own butler, confidant, or advisor to work with. Note that the unification of data will need to strike the right balance between thinking about value to the end-user, acknowledging the 24x7 real-time environment in which they live, modernizing the company’s infrastructure in manageable steps, and transforming the business over

time. Once data has been unified (say, by a high-value use case in the customer journey), companies then need to think about a robust delivery system for interventions, whether that be through using nudges, boosts, gamification, personalized CX, AI agents, or trained professionals.

As a hypothetical example in the financial area, if a provider has behavioral data and an understanding of an individual’s priorities and motivations for saving toward retirement, they may be able to merge such information with market info (e.g., drops in interest and mortgage rates) and traditional account data to highlight refinancing opportunities, which could result in recurring financial savings that could be redirected toward retirement. On the other hand, for a user that has not done much to reflect on retirement goals (e.g., either through their digital behavior or lack thereof), there could be different personalized nudges to help them take different actions (e.g., increasing their savings by 2 pennies for every dollar they earn, starting next year).

Conclusion

In this editorial, we have attempted to highlight some of the increasingly important new areas of behavioral economics research, as well as suggest additional directions for the field to expand. First, we have argued that misalignment between stated and revealed preferences, a foundational issue for studying behavior, is both important and understudied. There is misalignment throughout the chain, such as consumers choosing solutions that are misaligned with goals (e.g., gluten-free diets being used for weight loss goals instead of gluten sensitivity), marketers using vague language for virtuous signaling (e.g., labeling products as “natural”), and policymakers being able to understand how customers make sense of available information and make decisions (e.g., reading food labels in detail versus fuzzy tracing and getting the correct gist through lower-effort, cognitive processing). Policymakers who acknowledge these misalignments may then be able to build new approaches (labeling, choice architecture) that help guide consumers’ choices toward outcomes that better match their true preferences (Madrian, 2014).

Second, we have argued that, since the analysis of choice itself is limiting, advancing the field requires the integration of non-choice data to understand decision processes better. For example, not only have

longer response times in laboratory environments been indicative of cognitive load and problem difficulty, but response times in field environments also have been indicative of market behavior (e.g., good offers on eBay from sellers take longer for buyers to reject). Newer eye-tracking technologies have become more scalable than in the past, thus more readily allowing for the use of non-choice data to predict choice (e.g., larger-later options being selected after eye fixation patterns indicating greater user comparisons between intertemporal options versus integrating information within an option).

Third, to advance the field, we have argued that future successful behavioral economics intervention research will need to be more attentive to context and the heterogeneity of individual decision-makers. Whereas many early behavioral studies have addressed single biases in specific choice contexts that then lead to one-size-fits-all interventions as solutions (e.g., defaults for increasing savings), there are numerous situations that are less clearly defined and require more tailored interventions (e.g., financial decisions in retirement that need nudges specific to both the decision and the person making it). Studying the interactions between possible interventions and each segment's psychological characteristics remains methodologically challenging.

Finally, we have discussed three main points about translating behavioral economics research to industry: 1) academic-industry collaborations should explore heterogeneity for high-potential use cases; 2) the organizational DNA within industry should be adapted to implement behavioral insights better (such as widely adopting the notion of behavioral audits as a starting point); and 3) organizations need to create behavioral architecture foundations for a hyper-personalized world (such as systematically organizing non-choice and choice data, as described earlier, and incorporating behavioral models). While we have focused on industry adoption, it is worth noting that policymakers also play a role in adopting behavioral insights and would be subject to these same issues of heterogeneity and personalization when building behaviorally-informed regulation (Johnson & Leary, 2017; Chin et al., 2022).

In the remainder of this annual guide to behavioral economics research, we hope you will keep these four themes in mind as you read about the exciting

new projects being explored by modern researchers. Whether the research is looking at how preferences are expressed, exploring new methodologies for tracking choices, expanding the boundaries of how context and heterogeneity are accounted for, or bringing insights to a broader audience, the projects described in these pages will continue to lead the way for researchers yet to come.

THE AUTHORS

Suzanne B. Shu is the Dean of Faculty and Research for the Cornell University SC Johnson College of Business and the John S. Dyson Professor in Marketing at the Dyson School of Applied Economics and Management. The types of decisions analyzed in her research include consumer self-control problems and consumption timing issues, with important implications for both negative behaviors (such as procrastination) and positive behaviors (such as saving). Her work on financial decisions has focused specifically on decumulation during retirement (annuities, Social Security claiming) as well as on perceived fairness for financial products. She is also an NBER Research Associate and has served on the Academic Research Council for the Consumer Financial Protection Bureau and the Netspar Scientific Council.

Geoffrey Fisher is an Associate Professor of Marketing at Cornell University's SC Johnson College of Business. His research bridges marketing, behavioral economics, and neuroeconomics, with a focus on how people make multi-attribute decisions. His current work integrates process-tracing methods into models of decision-making, revealing how cognitive and attentional dynamics shape consumer behavior.

David Just is the Susan Eckert Lynch Professor in Science and Business in the Charles H. Dyson School of Applied Economics and Management at Cornell University and a fellow of the Agricultural and Applied Economics Association. His research applies behavioral economics to policy issues in food and agriculture. Much of his recent research examines issues surrounding food insecurity and consumer behavior related to private and public food assistance programs.

Stephen Shu is a Professor of Practice of Behavioral Economics at Cornell University, Dyson School of

Applied Economics and Management. He also serves as the Academic Director for the Master of Professional Studies (MPS) and CEMS Programs at Cornell. He has more than three decades of industry experience and is perhaps most recognized for helping to incubate some of the very first behavioral economics initiatives around the world in the commercial sector. His expertise is in the area of behavioral finance, and he has worked with numerous companies in the retirement, wealth, investment, banking, advisory, and FinTech spaces.

REFERENCES

- Ahern, A. L., Olson, A. D., Aston, L. M., & Jebb, S. A. (2011). Weight Watchers on prescription: An observational study of weight change among adults referred to Weight Watchers by the NHS. *BMC Public Health*, 11, 1–5. <https://doi.org/10.1186/1471-2458-11-434>
- Alós-Ferrer, C., Fehr, E., & Netzer, N. (2021). Time will tell: Recovering preferences when choices are noisy. *Journal of Political Economy*, 129(6), 1828–1877. <https://doi.org/10.1086/713732>
- Alós-Ferrer, C., & Garagnani, M. (2024). Improving risky-choice predictions Using response times. *Journal of Political Economy Microeconomics*, 2(2), 335–354. <https://doi.org/10.1086/728666>
- Arieli, A., Ben-Ami, Y., & Rubinstein, A. (2011). Tracking decision-makers under uncertainty. *American Economic Journal: Microeconomics*, 3(4), 68–76. <https://doi.org/10.1257/mic.3.4.68>
- Bakr, O., Afsar-Manesh, N., Raja, N., Dermenchyan, A., Goldstein, N. J., Shu, S. B., & May, F. P. (2020). Application of behavioral economics principles improves participation in mailed outreach for colorectal cancer screening. *Clinical and Translational Gastroenterology*, 11(1), e00115. <https://doi.org/10.14309/ctg.0000000000000115>
- Balasubramanian, S. K., & Cole, C., (2002). Consumers' search and use of nutrition information: The challenge and promise of the nutrition labeling and education act. *Journal of Marketing*, 66(3), 112–127.
- Banerjee, S., & Galizzi, M. M. (2024). People are different! And so should be behavioural interventions. In A. Samson (Ed.), *The Behavioral Economics Guide 2024* (pp. 109–118). <https://www.behavioraleconomics.com/be-guide/>
- Barahona, N., Otero, C., & Otero, S. (2023). Equilibrium effects of food labeling policies. *Econometrica*, 91(3), 839–868. <https://doi.org/10.3982/ECTA19603>
- Benartzi, S., Beshears, J., Milkman, K. L., Sunstein, C. R., Thaler, R. H., Shankar, M., ... & Galing, S. (2017). Should governments invest more in nudging?. *Psychological Science*, 28(8), 1041–1055. <https://doi.org/10.1177/0956797617702501>
- Blalock, S. J., & Reyna, V. F., (2016). Using fuzzy-trace theory to understand and improve health judgments, decisions, and behaviors: A literature review. *Health Psychology*, 35(8), 781–792. <https://doi.org/10.1037/hea0000384>
- Bryan, C. J., Tipton, E., & Yeager, D. S. (2021). Behavioural science is unlikely to change the world without a heterogeneity revolution. *Nature Human Behaviour*, 5(8), 980–989. <https://doi.org/10.1038/s41562-021-01143-3>
- Card, D., DellaVigna, S., Jiang, C., & Taubinsky, D. (2024). Understanding expert choices using decision time (SSRN Scholarly Paper 4851789). *Social Science Research Network*. <https://papers.ssrn.com/abstract=4851789>
- Cheung, S. L., Tymula, A., & Wang, X. (2022). Present bias for monetary and dietary rewards. *Experimental Economics*, 25(4), 1202–1233. <https://doi.org/10.1007/s10683-022-09749-8>
- Chin, A., Zimmerman, D., Johnson, H., & Shu, S. B. (2022). Decisions about overdraft coverage: Disclosure design and personal finances. *Journal of Experimental Psychology: Applied*, 28(4), 746–774. <https://psycnet.apa.org/doi/10.1037/xap0000460>
- Chiong, K., Shum, M., Webb, R., & Chen, R. (2024). Combining choice and response time data: A drift-diffusion model of mobile advertisements. *Management Science*, 70(2), 1238–1257. <https://doi.org/10.1287/mnsc.2023.4738>
- Clithero, J. A. (2018). Response times in economics: Looking through the lens of sequential sampling models. *Journal of Economic Psychology*, 69, 61–86. <https://doi.org/10.1016/j.joep.2018.09.008>
- Cotet, M., Zhao, W. J., & Krajbich, I. (2025). Deliberation during online bargaining reveals strategic information. *Proceedings of the National Academy of Sciences*, 122(7), e2410956122. <https://doi.org/10.1073/pnas.2410956122>

- Dalecke, S., & Karlsen, R. (2020). Designing dynamic and personalized nudges. In *Proceedings of the 10th international conference on web intelligence, mining and semantics*, June, 139–148.
- Delaney, S., & Bucher, A. (2022). If we build it right, they will come: Driving health outcomes with “precision nudging”. In A. Samson (Ed.) *The Behavioral Economics Guide 2022* (pp. 17–26). <https://www.behavioraleconomics.com/be-guide/>
- De La Rosa, W., Sharma, E., Tully, S. M., Giannella, E., & Rino, G. (2021). Psychological ownership interventions increase interest in claiming government benefits. *Proceedings of the National Academy of Sciences*, 118(35), e2106357118. <https://doi.org/10.1073/pnas.2106357118>
- Dotan, D., Pinheiro-Chagas, P., Al Roumi, F., & Dehaene, S. (2019). Track it to crack it: Dissecting processing stages with finger tracking. *Trends in Cognitive Sciences*, 23(12), 1058–1070. <https://doi.org/10.1016/j.tics.2019.10.002>
- Ellery, B. (2014, November 1). 62p AN HOUR: What women sleeping 16 to a room get paid to make Ed and Harriet’s £45 “This Is What A Feminist Looks Like” T-shirts. *Daily Mail*. <https://www.dailymail.co.uk/news/article-2817191/62p-HOUR-s-women-sleeping-16-room-paid-make-Ed-Harriet-s-45-Feminist-Looks-Like-T-shirts.html>
- Frydman, C., & Krajbich, I. (2022). Using response times to infer others’ private information: An application to information cascades. *Management Science*, 68(4), 2970–2986. <https://doi.org/10.1287/mnsc.2021.3994>
- Greenberg, A. E., Hershfield, H. E., Shu, S. B., & Spiller, S. A. (2023). What motivates social security claiming age intentions? Testing behaviorally informed interventions alongside individual differences. *Journal of Marketing Research*, 60(6), 1052–1070. <https://doi.org/10.1177/00222437221147221>
- Goldstein, D. G., Johnson, E. J., Herrmann, A., & Heitmann, M. (2008). Nudge your customers toward better choices. *Harvard Business Review*, 86(12), 99–105. <https://hbr.org/2008/12/nudge-your-customers-toward-better-choices>
- Grunert, K. G., & Wills, J. M. (2007). A review of European research on consumer response to nutrition information on food labels. *Journal of Public health*, 15, 385–399.
- Habowski, S., Ziegenfuss, T., Sandrock, J., Raub, B., Kedia, W., & Lopez, H. (2017). A prospective evaluation of a commercial weight loss program on body weight and body circumferences in overweight/obese men and women. *The FASEB Journal*, 31, lb291–lb291.
- Hastak, M., & Mazis, M. B. (2011). Deception by implication: A typology of truthful but misleading advertising and labeling claims. *Journal of Public Policy & Marketing*, 30(2), 157–167. <https://doi.org/10.1509/jppm.30.2.157>
- Harvard Health Publishing (2022). *Diet & weight loss*. <https://www.health.harvard.edu/topics/diet-and-weight-loss>
- Heal, S., Papp, P., & Roychoudhury, P. (2022). Partnership between data science and behavioural economics. In A. Samson (Ed.), *The Behavioral Economics Guide 2022* (pp. 27–37). <https://www.behavioraleconomics.com/be-guide/>
- Hershfield, H. E., Shu, S., & Benartzi, S. (2020). Temporal reframing and participation in a savings program: A field experiment. *Marketing Science*, 39(6), 1039–1051.
- Horton, J. J. (2023). Large language models as simulated economic agents: What can we learn from homo silicus? (No. w31122). *National Bureau of Economic Research*. <https://www.nber.org/papers/w31122>
- Hunter, R. F., Tang, J., Hutchinson, G., Chilton, S., Holmes, D., & Kee, F. (2018). Association between time preference, present-bias and physical activity: Implications for designing behavior change interventions. *BMC Public Health*, 18, 1–12. <https://doi.org/10.1186/s12889-018-6305-9>
- IFIC. (2024). *IFIC food & health survey*. International Food Information Council, Washington, D.C. <https://foodinsight.org/wp-content/uploads/2024/06/2024-IFIC-Food-Health-Survey.pdf>
- Johnson, H., & Leary, J. (2017). Policy watch: Research priorities on disclosure at the Consumer Financial Protection Bureau. *Journal of Public Policy & Marketing*, 36(1), 184–191. <https://doi.org/10.1509/jppm.17.025>
- Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics & bias-*

- es: The psychology of intuitive judgment (pp. 49–81). Cambridge University Press. <http://dx.doi.org/10.1017/CBO9780511808098.004>
- Krajich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, 13(10), 1292–1298. <https://doi.org/10.1038/nn.2635>
- Lim, J. S., Chock, T. M., & Golan, G. J. (2020). Consumer perceptions of online advertising of weight loss products: The role of social norms and perceived deception. *Journal of Marketing Communications*, 26(2), 145–165. <https://doi.org/10.1080/13527266.2018.1469543>
- Madrian, B. C. (2014). Applying insights from behavioral economics to policy design. *Annual Review of Economics*, 6(1), 663–688. <https://doi.org/10.1146/annurev-economics-080213-041033>
- Meeker, D., Linder, J. A., Fox, C. R., Friedberg, M. W., Persell, S. D., Goldstein, N. J., Knight, T. K., Hay, J. W., & Doctor, J. N. (2016). Effect of behavioral interventions on inappropriate antibiotic prescribing among primary care practices: A randomized clinical trial. *Jama*, 315(6), 562–570. <https://doi.org/10.1001/jama.2016.0275>
- Michie, S., Van Stralen, M. M., & West, R. (2011). The behaviour change wheel: A new method for characterising and designing behaviour change interventions. *Implementation Science*, 6, 1–12. <https://doi.org/10.1186/1748-5908-6-42>
- Milkman, K. L., Gandhi, L., Patel, M. S., Graci, H. N., Gromet, D. M., Ho, H., ... & Duckworth, A. L. (2022). A 680,000-person megastudy of nudges to encourage vaccination in pharmacies. *Proceedings of the National Academy of Sciences*, 119(6), e2115126119. <https://doi.org/10.1073/pnas.2115126119>
- Mogler, B. K., Shu, S. B., Fox, C. R., Goldstein, N. J., Victor, R. G., Escarce, J. J., & Shapiro, M. F. (2013). Using insights from behavioral economics and social psychology to help patients manage chronic diseases. *Journal of General Internal Medicine*, 28, 711–718. <https://doi.org/10.1007/s11606-012-2261-8>
- Munshi, S., & Ramani, P. (2023). The next frontier of personalization: Behavioral science customer segmentation in financial services. In A. Samson (Ed.), *The Behavioral Economics Guide 2023* (pp. 86–93). <https://www.behavioraleconomics.com/be-guide/>
- Nixon, P., & Gilbert, E. (2024). How machine learning can reduce the behaviour tax by informing hyper-personalised nudges. In A. Samson (Ed.), *The Behavioral Economics Guide 2024*, (pp. 74–81). <https://www.behavioraleconomics.com/be-guide/>
- Ong, A. S. J., Frewer, L., & Chan, M. Y. (2017). Cognitive dissonance in food and nutrition: A review. *Critical Reviews in Food Science and Nutrition*, 57(11), 2330–2342. <https://doi.org/10.1080/10408398.2015.1013622>
- Prada, M., Godinho, C., Rodrigues, D. L., Lopes, C., & Garrido, M. V. (2019). The impact of a gluten-free claim on the perceived healthfulness, calories, level of processing and expected taste of food products. *Food Quality and Preference*, 73, 284–287. <https://doi.org/10.1016/j.foodqual.2018.10.013>
- Priven, M., Baum, J., Vieira, E., Fung, T., & Herbold, N. (2015). The influence of a factitious free-from food product label on consumer perceptions of healthfulness. *Journal of the Academy of Nutrition and Dietetics*, 115(11), 1808–1814. <https://doi.org/10.1016/j.jand.2015.03.013>
- Quinn, C. C., Shardell, M. D., Terrin, M. L., Barr, E. A., Ballew, S. H., & Gruber-Baldini, A. L. (2011). Cluster-randomized trial of a mobile phone personalized behavioral intervention for blood glucose control. *Diabetes Care*, 34(9), 1934–1942. <https://doi.org/10.2337/dc11-0366>
- Rahman, S., Zasadzinski, L., Zhu, L., Edirisinghe, I., & Burton-Freeman, B. (2020). Assessing consumers' understanding of the term “Natural” on food labeling. *Journal of Food Science*, 85(6), 1891–1896. <https://doi.org/10.1111/1750-3841.15128>
- Reeck, C., Wall, D., & Johnson, E. J. (2017). Search predicts and changes patience in intertemporal choice. *Proceedings of the National Academy of Sciences*, 114(45), 11890–11895. <https://doi.org/10.1073/pnas.1707040114>
- Reutskaja, E., Nagel, R., Camerer, C. F., & Rangel, A. (2011). Search dynamics in consumer choice under time pressure: An eye-tracking study. *American Economic Review*, 101(2), 900–926. <https://doi.org/10.1257/aer.101.2.900>
- Schotter, A., & Trevino, I. (2021). Is response time predictive of choice? An experimental study of

- threshold strategies. *Experimental Economics*, 24(1), 87–117. <https://doi.org/10.1007/s10683-020-09651-1>
- Schultz, P. W., Nolan, J. M., Cialdini, R. B., Goldstein, N. J., & Griskevicius, V. (2007). The constructive, destructive, and reconstructive power of social norms. *Psychological Science*, 18(5), 429–434.
- Semmelmann, K., & Weigelt, S. (2018). Online webcam-based eye tracking in cognitive science: A first look. *Behavior Research Methods*, 50, 451–465. <https://doi.org/10.3758/s13428-017-0913-7>
- Shah, A. M., Osborne, M., Lefkowitz Kalter, J., Fertig, A., Fishbane, A., & Soman, D. (2023) Identifying heterogeneity using recursive partitioning: Evidence from SMS nudges encouraging voluntary retirement savings in Mexico, *PNAS Nexus*, 2(5), pgad058. <https://doi.org/10.1093/pnasnexus/pgad449>
- Shapiro, M. F., Shu, S. B., Goldstein, N. J., Victor, R. G., Fox, C. R., Tseng, C. H., Vangala, S., Mogler, B. K., Reed, S. B., Villa, E., & Escarce, J. J. (2020). Impact of a patient-centered behavioral economics intervention on hypertension control in a highly disadvantaged population: A randomized trial. *Journal of General Internal Medicine*, 35(1), 70–78. <https://doi.org/10.1007/s11606-019-05269-z>
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II Perceptual learning, automatic attending and a general theory. *Psychological Review*, 84, 127–190. <http://dx.doi.org/10.1037/0033-295X.84.2.127>
- Shu, S. B., Zeithammer, R., & Payne, J. W. (2016). Consumer preferences for annuity attributes: Beyond net present value. *Journal of Marketing Research*, 53(2), 240–262.
- Shu, S. B. (2018). Psychological ownership in financial decisions. *Psychological ownership and consumer behavior*, 165–176.
- Shu, S. D., Hershfield, H., Mason, R., & Benartzi, S. (2022). Reducing savings gaps through pennies versus percent framing. SSRN. <https://dx.doi.org/10.2139/ssrn.4042843>
- Sopher, J. (2004). Weight-loss advertising too good to be true: Are manufacturers or the media to blame? *Cardozo Arts Ent Law J*, 22(3), 933–964.
- Spiegler, R. (2020). Behavioral implications of causal misperceptions. *Annual Review of Economics*, 12(1), 81–106. <https://doi.org/10.1146/annurev-economics-072219-111921>
- Stillman, P. E., Krajbich, I., & Ferguson, M. J. (2020). Using dynamic monitoring of choices to predict and understand risk preferences. *Proceedings of the National Academy of Sciences*, 117(50), 31738–31747. <https://doi.org/10.1073/pnas.2010056117>
- Stein, A. J., & Santini, F. (2022). The sustainability of “local” food: A review for policy-makers. *Review of Agricultural, Food and Environmental Studies*, 103(1), 77–89. <https://doi.org/10.1007/s41130-021-00148-w>
- Sullivan, N., Hutcherson, C., Harris, A., & Rangel, A. (2015). Dietary self-control is related to the speed with which attributes of healthfulness and tastiness are processed. *Psychological Science*, 26(2), 122–134. <https://doi.org/10.1177/0956797614559543>
- Thaler, R. H., & Sunstein, C. R. (2009). *Nudge: Improving decisions about health, wealth, and happiness*. Penguin.
- Trudel, R., Murray, K. B., Kim, S., & Chen, S. (2015). The impact of traffic light color-coding on food health perceptions and choice. *Journal of Experimental Psychology: Applied*, 21(3), 255–275.
- Vladeck, D. C. (2000). Truth and consequences: The perils of half-truths and unsubstantiated health claims for dietary supplements. *Journal of Public Policy & Marketing*, 19(1), 132–138. <https://doi.org/10.1509/jppm.19.1.132.16948>
- Wang, J. T., Spezio, M., & Camerer, C. F. (2010). Pinocchio's pupil: Using eyetracking and pupil dilation to understand truth telling and deception in sender-receiver games. *American Economic Review*, 100(3), 984–1007. <https://doi.org/10.1257/aer.100.3.984>
- Yang, X., & Krajbich, I. (2021). Webcam-based online eye-tracking for behavioral research. *Judgment and Decision-making*, 16(6), 1485–1505. <https://doi.org/10.1017/S1930297500008512>



APPLICATIONS

A New Tool for Behavioral Scientists: Exploring the Potential of Synthetic Participants to Accelerate How We Study Attitudes and Behaviors

ANTONIO HANNA-AMODIO¹, CHIARA CAPPELLINI, RASHA ATTAR
AND AILIDH FINLAYSON

Behavioral Science Group

This study explores whether synthetic participants—AI-generated proxies built from real demographic data—can offer a practical tool for behavioral scientists. Conducted by the UAE’s Behavioral Science Group (BSG), the project used the open-source Falcon model to simulate responses from a nationally representative sample ($n = 7,652$), and compared experimental results against real data from two online experiments ($n = 4,577$). Results suggest that synthetic participants can accurately approximate public attitudes and opinions in diverse, non-WEIRD populations. However, predicting behavioral outcomes was more mixed. While synthetic participants often correctly identified the best and worst performing interventions, they tended to overestimate effect sizes. In one case, though, their predictions beat expert forecasts. Synthetic participants aren’t a silver bullet, but they do show promise as a new tool in the behavioral science toolkit. As methods evolve, they could help policymakers sense-check early ideas, prioritize what to test, and accelerate the design of effective interventions.

What Role Can Synthetic Participants Play in Behavioral Science?

Artificial intelligence (AI) is set to transform how we study human behavior (Flahavan, 2024). One promising application is synthetic participants, which can be understood simply as using data from human respondents to generate AI-powered proxies that can simulate opinions and behaviors.

For example, a demographic dataset can be used to generate synthetic participants that simulate how real people might respond to questions on topics of interest. Each synthetic participant is created based on known characteristics, such as demographics and attitudinal and behavioral data. Combined, these participants form a synthetic sample that can provide policymakers with near-instant insights.

Synthetic participants offer a new way to explore attitudes, pre-test interventions, and study behavioral patterns at scale. Early research suggests that AI-generated responses can approximate real human attitudes and behaviors

(Argyle et al., 2023; Horton, 2023).

The appeal of synthetic participants is intuitive. They allow researchers to generate instant insights relatively cheaply (Argyle et al., 2023). Open-source datasets can be used to create synthetic datasets, in turn answering new research questions. Real-world fieldwork, in contrast, often involves lengthy participant recruitment and attrition. If found to be sufficiently accurate, synthetic participants could become a valuable new methodology, augmenting, rather than automating, typical working methods used in social science (Handa et al., 2025).

Despite their potential, their use and adoption remain underexplored, and more research is needed to advance this approach. For instance, most work in this space has focused on opinion-polling rather than behavioral responses, which means little is known about whether synthetic participants can predict the effects of policies and interventions that aim to change behavior. And yet, behavioral science is ultimately concerned with understanding and

¹ Corresponding author: antonio.hanna-amodio@bsg.ae

shifting behavior, not just measuring attitudes. This, therefore, is one important gap to explore.

Understanding UAE Attitudes and Behavior With Falcon

This paper outlines a project led by the Behavioral Science Group (BSG), a specialist team within the UAE's Office of Development Affairs dedicated to creating positive societal impact through the application of behavioral science. The BSG draws on data about national attitudes and behaviors (collected through surveys, trials, and collaborative research) to inform and test new methodological approaches, such as the use of synthetic participants

The project drew on large-scale government data sources in the UAE ($n=7,652$), combined with data from online experiments conducted with nationally representative samples ($n=4,577$). Together, these provided a comprehensive foundation for understanding behavioral trends and preferences within the UAE, a highly heterogeneous country where around 90% of the population are expatriates from over 200 nations.

The study was conducted in partnership with AI71², using Falcon 180B³—the UAE's leading large language model. Crucially, Falcon is open source, which matters for policymakers because, unlike closed models, open-source systems can be deployed on local hardware, thereby offering greater control and data security. This is no small advantage in a world where data security is an increasingly important consideration. Notably, this specific study was conducted as a pioneer use case for the AI71 Platform, which offers greater ease and flexibility to researchers.

The use of Falcon has additional benefits. Most existing research has relied on Open AI's GPT-based models, which are trained primarily on Western-centric datasets (Argyle et al., 2023; Shrestha et al., 2025). Falcon has been shown to compare favorably to other open-source models (Kemper, 2024), thus offering a potentially more regionally relevant alternative for generating and evaluating synthetic participants.

Our Three Guiding Questions for Exploration

Creating, tweaking, and testing synthetic audiences is still an evolving field. Considering the nascent but growing literature, this project explored both synthetic participants' potential and current limitations in terms of where they might already offer value for policymakers and where further research is needed. We therefore focused on key research questions to generate practical insights for real-world applications.

Question 1: Can Synthetic Participants Simulate the Attitudes of Non-WEIRD Audiences?

Synthetic participants have shown promise in accurately simulating the attitudes of real-life populations (Argyle et al., 2023). However, existing research has largely focused on WEIRD⁴ populations, especially in the US, raising concerns about how well synthetic participants generalize to other contexts (von der Heyde et al., 2024). This trend is consistent with a wider over-reliance on WEIRD samples in the social sciences (Henrich, 2010; Muthukrishna et al., 2020).

A notable exception is Shrestha et al. (2025), who compared synthetic and human survey responses in the UAE, Saudi Arabia, and the United States. Using GPT-4, they found that alignment between synthetic and real responses was generally weaker in Gulf regions compared to the US, suggesting that synthetic predictions may generalize less well across different cultural contexts. These findings highlight the need for further research to understand better if accuracy might decline in non-Western populations and, further, whether alternative models or conditioning techniques can improve reliability in diverse settings.

This study aims to take a step forward by creating a large, representative synthetic sample of the UAE population and evaluating its predictive accuracy, relative to existing survey data.

Question 2: Can Synthetic Participants Accurately Reflect Public Views on Policy?

As practitioners, our primary interest lies in understanding how synthetic participants can be used to inform real-world policymaking. One

² <https://ai71.ai/>

³ <https://falconllm.tii.ae/falcon-models.html>

⁴ Western, Educated, Industrialized, Rich and Democratic

particularly exciting use case is assessing public sentiment towards future policy. Traditionally, this requires surveys or focus groups, which are often time-consuming and costly, especially when quick feedback is needed.

If synthetic participants can reliably simulate how a population might respond to an existing or a proposed policy, this could offer a powerful tool to sense-check opinion, signal areas of likely resistance or support, and guide early decision-making—prior to real-world testing.

Question 3: Can Synthetic Participants Predict Experimental Outcomes?

Most studies evaluate AI-generated responses in attitudinal settings, making it unclear whether synthetic participants can accurately predict real-world behavioral responses to interventions (Bisbee et al., 2024; Durmus et al., 2024).

This study moves beyond attitudes by testing whether synthetic participants can predict similar outcomes from two previously run online experiments conducted by the BSG with nationally representative samples.

The first experiment focused on health promotion, testing messages designed to encourage sign-ups

to an exercise event ($n=4,397$). The second targeted climate behavior, using behavioral framings to encourage reduced air-conditioning use ($n=4,577$).

Methods for Creating Synthetic Participants

Single-Step Prompting vs Pen Portrait Prompting

Generating synthetic participants is currently as much an art as it is a science. While research is evolving, we follow early best practices identified in the literature. Our study employs two distinct methods for generating synthetic participants: single-step prompting and pen portrait prompting.

Single-step prompting: Synthetic participants are generated using a single prompt containing structured input variables ($n = 7,652$). These variables were drawn from nationally representative UAE government datasets and include demographic and attitudinal information spanning wellbeing, social values, and climate attitudes. Each prompt is phrased in the second person (e.g., ‘You live in the UAE. You are female. You are 29 years old’).

The model is then asked to generate responses to survey questions or behavioral interventions based on that identity. This method is relatively more computationally efficient and requires less time.




Figure 1: Raw data prompting vs Pen Portraits approach.

Pen portrait prompting: Under this approach, the Falcon LLM is given the same input variables but instructed to generate a narrative “pen portrait”: a brief but detailed description of an individual based on these attributes. These synthetic personas are then fed into the model, which is prompted to predict how that individual would respond to survey questions or behavioral interventions.

A growing body of research suggests that this type of conditioning—providing a broader personal context—can improve the realism and reliability of LLM-generated responses (Argyle et al. 2013; Flahavan, 2023). This two-step process is comparatively more resource-intensive; however, it was a hypothesis that it could produce more accurate responses.

Evaluating Synthetic Participants

Currently, there is no widely accepted metric for evaluating the accuracy of synthetic participants. Our evaluation approach compared synthetic response distributions to the real data distributions. This is a more rigorous test than simply comparing average scores. For example, on a response scale, if 40% of real respondents strongly agree and 40% strongly disagree, a model producing mostly neutral

responses might deceptively seem to have a high average score, albeit whilst overlooking important opinion polarization.

Our protocol for evaluating synthetic participants combined three approaches:

Simple results comparison: Overlaying bar charts and density plots of synthetic vs real survey data to check for alignment and divergence. Checking whether distributions are directionally correct provided an initial and intuitive first-pass assessment of accuracy.

Similarity score calculation: This metric was derived directly from the Jensen-Shannon (JS) divergence, a statistical measure that compares probability distributions between real and synthetic data (Apellániz et al., 2024).

Specifically, the similarity score is $1 - \text{JS divergence scores}$, which range from 0 to 1. A score of 1 indicates perfect alignment between synthetic and real distributions, while 0 represents complete divergence. This bounded similarity score allows for easier comprehension and comparison across different questions and datasets.

Reproducing experimental results: The final test tried to simulate the results of two online RCTs conducted with real-life representative samples from the UAE. We tested whether synthetic participants can (1)

Meet Ahmed, a 30-year-old Emirati man living in Abu Dhabi with his wife and two children. He was born and raised in the UAE and is proud of his Emirati heritage. Ahmed is outgoing and friendly, always willing to help others.

Ahmed comes from a large family and is very close to his parents and siblings. He values family above all else and loves spending time with his wife and children. Ahmed is also a devout Muslim, and his faith is an important part of his life.

Ahmed works as a manager at a local construction company, where he is highly respected by his colleagues for his leadership skills and work ethic. He completed his secondary education in Abu Dhabi and is always looking for ways to improve himself personally and professionally.

Ahmed enjoys playing football and going on family outings to the beach in his free time. He is also an avid reader and enjoys learning about new cultures and ideas.

Figure 2: Pen Portrait example.

correctly identify the most and least effective interventions and (2) approximate the relative effect sizes observed in online experiments with real participants.

The synthetic participants were randomly assigned to control or treatments and then prompted to predict actual behavioral outcomes, such as clicking on a link. In the original RCTs, the efficacy of interventions was evaluated by running an appropriate regression model, hence controlling for demographics. The same analysis was replicated on the synthetic datasets.

Results

Result 1: Synthetic Participants Show Promise in Simulating Attitudinal Responses of Non-WEIRD Audiences

Overall, synthetic participants produced results that closely mirrored their human counterparts. Across a range of questions, AI-generated responses broadly reflected real-world distributions, though some variation did emerge, as illustrated by the figures that follow. The fact that these responses came from a nationally representative sample of the UAE—a population that is both highly diverse and predominantly non-WEIRD—indicates the potential for using synthetic methods to replicate diverse, heterogeneous samples.

Contrary to expectations, pen portrait prompting did not, on average, result in greater accuracy than single-step prompting. The single-step prompting approach produced a similarity score of 0.734, while the pen portrait prompting approach produced a score of 0.730. Given that pen portrait prompting requires additional computational resources and time

(e.g., increased API usage and processing costs), its practical value, in this case, was limited.

Application: Simulating Perceptions on Artificial Intelligence

An interesting (and appropriately meta) case study explored whether AI-generated participants can predict what real humans think about AI. When asked whether AI will mostly help or harm people in the UAE over the next 20 years, synthetic participants and real people produced strikingly similar distributions. In both cases, the vast majority of respondents indicated that AI will mostly help people, with only a small proportion believing it will cause harm.

Result 2: Synthetic Participants Show Promise in Simulating a Public Response to Behavior Change Policies

A key advantage of synthetic participants is their ability to rapidly gauge public opinion on live policy issues, which in turn creates opportunities to sense-check levels of support or opposition quickly. Looking forward, one could imagine testing how public sentiment shifts in response to changes in policy framing, helping policymakers better anticipate how their proposals will be received.

Application: Simulating Public Opinion on the Single-Use Plastic Bag Ban

The UAE's ban on single-use plastic bags, announced in 2023, provided a relevant test case. Here, too, synthetic participants demonstrated a response similar to human respondents, i.e., overwhelming support for the policy. Interestingly, synthetic

These days, there are machines that are able to perceive, synthesize and infer information on their own to make decisions - known as artificial intelligence (AI). Do you think that AI will mostly HELP or mostly HARM people in the UAE in the next 20 years?



Figure 3: Public perceptions of AI.




Figure 4: Support for single-use plastic ban.

participants produced slightly more extreme predictions, forecasting higher support and lower opposition than actual public opinion (see Figure 4). This suggests that synthetic responses align well with general trends but may amplify dominant sentiment—a pattern observed in other studies (Bisbee et al., 2024).

Interestingly, in our study, synthetic participants performed less accurately on climate-related questions than on other topics. Unlike the general pro-sustainability bias reported by Shrestha et al. (2025), the pattern here was less consistent. In one case, synthetic participants expressed a stronger belief in climate change than real respondents; in another, they underestimated how many people were willing to make significant lifestyle changes to address it.

Result 3: Predicting Behavioral Outcomes Produced More Mixed Results

Results for behavioral measures were more mixed than those observed in attitudinal and opinion-based simulations. Positively, our synthetic panel could broadly distinguish between more and less effective treatments.

For example, they correctly identified the best and worst-performing treatment arms in three out of four outcomes we looked at. In the other case, however, the predicted top performer was actually the least effective treatment in the real-world experiment.

Application: Replicating an Online Experiment Testing Messages to Reduce AC Use

Here, we spotlight one of the two experiments in more detail to illustrate our findings. The original study, with 4,577 UAE residents, aimed to reduce air conditioning (AC) use.

Participants were randomly assigned to receive one of several messages, or no message at all. These included two leadership-framed messages (“Support the national vision”) and varied the call to action between “set your AC to 24°C” or “raise it by 1°”, an easy message offering simple temperature rules of thumb (for when home, away, or in empty rooms), and a health message warning that excessive AC use could cause dry skin or colds.

As illustrated in Figure 5, the results demonstrate the broader pattern we observed: synthetic participants consistently ranked the control group as least effective and showed strong alignment with real human participants for the most effective interventions. However, synthetic participants consistently overestimated effect sizes across all four outcomes, often by a large margin.

Notably, these experimental findings should be interpreted as preliminary. They are based on two online experiments, with most outcomes measuring stated intentions rather than actual behavior. These simulations obviously also lack a “behavioral environment,” i.e., the real-world situational cues and constraints that shape how people act in the moment (Lewin, 1951). More research is therefore needed to understand when and how synthetic participants can reliably predict the real-world effects of behavioral interventions.

Policy Implications

Can Synthetic Audiences Be a New Addition in the Behavioral Scientist Toolkit?

Emerging evidence, including our study, suggests that one of the most practical applications of synthetic participants is rapidly assessing public opinion and




Figure 5: Actual vs AI predicted results for AC experiment. X axis labels refer to different framing treatments.

attitudes, especially in contexts constrained by time or resources.

While never a substitute for real-world data collection, synthetic responses can provide an additional layer of insight to inform early decision-making. They do so by offering a way to sense-check assumptions, narrow down options, and guide the design of more targeted research.

In this way, synthetic participants can play a valuable role during the scoping and pre-trial phase, helping ensure that limited resources are directed where they are likely to have the greatest impact.

When Is Good “Good Enough”?

Our results suggest the value of this approach lies in providing directionally correct insights, helping to determine general sentiment on a given issue. Indeed, accuracy seemed to improve when response scales were collapsed. In other words, aggregating categories to focus on broad sentiment, rather than fine distinctions, can sketch a quick picture of the overall story.

For example, when assessing life satisfaction (see Figure 6), synthetic participants more closely aligned with real-world data when responses were grouped into ‘Satisfied’ vs ‘Not Satisfied’, rather than using a more granular scale. This is perhaps unsurprising and aligns with broader findings in behavioral science. As Kahneman et al. (2021, p.199) note, ‘ambiguity

in the wording of scales is a general problem’, thus making precise measurement of similar sentiments inherently noisy.

Of course, in real-world applications, the optimal grouping may not be known in advance. Nevertheless, this points to a practical implication: simpler scales that capture broad sentiment—such as positive, neutral, or negative—may offer more reliable early-stage insights.

Better Than Practitioner Intuition Alone?

When designing behavioral interventions, decisions are often guided by human intuition. Sometimes, that intuition is grounded in deep expertise, but assumptions can be flawed, and results are often surprising. On the spectrum of robustness, intuition alone sits at the lower end, while empirical evaluation—such as randomized trials—remains the gold standard. Synthetic participants may sit somewhere in between.

In this study, synthetic participants were evaluated by comparing their responses to real-world data, mostly on opinions and attitudes. However, if their primary value is in the early stages of intervention design, the more meaningful comparison may be to practitioner judgment: specifically, decisions about which ideas to pursue and which are expected to be most effective. If synthetic data can support or enhance that judgment, it could become a valuable tool for refining and prioritizing ideas before field trials.




Figure 6: Illustration of collapsing response scales.

Human vs Machine Predictions of Experimental Outcomes

For the air-conditioning study, we had access to pre-trial predictions from five behavioral scientists and policymakers on the efficacy of the interventions. As shown in Figure 5, both synthetic datasets accurately identified the top 3 interventions, with pen portrait prompting accurately ranking all interventions.

In comparison, none of the five practitioners correctly predicted which arm would be the top performer. Additionally, three of the five inaccurately forecasted that telling people about potential health implications would be the most effective motivator for them to reduce their AC use. In reality, this was the worst performing of the behavioral messages.

In short: synthetic participants were more accurate in this scenario at identifying the arms that would perform better than the practitioners.

Synthetic Pre-Testing: A Tool to Refine and Short-List Ideas

Looking ahead, synthetic participants could play a key role in behavioral policy design. While our study found mixed results, other research has suggested potential value in using AI-generated participants to predict intervention effectiveness (Hewitt et al., 2024). This approach could be integrated as an additional step in the typical behavioral science project workflow.

Consider an illustrative example of a campaign aimed at boosting childhood vaccination rates—at a time when immunization rates have fallen globally (UNICEF, 2023). The process could begin with a literature review to identify broad themes of effective vaccination messaging (Step 1), followed by AI-generated synthetic participants pre-testing 20 message variants (Step 2). The top 10 messages could then be tested in an online experiment with human participants (Step 3) before selecting the best five performers for real-world field testing (Step 4).

Synthetic participants can aid early-stage exploration, without replacing direct observation or field testing, which often reveal gaps between intention and action. That said, incorporating synthetic pre-testing could help researchers and policymakers quickly filter out weaker interventions, saving time and resources before conducting large-scale trials.

Conclusion

Synthetic participants are no silver bullet, but they do offer genuine promise. In this study, we found that AI-generated responses can meaningfully approximate human survey data on attitudes and opinions for non-WEIRD populations. Predicting behavioral outcomes was more mixed—yet encouraging in parts—revealing both this emerging approach's potential and its current limitations.

While real-world data collection is and will remain essential, synthetic participants may accelerate and augment this process by providing a fast, low-cost way to sense-check public sentiment or pre-test interventions early in the policy design process. As methods improve, so will the opportunities to embed synthetic participants into the behavioral insights toolkit in practical and complementary ways.

THE AUTHORS

Antonio Hanna-Amodio is a policy advisor at the Behavioral Science Group (BSG), where he leads work at the intersection of artificial intelligence and human behavior. Beyond this, he designs behaviorally informed policies to address societal challenges in diverse cultural and contextual settings. He holds a BA in Economics and World Philosophies from SOAS University of London and an MSc in Behavioural Science from the London School of Economics and Political Science.

Chiara Cappellini is the Senior Behavioral Scientist for the Behavioral Science Group, where she oversees strategic policy direction. She took on this exciting role after years of working with BIT designing and delivering behavior change projects globally. Prior to BIT, Chiara worked as an evaluator in the public, private, and third sectors. She holds a BSc in Social Policy from the London School of Economics, and an MSc in International Development and Research Methods from the University of Edinburgh.

Rasha Attar is Director of the BSG and the wider insights portfolio at the Office of Development Affairs. She has over 15 years of experience delivering research and impact for the UAE government and spearheaded founding the BSG as the central government's behavioral science unit. She holds a degree in Economics from the American University of Beirut, has completed the Harvard Kennedy School Behavioral Insights and Public Policy Program, and holds an Executive MSc in Behavioral Science at the London School of Economics.

Ailidh Finlayson is a researcher at the Behavioral Science Group (BSG), where she uses advanced quantitative methods to design experiments and generate insights across a range of policy areas. Her work includes the development and analysis of randomized controlled trials as well as other

experimental designs. She holds an MRes in Advanced Quantitative Methods from the University of Bath and a BA in Experimental Psychology from the University of Oxford.

REFERENCES

- Apellániz, P. A., Jiménez, A., Galende, B. A., Parras, J., & Zazo, S. (2024). Synthetic tabular data validation: A divergence-based approach. *arXiv Preprint*. <https://arxiv.org/abs/2405.07822>
- Argyle, L. P., Busby, E. C., Fulda, N., Gubler, J. R., Rytting, C., & Wingate, D. (2023). Out of one, many: Using language models to simulate human samples. *Political Analysis*, 31(3), 337–351.
- Arora, N., Chakraborty, I., & Nishimura, Y. (2024). Revolutionizing marketing research with a large language model: A hybrid AI-human approach. SSRN. <https://doi.org/10.2139/ssrn.4683054>
- Bisbee, J., Clinton, J. D., Dorff, C., Kenkel, B., & Larson, J. M. (2024). Synthetic replacements for human survey data? The perils of large language models. *Political Analysis*, 32(4), 401–416. <https://doi.org/10.1017/pan.2024.5>
- Durmus, E., Nyugen, K., Liao, T., Schiefer, N., Askill, A., Bakhtin, A., Chen, C., Hatfield-Dodds, Z., Hernandez, D., Joseph, N., Lovitt, L., McCandlish, S., Sikder, O., Tamkin, A., Thamkul, J., Kaplan, J., Clark, J., & Ganguli, D. (2023). Towards measuring the representation of subjective global opinions in language models. *arXiv Preprint*. <https://arxiv.org/abs/2306.16388>
- Flahavan, E. (2024). BIT's roadmap for AI & BI: How AI is changing and improving the methodologies we use. *Behavioural Insights Team Blog*. <https://www.bi.team/blogs/bits-roadmap-for-ai-bi/>
- Flahavan, E. (2023). Can AI accurately simulate cross-national cultural values? An exploration. *Medium*. <https://medium.com/@ed.flahavan/can-ai-accurately-simulate-cross-cultural-values-an-exploration-e4637bfd8901>
- Handa, K., Tamkin, A., McCain, M., Huang, S., Durmus, E., Heck, S., Mueller, J., Hong, J., Ritchie, S., Belonax, T., Troy, K. K., Amodei, D., Kaplan, J., Clark, J., & Ganguli, D. (2025). Which economic tasks are performed with AI? Evidence from millions of Claude conversations. *arXiv Preprint*. <https://arxiv.org/pdf/2503.04761>

- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2-3), 61-135. <https://doi.org/10.1017/S0140525X0999152X>
- Hewitt, L., Ashokkumar, A., Ghezze, I., & Willer, R. (2024). Predicting results of social science experiments using large language models. *Preprint*. Stanford University & New York University.
- Horton, J. J. (2023). Large language models as simulated economic agents: What can we learn from Homo Silicus? *National Bureau of Economic Research*. <https://www.nber.org/papers/w31122>
- Kahneman, D., Sibony, O., & Sunstein, C. R. (2021). *Noise: A flaw in human judgment*. HarperCollins.
- Kemper, J. (2024). Falcon 3 series sets new benchmarks for open-source LLMs on a single GPU. *The Decoder*. <https://the-decoder.com/abu-dhabi-tii-releases-falcon-3-setting-new-benchmarks-for-efficient-open-source-ai-models/>
- Lewin, K. (1951). *Field theory in social science: Selected theoretical papers* (D. Cartwright, Ed.). Harpers.
- Muthukrishna, M., Bell, A. V., Henrich, J., Curtin, C. M., Gedranovich, A., McInerney, J., & Thue, B. (2020). Beyond Western, Educated, Industrial, Rich, and Democratic (WEIRD) psychology: Measuring and mapping scales of cultural and psychological distance. *Psychological Science*, 31(6), 678-701. <https://doi.org/10.1177/0956797620916782>
- Neekhara, B., Kapoor, K., & Gupta, D. (2023). SYNTHPOP++: A hybrid framework for generating a country-scale synthetic population. *arXiv Preprint*. <https://arxiv.org/abs/2304.12284>
- Santurkar, S., Durmus, E., Ladhak, F., Lee, C., Liang, P., & Hashimoto, T. (2023). Whose opinions do language models reflect? *arXiv Preprint*. <https://arxiv.org/abs/2303.17548>
- Shrestha, P., Krpan, D., Koaik, F., Schnider, R., Sayess, D., & Binbaz, M. S. (2025). Beyond WEIRD: Can synthetic survey participants substitute for humans in global policy research? *Behavioral Science & Policy*. <https://doi.org/10.1177/237946072413117>
- UNICEF. (2023). New data indicates declining confidence in childhood vaccines of up to 44 percentage points in some countries during the COVID-19 pandemic. *UNICEF Press Release*. https://www.unicef.org/press-releases/sowc_2023_immunization
- von der Heyde, L., Haensch, A.-C., & Wenz, A. (2024). Vox Populi, Vox AI? Using language models to estimate German public opinion. *arXiv Preprint*. <https://doi.org/10.48550/arXiv.2407.08563>

Symbolic Decision Support: Enhancing Financial Help-Seeking and Investment Decisions Using GenAI Imagery

CHRISTELLE T. NGNOUMEN¹ AND KIM H. NGUYEN

Voya Financial

The use of financial advice can help individuals make smarter financial decisions. So why aren't more people using financial information that is available to them? In an online experiment, we examined the effectiveness of imagery in promoting the use of financial guidance, and we established that imagery that addresses barriers to financial help-seeking can influence decisions to adopt financial guidance. Imagery was generated with the assistance of Generative AI (GenAI). We also determined that the use of guidance reduces investment mistakes on a hypothetical investment allocation task. This work demonstrates a collaborative framework between behavioral designers and GenAI, wherein individuals guide the conceptual and experimental direction while GenAI provides multiple visual options. The chapter closes with implications for financial investment and discusses the application of human-AI collaboration in content design and testing.

Introduction

Picture this: an employee at a new job is enrolling in a 401(k) plan and is scrolling through endless streams of text, data, and jargon. After selecting a portion of their paycheck to contribute to retirement, the employee is prompted to invest their contributions. They are presented with a list of investments to choose from. While some of the investments on the list are familiar, others are not. The employee becomes overwhelmed by the options. Before they can reach for any initial or familiar information to reduce their discomfort and quickly click “next,” an icon offering guidance appears. The employee decides to use this guidance and is encouraged to protect their portfolio from market fluctuations by spreading risk across different asset classes. A thoughtfully designed icon helps the employee take their next best step. It becomes the difference between satisficing, to quickly click “next,” and optimizing, to achieve the best possible outcome.

Over the past fifty years, a significant shift has occurred in the United States retirement system with employers moving away from offering defined benefit plans (i.e., pensions) to providing defined

contribution (DC) plans such as the 401(k). About half of private industry workers in the US participate in a DC plan (Cerulli Associates, 2024; Employee Benefit Research Institute, 2023). This shift has transferred important retirement decisions to individual employees, including when to start saving, how much to save, how long to save for, and how to invest their savings. As of 2024, there were more than 720,000 employer-sponsored 401(k) plans in the US, covering over 86 million active participants and millions of former employees and retirees (Investment Company Institute, 2024). Even with the advent and popular application of target date funds—a type of investment that automatically adjusts its asset allocation over time to help investors reach their retirement goals, facilitating a “set and forget” approach—a significant portion of private industry workers still choose to invest their savings themselves, thereby highlighting the importance of guidance and support for their investment decisions.

Only about 46% of US consumers have adopted financial technology despite high consumer awareness levels (Statista, 2024). While financial guidance is widely recognized as important, actual adoption

1 Corresponding author: christelle.ngnoumen@voya.com

rates vary depending on factors such as income levels and access to financial education resources (Amnas et al., 2023). The use of financial guidance can help consumers navigate complex decisions related to their savings, investments, and other financial topics, and it has been shown to improve investment performance, increase portfolio diversification, and increase financial literacy (Collins, 2012; Marsden et al., 2011; Mihaylov et al., 2015). Despite the extensive amount of evidence demonstrating a positive impact of financial guidance on financial wellness outcomes (Lusardi & Messy, 2023; Lusardi & Streeter, 2023; Mercado et al., 2024), there are many barriers to both seeking and using this help.

Barriers to Financial Help-Seeking

Common barriers to seeking financial advice include concerns about associated costs, misunderstanding the value proposition (i.e., what financial help involves and what one can gain from it), the motivation to save time, embarrassment (i.e., feeling ashamed about one's financial situation or about appearing incompetent), prescriptive norms (i.e., perceptions of what should or should not be done to manage finances), low financial literacy, and not trusting financial advisors (Westermann et al., 2020). There is therefore a need for evidence-based and innovative interventions to support people in seeking financial help.

Using Imagery to Drive Smarter Financial Decisions

Given the widespread use of visuals, it is surprising that finance literature has paid limited attention to imagery (Ronen et al., 2023). While text-based information is more commonly used on financial platforms, and more readable text facilitates comprehension and informed decisions (Loughran & McDonald, 2016; Tan et al., 2015), visual imagery can play a similar—if not more critical—role in facilitating financial decision-making. Research has highlighted the role of visual factors like salience, position, and size in guiding attention during decision-making and on decision outcomes (Orquin et al., 2021). Additional research has found an influence of color on financial decisions; for example, when financial losses are displayed in red, individuals tend to become more cautious and expect lower future stock returns (Bazley

et al., 2021). Research has also found that people anchor their evaluations to specific items in financial documents, leading to biased choices; in this case, for instance, if an annual fee is prominently displayed at the top of a document, it can disproportionately influence decision-making, even if other charges compensate for that amount (Ceravolo et al., 2022).

Visual imagery can help increase the visibility and adoption of financial help services. Beyond the demonstrated ease with which visualizations captivate our focus and direct our bottom-up attention (Carrasco, 2011; Corbetta & Shulman, 2002; Desimone & Duncan, 1995; Kahneman, 2011; Kastner & Ungerleider, 2000), additional benefits of leveraging this content form include its ability to act as a visual cue that can be quickly recognized and interpreted without relying solely on text, particularly when people skim copy. Symbolic imagery, such as icons, helps reduce cognitive load and facilitate decision-making by simplifying complex ideas in a glance. It can also improve accessibility by providing information intuitively, particularly for individuals with cognitive impairments or language barriers, and by allowing for descriptive text readable by screen readers and individuals with visual impairments (Homer & Gauntt 1992).

In an online experiment designed to simulate the digital experience of investment selection, the following research examined whether imagery addressing barriers to financial help-seeking can influence decisions to adopt financial guidance and optimize investment allocations. GenAI assisted in the creation of symbolic visual imageries.

Experiment Design

After progressing through a survey invitation email on a third-party research platform (*Qualtrics*), 700 participants (recruited from a research panel) were randomly assigned to one of eight experimental conditions that varied in terms of the imagery presented alongside a financial guidance prompt. The independent variable was the *symbolic imagery* associated with a guidance prompt shown prior to completing a hypothetical investment allocation task. Symbolic imagery consisted of five GenAI-based images designed to address five common barriers to financial help-seeking: financial literacy, overconfidence,




Figure 1: Experimental design. The eight experimental conditions and their respective symbolic imagery.

embarrassment, time, and prescriptive norms. The major dependent variables included *Guidance Adoption* (whether a participant decided to use guidance during the investment allocation task) and *Investment Mistakes* (the types and frequencies of investment mistakes made on the investment allocation task) (Figure 1).

Participants

The study included part- and full-time employees with access to employer benefits. Participant ages ranged from 23 – 88 years ($Meanage = 58$ years; $SDage = 11.37$). The sample was 59% male, 25% female, and 16% did not disclose their gender identification. In the sample, 56% were White, 6% Asian, 6% Black or African American, 6% Hispanic or Latino, 3% Multiracial, 1% Native and Indigenous Peoples, and 23% did not disclose their race and/or ethnicity. In terms of education, 14% had a high school degree or some college education, 40% had a college degree, and 35% had a graduate degree. Finally, 7% had an annual household income below \$50,000, 46% had an annual household income between \$50,000 and \$150,000, and 46% had an annual household income above \$150,000.

Hypothetical Investment Allocation Task

In an expanded adaptation of Hung and Yoong's (2013) computerized task, participants viewed a list

of ten investment funds for investing their money in a retirement savings account. The investment menu contained a range of fund options, including a money market fund, a bond market fund, a balanced fund, a large cap value fund, a large cap growth fund, a small cap value fund, a global fund, a real estate investment fund, a target date fund, and a commodity fund. Ten investment options were included to mirror the average number of investment options offered by 401(k) plans (Investment Company Institute, 2024). The average 10-year performance for each fund type was provided, and fees were kept the same across all funds. Participants were asked what percentage of their savings they would like to allocate to each fund. Prior to making their allocations, participants in seven of the eight conditions were asked if they would like to receive guidance while making their choices. The same text was used across all seven conditions (e.g., *Would you like to receive some guidance while making these choices?* [Yes; No]).

In six of the eight conditions, symbolic imagery related to barriers to financial help-seeking accompanied the guidance prompt. If a participant declined the use of this guidance, they would move on to a screen with the same ten investment funds and have a chance to enter any value between 0% and 100% for each fund. Their total allocations needed to amount to 100% across all funds. If a participant

Workers who enroll in an employer-sponsored retirement plan have a chance to contribute a portion of their paycheck for retirement, and select how to invest their contributions.

Suppose you were offered the following selection of funds for investing your money in a retirement savings account. Below is a table with information about the available funds and the annual rate of return for each fund over the past ten years. Annual fees are the same for all funds.

On the next screen we'll ask you what percentage of your money you would like to allocate to each fund.

Fund Choices	10 Year Return
Money Market Fund	5.03%
Bond Market Index Fund	4.45%
Balanced Fund	8.73%
Large Cap Value/Blend Fund	8.34%
Large Cap Growth Fund	18.68%
Small Cap Value Index Fund	10.50%
Global/International Equity Index Fund	10.50%
Real Estate Investment Trust Fund (REIT)	8.34%
Target Date Fund	8.41%
Commodity Fund (Energy)	8.87%

symbolic imagery
varied across
conditions



Would you like to receive some guidance while making these choices?

Yes ☐ No ☐

 Below are some general guidelines for investing:

- A zero balance in stocks is not recommended.
- A stock balance of less than 40% is considered conservative.
- Holding more than 95% in stocks is considered overly aggressive.
- A portfolio that is 100% in a single asset class is considered under diversified.

guidelines appeared only if a participant selected 'Yes' to receive guidance

Indicate the percentage of your money you would like to allocate to each fund. For each fund you can enter any value between 0% and 100%. Your total allocations must be 100%.

Money Market Fund	0 %
Bond Market Index Fund	0 %
Balanced Fund	0 %
Large Cap Value/Blend Fund	0 %
Large Cap Growth Fund	0 %
Small Cap Value Index Fund	0 %
Global Equity Index Fund	0 %
Real Estate Investment Trust Fund	0 %
Target Date Fund	0 %
Commodity Fund (Energy)	0 %
Total	0 %

Figure 2: A representation of the hypothetical investment allocation task, including instructions, ten investment fund types, and symbolic imagery addressing overconfidence-related barriers.

opted to receive guidance, they would move on to a screen similar to the one allocated to those who declined, albeit with the addition of a section on top of the screen that provided four guidelines for investing. The four guidelines were informed by Mottola and Utkus' (2009) identification of commonly accepted investment portfolio mistakes, including holding a zero balance in equities (i.e., not investing), holding an equity balance of less than 40% (i.e., being under-conservative), holding an equity balance more than 95% percent (i.e., being overly aggressive), and holding a portfolio that is 100% in a single asset class (i.e., being under-diversified) (Figure 2).

Questionnaire

The hypothetical investment allocation task was followed by a questionnaire assessing *Age, Gender, Race and Ethnicity, Education, Income, Investing Experience, Risk Tolerance, Financial Literacy, Financial Self-Efficacy, and Barriers to Advice-Seeking*. Demographic and psychographic factors did not differ by condition.

GenAI Image Development and Pilot Testing

Experimental imagery was created using CoPilot's Design tool. A specific set of prompts outlining each image's salient information and functional properties was designed and submitted by the authors to generate simplistic symbolic imagery. All prompts included instructions to generate imagery that was "simplistic," "grayscale," "round in shape," "included a dollar sign," and "included a supportive hand" to minimize differences across the images that were not related to key symbolic properties brought to the foreground. This helped increase the salience of key symbolic properties in the visualizations and was informed by research illustrating that viewers are biased by salient information in a visualization. It was also informed by research suggesting that images with groupings of three components or fewer facilitate feature integration and comprehension (Fabrikant et al., 2010; Hegarty et al., 2010; Padilla et al., 2017; Schirillo & Stone, 2005; Stone et al., 2003; Treisman & Gelade, 1980).

An original set of 12 images was evaluated by ten independent raters in terms of their perceived *Attractiveness, Familiarity, Credibility, and Quality*. Each rater evaluated all images (in randomized order) along the four dimensions. The five images

used in the experiment included those that scored comparably. The raters were also asked to describe what each image meant or represented to them, including the associations, thoughts, or feelings each image evoked. Pilot testing of the content suggested distinct, symbolic themes across the five images, corresponding to (1) *financial literacy* (e.g., "budgeting and money management;" "help with calculating needs and saving;" "help calculating your financial future"), (2) *overconfidence* (e.g., "a checklist where saving money is one [of many] components;" "a to-do list with my financial tasks;" "steps and help to improve my financial situation;" "knowing what to do next based on a list or plan"), (3) *embarrassment* (e.g., "unlocking a hidden door where money questions become clear;" "opening doors to a better financial future;" "a bright future"), (4) *time* (e.g., "this image represents revenue growth over time;" "organizing my finances in a timely way"), and (5) *prescriptive norms* (e.g., "the thumbs up reminds me of [social media platform];" "savings and positive outlook because of the thumbs up;" "better prepared for a wholesome financial future;" "team-based financial planning") (Figure 3).

The first hypothesis tested in this experiment was that guidance prompts, presented with imagery addressing barriers to financial help-seeking, increase guidance adoption during a hypothetical investment allocation task compared to a control. The second hypothesis tested was that the use of guidance optimizes decisions (i.e., reduce mistakes) on the hypothetical investment allocation task.

Results

There was a significant relationship between symbolic imagery and the decision to use guidance (or not), $\chi^2(5, N = 535) = 18.05, p = .003$ (Cramer's $V = 0.2$; $1 - \beta = .96$). Post-test comparisons (with Bonferroni correction) of the five symbolic imagery conditions against a control revealed a significantly greater proportion of guidance adoption within the Overconfidence condition compared to the No Icon Control condition ($p < .05$). The Prescriptive Norms condition drove significantly less guidance adoption compared to the No Icon Control condition ($p < .05$). The guidance adoption rates for the Time, Embarrassment, and Financial Literacy conditions were not significantly different from the No Icon Control condition, which



Figure 3: Symbolic imagery with corresponding functional properties, GenAI prompts, and evidence guiding content creation.

was not significantly different from chance (i.e., 50%) (Table 1).

The use of guidance was associated with fewer investment mistakes on the hypothetical investment allocation task, $\chi^2(1, N = 617) = 13.53, p < .001$ (Cramer's $V = 0.15$; $1-\beta = .96$). More specifically, individuals who used guidance demonstrated greater diversification ($p < .001$) and a lower proportion of overly aggressive portfolios compared to those who did not use it ($p < .01$).

Several demographic and psychographic factors predicted financial guidance adoption. Females were more likely to choose guidance compared to males, $p < .001$. Individuals whose highest level of education was high school or some college were more likely to choose guidance compared to those with college degrees or higher, $p < .05$. Individuals with little to no prior experience in investing were also more likely to choose guidance compared to those with intermediate or advanced experience, $p < .001$. Individuals with moderate risk tolerance were more likely to choose guidance compared to those with conservative or aggressive risk tolerance, $p < .001$. Those with lower financial literacy were more likely to choose guidance compared to those with higher

financial literacy, $p < .001$. Similarly, individuals with lower financial self-efficacy were more likely to choose guidance compared to those with higher financial self-efficacy, $p < .001$.

General Discussion

This research complements and expands on prior literature on the effects of visual imagery—as a content format incremental to text—on investors' decisions (Luffarelli et al., 2019; Ronen et al., 2023). It provides evidence of the impact of symbolic imagery on guidance adoption, and of guidance adoption in optimizing investment decisions, and it is one of few studies to incorporate GenAI as a collaborative tool in content intervention development, gradually harnessing AI's full potential in behavioral design and experimentation. GenAI helped speed up the experimental design process and the testing of behaviorally informed content to optimize financial choices. Visual imagery that is informed by an understanding of behavioral barriers can enhance user engagement with financial platforms, leading to increased interaction, better decision-making, and improved outcomes.

Table 1: Guidance Decision by Symbolic Imagery

Symbolic Imagery	Guidance Decision		Total (Row)
	Yes	No	
No Icon (Control)	54.4% (49)	45.6% (41)	100% (90)
Financial Literacy	44.1% (45)	55.9% (57)	100% (102)
Overconfidence *	67.5% (56)	32.5% (27)	100% (83)
Embarrassment	49.5% (45)	50.5% (46)	100% (91)
Time	56.5% (48)	43.5% (37)	100% (85)
Prescriptive Norms *	38.1% (32)	61.9% (52)	100% (84)
Total	51.4% (275)	48.6% (260)	100% (535)

Note: Counts in parentheses; $\chi^2 (5, N = 535) = 18.05, p = .003$

* Overconfidence > Control, $p < .05$; Prescriptive Norms < Control, $p < .05$

Addressing Overconfidence Bias in Financial Help-seeking

Prior work finds that visualizations can be used to reduce decision-making biases, including anecdotal evidence bias, side effect aversion, and risk aversion (Fagerlin et al., 2005; Waters et al., 2007; Weinstein et al., 2006). This work presents evidence that visualizations can also be used to mitigate confidence-related barriers to financial help-seeking. Overconfidence bias is a well-established cognitive error whereby individuals overestimate their abilities and knowledge, leading to poor decision-making, including under-diversification, excessive trading, and taking excessive risks. Nearly half (49%) of the current study's sample indicated that confidence in managing their finances on their own was a top barrier to seeking financial help. Self-reported confidence in one's ability to save and invest was related to the number of investment mistakes made on the hypothetical allocation task, suggesting that individuals with higher confidence demonstrated a greater proportion of mistakes (35%) compared to those with lower confidence (29%), $p = .05$. Behavioral designers and choice architects can address overconfidence bias by introducing reflective periods before critical decisions, encouraging the diversification of sources of information, and promoting self-checking.

Limitations

This study has limitations that should be considered when interpreting the findings. Firstly, convenience sampling limits the generalizability of the results to groups with similar characteristics to the research sample. Additionally, reliance on a hypothetical task means responses may differ from real-world behavior. That said, hypothetical performances seemed to align with behavioral patterns observed in the industry and replicated prior research findings (Hung & Yoong, 2013). Future research should use larger and more diverse samples to mitigate these issues. Despite these limitations, the study provides valuable insights into promoting financial help-seeking and financial guidance adoption.

Cognitive theory (Vessey, 1991) suggests that when viewing imagery, people compare a learned mental schema to the visual image. Visualizations that do not match the mental schema require cognitive transformations to align them as well as mental representations. When a viewer is forced to mentally transform a visualization to match their mental schema, this increases the demands on their working memory as well as mental processing steps and task completion time (Evans & Stanovich, 2013; Kahneman & Frederick, 2002; Lohse, 1997). In the current study, individuals in the condition involving an icon signaling prescriptive norms adopted financial guidance

at a significantly lower rate than individuals in a control group. While task decision times, perceived attractiveness, credibility, familiarity, and quality did not significantly differ across the icons used in the study, it is possible that another domain not captured during pre-testing of this icon introduced a greater degree of abstraction (i.e., interpretation) and impacted the participants' willingness to adopt financial guidance.

Design and Content Contributions

This research has several practical implications for behavioral designers and content creators. Based on the current findings, we recommend that designers of financial platforms aim to complement text-based content with symbolic, visual imagery that addresses known barriers to action. Furthermore, content and behavioral designers can capitalize on Type 1 processing to help viewers seek and adopt financial guidance by highlighting symbolic associations that are already held in long-term memory and are easier to apply automatically.

This early example of combining human judgment and GenAI creates opportunities for behavioral designers with design interests—but without specialized training—to speed up their iterative testing and learning efforts, particularly when design resources are limited. The availability of capabilities like GenAI can speed up design workflow and contribute to design practice, especially in environments that prioritize testing and learning over perfectionism. However, rather than replacing design professionals, this practice recommends still consulting with them along the way. This study involved consulting with our design partners to inform the conceptualization of imagery leveraging iconographic properties and modifying GenAI-based imagery. This may naturally lead us to the question of how to keep human judgment at the center of human-AI collaboration. We believe the answer lies in designing prompts.

Effective Prompt Design

Ultimately, placing human judgments at the center of the process starts with effective prompts that are structured with clarity and include multiple rounds of refinement. A couple of examples gained from this

experiment include:

- Specifying design elements: instructions detailing color schemes, number of components, relationships among components, and aesthetic preferences.
- Using affirmative phrasing in prompts: observations from the prompt creation process seem to suggest prompts that detail desired actions (e.g., *design an icon that has/does x, y, z*) were more effective than those focused on actions to avoid (e.g., *design an icon that does not/is not x, y, z*). Research in behavioral psychology supports this idea, i.e., that explicit, positive instructions lead to better outcomes than avoidance-based directives, since the former provides a clear roadmap for action while the latter may introduce more room for interpretation.
- Paying attention to (and correcting for) unintended features or biases: notable gender biases in the AI outputs for finance-related content were observed, and they frequently leaned toward male-coded imagery (e.g., depictions of figures with collared shirts, ties, or briefcases).

Implications for Communication

Our findings suggest broader implications for how we communicate in digital environments. Visual imagery already serves as an important communication cue that can be processed quickly and efficiently. As AI technology continues to advance, these visual images will become even more common and sophisticated. To avoid silos and ensure consistent interpretation, this growth will require proactive collaboration to not only enable replication and refinement, but also foster the collective development of best practices in AI-assisted visual communication design and testing. Furthermore, we recommend that future research compares GenAI-generated images and those created by humans. This comparison could provide not only valuable insights into the effectiveness and performance of these images, but also recommendations for enhancing collaboration between AI and human creators. By sharing this study, we encourage fellow researchers and practitioners to explore and share their applications of AI in behavioral science.

THE AUTHORS

Christelle Ngnoumen is AVP of Behavioral Science Research at Voya Financial, where she leads research and experimentation efforts focused on helping people boost their savings and benefits utilization and on optimizing the customer experience. As a behavioral scientist and research leader, she leverages behavioral science to design digital experiences and communications that help people make better decisions and improve their wealth and health outcomes. Prior to Voya, she led behavioral science at Headspace, a leading provider of wellbeing and mental health solutions. She has also worked with companies and government agencies on digital experiences and services to improve social outcomes. Christelle holds a Ph.D. in Experimental Psychology from Harvard University.

Kim Nguyen is a customer experience consultant at Voya Financial. Kim leads and supports research, mapping, and testing initiatives. She integrates behavioral science to transform customer insights into actionable recommendations, bridging the gaps between current and future state experiences, and applying academic research to industry practices. She has worked in domains such as journey management, financial literacy, benefit communications, eDelivery adoption, and retirement enrollment. Kim's professional interests center on exploring the barriers people face in making financial decisions, helping them achieve their financial goals and cultivate a healthier relationship with money, thereby increasing their well-being. She holds a Master's degree in Behavioral and Decision Sciences from the University of Pennsylvania and a Bachelor's degree in Psychology from Penn State University.

REFERENCES

- Amnas, M. B., Selvam, M., Raja, M., Santhoshkumar, S., & Parayitam, S. (2023). Understanding the determinants of FinTech adoption: Integrating UTAUT2 with trust theoretic model. *Journal of Risk and Financial Management*, 16(12), 505. <https://doi.org/10.3390/jrfm16120505>
- Barber, B. M., & Odean, T. (2001). Boys will be boys: Gender, overconfidence, and common stock investment. *The Quarterly Journal of Economics*, 116(1), 261–292. <https://doi.org/10.1162/003355301556400>
- Bazley, W. J., Cronqvist, H., & Mormann, M. (2021). Visual finance: The pervasive effects of red on investor behavior. *Management Science*, 67(9), 5616–5641. <https://doi.org/10.1287/mnsc.2020.3747>
- Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, 51(13), 1484–1525. <https://doi.org/10.1016/j.visres.2011.04.012>
- Ceravolo, M. G., Farina, V., Fattobene, L., Leonelli, L., & Raggetti, G. (2022). Anchoring effect in visual information processing during financial decisions: An eye-tracking study. *Journal of Neuroscience, Psychology, and Economics*, 15(1), 19–30. <https://doi.org/10.1037/npe0000153>
- Cerulli Associates. (2024). *U.S. Retirement End-Investor 2023: Personalizing the 401(k) Investor Experience*. (Report No. WA-RD 896.4). Report and recommendations from Cerulli Associates. https://pro.europeana.eu/files/Europeana_Professional/Europeana_Network/metadata-quality-report.pdf
- Collins, J. M. (2012). Financial advice: A substitute for financial literacy? *Financial Services Review*, 21(4), 307–322.
- Corbetta, M., Kincade, J. M., & Shulman, G. L. (2002). Neural systems for visual orienting and their relationships to spatial working memory. *Journal of Cognitive Neuroscience*, 14(3), 508–523. <https://doi.org/10.1162/089892902317362029>
- Davis, M. M., Modi, H. H., Skymba, H. V., Finnegan, M. K., Haigler, K., Telzer, E. H., & Rudolph, K. D. (2022). Thumbs up or thumbs down: Neural processing of social feedback and links to social motivation in adolescent girls. *Social Cognitive and Affective Neuroscience*, 18(1), nsac055. <https://doi.org/10.1093/scan/nsac055>
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18, 193–222. <https://doi.org/10.1146/annurev.ne.18.030195.001205>
- Ely, J. W., Graber, M. L., & Croskerry, P. (2011). Checklists to reduce diagnostic errors. *Academic Medicine*, 86(3), 307–313. <https://doi.org/10.1097/ACM.0b013e31820824cd>

- Employee Benefit Research Institute. (2023, January). *Workplace Retirement Plans: By the Numbers*.
- Employee Benefit Research Institute. (2024, April). *401(k) Plan Asset Allocation, Account Balances, and Loan Activity in 2022*. [https://www.ebri.org/publications/research-publications/issue-briefs/content/401\(k\)-plan-asset-allocation--account-balances--and-loan-activity-in-2022](https://www.ebri.org/publications/research-publications/issue-briefs/content/401(k)-plan-asset-allocation--account-balances--and-loan-activity-in-2022)
- Evans, J. S. B., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science*, 8(3), 223–241.
- Fabrikant, S. I., & Hall, L. E. (2010). The role of visual and verbal information in the comprehension of statistical graphics. *Journal of Experimental Psychology: Applied*, 16(2), 147–1622.
- Giorgadze, N. (2008). The Greater Reality Behind Doors. *Ethnographica*, 19–28. <https://soc.kuleuven.be/antropologie/ethnographica/2008/3Giorgadze.pdf>
- Hegarty, M., Canham, M. S., & Fabrikant, S. I. (2010). Thinking about the weather: How display salience and knowledge affect performance in a graphic inference task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36(1), 37–53. <https://psycnet.apa.org/doi/10.1037/a0017683>
- Homer, P. M., & Gauntt, S. G. (1992). The role of imagery in the processing of visual and verbal package information. *Journal Of Mental Imagery-New York-International Imagery Association*, 16, 123–123.
- Hung, A., & Yoong, J. (2013). Asking for help: Survey and experimental evidence on financial advice and behavior change. *The Market for Retirement Financial Advice*, 182, 182–212.
- Investment Company Institute. (2024). *401(k) Resource Center*. <https://www.ici.org/401k>
- Ismail, A., & Mavis, C. P. (2022). A new method for measuring CEO overconfidence: Evidence from acquisitions. *International Review of Financial Analysis*, 79. <https://doi.org/10.1016/j.irfa.2021.101964>
- James, W. (1890). *The principles of psychology*. Henry Holt and Co.
- Kahneman, D. (2011). *Thinking, fast and slow (Vol. 1)*. Farrar, Straus and Giroux.
- Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 49–81). Cambridge University Press.
- Kastner, S., & Ungerleider, L. G. (2000). Mechanisms of visual attention in the human cortex. *Annual Review of Neuroscience*, 23, 315–341. <https://doi.org/10.1146/annurev.neuro.23.1.315>
- Luffarelli, J., Mukesh, M., & Mahmood, A. (2019). Let the logo do the talking: The influence of logo descriptiveness on brand equity. *Journal of Marketing Research*, 56(5), 862–878. <https://doi.org/10.1177/0022243719845000>
- Lusardi, A., & Messy, F.-A. (2023). The importance of financial literacy and its impact on financial wellbeing. *Journal of Financial Literacy and Wellbeing*, 1(1), 1–11. <https://doi.org/10.1017/flw.2023.8>
- Lusardi, A., & Streeter, J. L. (2023). Financial literacy and financial well-being: Evidence from the US. *Journal of Financial Literacy and Wellbeing*, 1(2), 169–198. <https://doi.org/10.1017/flw.2023.13>
- Malmendier, U., & Tate, G. (2005). Does overconfidence affect corporate investment? CEO overconfidence measures revisited. *European Financial Management*, 11(5), 649–659. <https://doi.org/10.1111/j.1354-7798.2005.00302.x>
- Marsden, M., Zick, C. D., & Mayer, R. N. (2011). The value of seeking financial advice. *Journal of Family and Economic Issues*, 32, 625–643. <https://doi.org/10.1007/s10834-011-9258-z>
- Mercado, C., Bullard, K. M., Bolduc, M. L. F., Banks, D., Andrews, C., Freggens, Z. R. F., & Njai, R. (2024). Exploring associations of financial well-being with health behaviours and physical and mental health: A cross-sectional study among US adults. *BMJ Public Health*, 2, e000720.
- Mihaylov, G., Yawson, A., & Zurbrugg, R. (2015). The decision to seek advice in the self-directed retirement fund industry. *Applied Economics*, 47(32), 3367–3381. <https://doi.org/10.1080/00036846.2015.1013620>

- Orquin, J. L., Lahm, E. S., & Stojić, H. (2021). The visual environment and attention in decision making. *Psychological Bulletin*, 147(6), 597–617. <https://doi.org/10.1037/bul0000328>
- Padilla, L. M., Creem-Regehr, S. H., Hegarty, M., & Stefanucci, J. K. (2018). Decision making with visualizations: A cognitive framework across disciplines. *Cognitive Research: Principles and Implications*, 3, 1–25. <https://doi.org/10.1186/s41235-018-0120-9>
- Ronen, J., Ronen, T., Zhou, M. J., & Gans, S. E. (2023). The informational role of imagery in financial decision making: A new approach. *Journal of Behavioral and Experimental Finance*, 40, 100851. <https://doi.org/10.1016/j.jbef.2023.100851>
- Schirillo, J. A., & Stone, J. V. (2005). The influence of spatial frequency on the perception of facial expressions. *Journal of Vision*, 5(8), 671–6822.
- Stone, J. V., Yates, J. F., & Parker, A. M. (2003). The role of visual and verbal information in the comprehension of statistical graphics. *Journal of Experimental Psychology: Applied*, 9(2), 147–1622.
- Statista. (2023). *Fintech adoption rate*. <https://www.statista.com/statistics/1331176/fintech-adoption-rate/>
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12(1), 97–136.
- Vessey, I., & Galletta, D. (1991). Cognitive fit: An empirical study of information acquisition. *Information Systems Research*, 2(1), 63–84.
- Westermann, S., Niblock, S. J., Harrison, J. L., & Kortt, M. A. (2020). Financial advice seeking: A review of the barriers and benefits. *Economic Papers: A Journal of Applied Economics and Policy*, 39(4), 367–388. <https://doi.org/10.1111/1759-3441.12294>
- Yalcin, G., Lim, S., Puntoni, S., & van Osselaer, S. M. J. (2022). Thumbs up or down: Consumer reactions to decisions by algorithms versus humans. *Journal of Marketing Research*, 59(4), 696–717. <https://doi.org/10.1177/00222437211070016>
- Zhu, M., Yang, Y., & Hsee, C. K. (2018). The mere urgency effect. *Journal of Consumer Research*, 45(3), 673–690. <https://doi.org/10.1093/jcr/ucy008>

One Nudge Is Not Enough: Exploring the Limits and Potentials of Behavioral Approaches for Sustainable Tourism Mobility

JÖRN BASEL¹, CHRISTIAN WEIBEL AND MALIN WENZKE

Lucerne School of Business

This article provides a literature-based overview of the effectiveness of behavioral interventions—so-called “nudges”—in promoting sustainable tourism mobility. Drawing on the model of self-regulated behavioral change and findings from behavioral economics, the article discusses how the effectiveness of nudges in this domain is influenced by a wide variety of factors and whether these kinds of subtle interventions can lead to lasting behavioral change. Rather than presenting new empirical data, the article synthesizes existing research and illustrates key findings through selected case examples. Special attention is given to whether one-time nudges, such as free public transport offers, are sufficient to induce long-term shifts in travel behavior among tourists. The overview shows that while nudges can potentially trigger initial changes, sustained behavioral transformation typically requires repeated or combined interventions tailored to a specific travel context.

Introduction

The classic study by Goldstein and colleagues (2008) “the majority of guests reuse their towels”, which aimed to encourage hotel guests to reuse towels and thus promote sustainable tourism behavior, showed that nudging in the form of a note stating that 75% of all guests in that room reuse their towels led to almost half of the guests actually reusing their towels. A change in behavior simply by means of a note! But is it really that simple?

In recent years, behavior-oriented approaches have increasingly become prominent (White et al., 2019). The focus is on systematically (re)designing the choice environment to reduce travelers’ ecological footprints. The need for environmentally friendly behavior is widely recognized, which is why comprehensive approaches have been introduced by governments, private sector actors, and non-governmental organizations.

(Green) Nudges (Zaneva & Dumbalska, 2020) have a clear advantage because they do not involve regulations (like bans) or pure monetary incentives. These two features make them particularly appealing to the tourism industry, where intense competition

means there is always a fear that certain sustainability measures could deter customers. Nudging therefore acts as a silver bullet, in that it encourages travelers to adopt more sustainable behavior at little cost, which ideally corresponds to their preferences anyway—a win-win situation. However, a closer look at the empirical findings on pure nudging approaches show that existing evidence is fragmented and, in some cases, even contradictory at first glance (Song et al., 2024; Souza-Neto et al., 2023).

This article therefore aims to present a classification of the existing findings and critically examine the effectiveness of nudges for behavioral change in tourism mobility, which is the central source of emissions within the tourism industry, based on selected evidence. It does so by addressing the following questions:

1. How effective are nudging approaches in promoting sustainable tourism mobility?
2. What factors determine the success or failure of nudging measures?
3. What complementary strategies are necessary to bring about long-term behavioral change?

¹ Corresponding author: joern.basel@hslu.ch

In order to assess nudging interventions thoroughly in the context of sustainable tourism mobility, a theoretical framework is needed—one that explains not only behavior itself, but also how it changes over time. The Phase Model of Self-Regulated Behavior Change offers such a framework. The following chapter introduces this model and illustrates how nudges can be systematically assigned to different phases of the behavior change process, allowing for a theory-driven synthesis of empirical findings on nudging in sustainable tourism mobility. Drawing on the principles of narrative and integrative reviews (Baumeister & Leary, 1997; Snyder, 2019; Tranfield et al., 2003), the article reviews empirical evidence on the effectiveness of nudges in the fields of tourism and mobility and, additionally, discusses key challenges through a case study from Engelberg, Switzerland. It also discusses the limitations of nudging and outlines complementary strategies to foster long-term sustainable tourism mobility behavior.

Behavior Change in Phases: Levers for Sustainable Tourism Mobility

The Phase Model of Self-Regulated Behavioral Change (see Fig. 1), with its four phases, offers a

structured approach to understanding and promoting shifts toward sustainable tourism mobility (Bamberg, 2013a, 2013b; Ohnmacht et al., 2017).

Building on this model, we now explore concrete levers that can be applied within each phase to guide individuals. Further, by aligning behavioral interventions with the psychological tasks of each phase, the potential impact of nudges and other measures can be maximized.

Pre-Decision: Building Awareness

One approach to building a goal intention is to establish norms and raise awareness of sustainable tourism mobility. Social norms are shared values conveyed by society that influence behaviors by indicating what is considered right or wrong (Ajzen, 1991; Bicchieri & Dimant, 2022). Personal norms and internalized social norms are also linked to pro-environmental behavior (Bamberg et al., 2007; Bamberg & Möser, 2007; Stern, 2000), in that the stronger the personal norm, the greater the likelihood that the intention to use public transportation to reach a destination, for example, will arise. The social norm should be tailored to the target group, with a possible descriptive message being: ‘X percent of




Figure 1: Phase Model of Self-Regulated Behavioral Change.

people in your city used public transportation for their weekend trip' (Gravert & Olsson Collentine, 2021).

Targeted measures to address norms could include:

- Social events (e.g., sustainable travel challenges at destinations)

- Commitment nudges (e.g., public pledges to use sustainable mobility)
- Advertising focused on moral values (e.g., environmental protection).

But what does the evidence tell us? Is one nudge enough (see Table 1)?

Table 1: Evidence of (In)Effectiveness in the Pre-Decision Phase

Evidence of Effectiveness	Evidence of Ineffectiveness
Steffen et al. (2024): A mobility study in Flanders (Belgium) demonstrated the effectiveness of social norm nudges, revealing the percentage of public transport and bicycle use by neighbors and disclosing the resulting emissions savings. The nudges led to an increase in the intention to use both public transportation and active travel modes.	Gravert & Olsson Collentine (2021): A mobility study in Sweden examined a social norm nudge stating that '72% of residents use public transport'. However, this nudge did not influence the choice of transport mode.
Result: Norm nudges may or may not work with certainty; context, wording, and presentation can influence their effectiveness.	

Pre-Action: Forming an Intention

Based on the Theory of Planned Behavior (Ajzen, 1991), measures that target attitudes and perceived behavioral control are promising for encouraging behavioral intentions.

"Attitude" measures the evaluation of assessment objects (e.g., people, objects, ideas), for example the question pertaining to whether public transportation is viewed positively (e.g., 'public transportation makes my trip stress-free') or negatively (e.g., 'traveling with luggage on public transportation is exhausting'). Knowledge plays an important role in the formation of beliefs (Bamberg & Möser, 2007). Furthermore, attitudes are influenced by communication campaigns, especially by emotional communication that evokes positive (e.g., happiness, wellbeing, trust) or negative feelings (e.g., fear, guilt, worry) (Bagozzi et al., 1999; Ohnmacht et al., 2017).

Measures to influence attitudes towards sustainable tourism mobility could include:

- Information on the positive impacts of behavior change (e.g., smartphone feedback on cost savings after a trip using public transport)
- Emotional communication that conveys feelings

- Image campaigns to enhance the public transport image or to promote the sustainability efforts of destinations.

But what does the evidence tell us? Is one nudge enough (see Table 2)?

"Perceived behavioral control" refers to a person's subjective assessment of their ability to perform a certain behavior, including the necessary skills, abilities, and resources. A contradiction between perceived control and attitude can arise, for example, when someone has a positive attitude toward public transportation but is unable to implement the desired behavior (e.g., due to excessive distance to the nearest stop) (Ajzen, 1991).

Measures to influence the perceived behavioral control could include:

- Economic incentives (e.g., free transport tickets)
- Service measures (e.g., platforms combining several sustainable means of transport in tourism destinations)
- Convenience nudges (e.g., improved placement of public transport information in apps)

But what does the evidence tell us? Is one nudge enough (see Table 3)?

Table 2: Evidence of (In)Effectiveness in the Pre-Action Phase (Attitude)

Evidence of Effectiveness	Evidence of Ineffectiveness
Aravind et al. (2024): A mobility study in Tennessee tested various types of nudges, including emotional nudges, to promote the use of public transportation. They used smileys and color schemes to convey information in an emotional way, resulting in an increase in the likelihood of choosing public transportation.	Dolnicar & Demeter (2024): A literature review on the effectiveness of nudges in tourism to influence attitude change concluded that nudges often fail in this regard. Possible reasons included tourists often not noticing the nudges, not processing the messages cognitively, or not changing their attitudes based on the message.
	Luger-Bazinger et al. (2023): A former Behavioral Economics Guide article evaluated various techniques for changing behavior that are used in mobility apps, including the dissemination of information. The authors concluded that more than half of the apps examined used measures to influence attitudes, but according to Semenscus et al. (2020), they could not be classified as effective techniques.

Result: Cognitive nudges alone, such as providing information, appear insufficient for changing attitudes.

Table 3: Evidence of (In)Effectiveness in the Pre-Action Phase (Perceived Behavioral Control)

Evidence of Effectiveness	Evidence of Ineffectiveness
Riggs (2017): A mobility study used nudges to test the elimination of parking spaces and the reduction of car trips at Cal Poly University. Both financial incentives (a gift or a sum of money) and social norms in the form of altruistically giving up a parking space for someone else were tested. The study identified both financial and social nudges as effective in this regard.	Sticher et al. (2024): This study tested financial incentives, specifically the raffling of mobility vouchers in a competition at a Swiss university to promote sustainable mobility alternatives among students and employees. This study did not identify financial incentives as being effective.
	Riggs (2017): Although the study found that financial and social incentives are individually effective, combining them was ineffective. This shows that more incentives do not necessarily lead to greater effectiveness.

Result: These findings demonstrate that the combination of nudges does not necessarily increase effectiveness and that combinations of interventions need to be tested and well-coordinated.

Action: Implementing Behavior

The transition from the pre-action phase to the action phase is marked by the formation of an implementation intention, which can be understood as a concrete ‘if-then plan’ (Gollwitzer, 1999). This approach specifies when, where, and how a behavior will be carried out; for example, ‘On my next trip to the Swiss mountains, I

will use public transportation instead of using my car’.

However, difficulties in action planning (e.g., motivation or implementation problems) may arise and need to be addressed. Strategies in this case could include imparting procedural knowledge, setting goals, or guiding consumers through choice architecture and default options.

Possible measures include:

- Information brochures (e.g., on public transport use at tourism destinations)
- Default setting (e.g., public transport as the main mean of transportation in routing

applications)

- Setting reference points by the community (e.g., reducing carbon emissions by X%).

But what does the evidence tell us? Is one nudge enough (see Table 4)?

Table 4: Evidence of (In)Effectiveness in the Action Phase

Evidence of Effectiveness	Evidence of Ineffectiveness
Anagnostopoulou et al. (2018): A literature review examined 44 articles testing persuasive mobility applications to promote sustainable mobility behavior. Under a third of these applications used goal-setting as a behavioral change technique, which was identified as effective.	Arnott et al. (2014): A meta-analysis examined 13 studies investigating the effectiveness of behavioral interventions to increase sustainable mobility alternatives. Most of the studies also included measures for action planning or the transfer of procedural knowledge. The meta-analysis found no significant effectiveness of these interventions.
Result: The effectiveness of nudges depends on various components, e.g., a technically stimulating, user-friendly, and appropriately combined implementation of interventions can boost effectiveness.	

Case Study: Engelberg, Switzerland

For environmentally friendly campaigns, it is crucial not to “preach to the converted,” as this can have an illusory effect, since it does not lead to any real change in attitudes or behavior. It therefore seems sensible to start where there is the greatest potential for behavioral change.

In their study, Kormos et al. (2020) selectively integrated behavior-guiding design elements, such as references to social norms or a reduction of ambiguity, into postcards with the aim of improving land use between landowners and tenants in the Canadian

province of British Columbia. In a similar way, this study aimed to use nudges to encourage people who had traveled to the destination by car to use public transport on their next visit to Engelberg (Basel et al., 2023). As an independent variable, five different postcard designs were created, and in each of which one (or more) nudges were integrated:

1. Placebo postcard (picture of a train journey, stating, ‘It’s worth traveling by train’, without additional information) (see Fig. 2a).
2. Information about carbon savings by traveling by public transport: ‘It’s worth traveling by train.



Figure 2: Postcard front design examples.

*For the environment!*⁹ (See Fig 2b).

3. Information about saving on parking fees, to activate loss aversion: *'It's worth traveling by train. For the environment and for you. It's better to invest the parking fee in coffee and croissants'*.
4. Statement from an outdoor influencer, communicating a social norm: *'We have a responsibility to ensure that our travel behavior is environmentally friendly'*.
5. All the above information together (carbon saving, saving on the parking fee, and reference to the social norm).

With the support of the NGO *Protect Our Winters Switzerland* (POW), 4,000 cards in total were randomly distributed to motorists departing the Swiss tourist destination Engelberg at the start of the 2022 winter season (see Fig. 3) (of which around 11% were solo travelers, 66% with family members, 20% with friends, and the remaining 3% with carpools). In addition to the corresponding nudge, the car drivers were offered the equivalent of a small snack (coffee and croissant) on their return journey by public transport as a small additional incentive highlighted on the backside of the postcard.

However, the response rate was marginal, and only around two handfuls of cards were redeemed. The hypothesis that cards with a nudge message should prove to be effective—especially in comparison with the card without an explicit message—could not be confirmed. The nudging used was not able to initiate a direct switch from the car to public transport. Using a QR code or pre-stamped on the backside of the

postcard, however, participants could also take part in a short survey and win an annual pass to Engelberg. A total of 282 postcards were returned by post, and 177 online questionnaires were completed in full. The results of this survey showed why a nudge was apparently not enough in this context. Many tourists continue to prefer cars because of comfort, flexibility and luggage transportation. These factors are the strongest predictors of sustainable transportation choices, so targeted measures such as improved public transport connections, luggage transport services, and attractive alternatives are needed. Infrastructure improvements and combined measures of incentives, information, and mobility solutions are more promising than pure nudging campaigns.

Post-Action: Forming a Habit

The goal of the habit phase is to maintain and internalize the new behavior. To establish a new routine, it is crucial to avoid setbacks (e.g., returning to driving in the Swiss mountains). Despite firm intentions to implement change, long-term behavioral change can be disrupted by distractions, old habits, or conflicting goals (Gollwitzer, 1999). Feedback mechanisms can be a possible way to prevent falling back into old patterns (Abrahamse et al., 2005; McCalley & Midden, 2002).

Examples of feedback mechanisms include:

- Gamification (e.g., app-based feedback on carbon savings)
- Thank-you SMS (e.g., from transport companies after using public transport)



Figure 3: Distribution procedure at the three parking lots in the destination Engelberg, Switzerland.

- Reminders (e.g., from transport companies or local authorities to use public transport).

But what does the evidence tell us? Is one nudge enough (see Table 5)?

Beyond Nudging: Why One Nudge Is Not Enough

The concept of nudging sounds promising, with some studies supporting the notion that it

Table 5: Evidence of (In)Effectiveness in the Post-Action Phase

Evidence of Effectiveness	Evidence of Ineffectiveness
Semenscu et al. (2020): A meta-analysis evaluated nudges aimed at reducing car traffic, finding that providing feedback on individual mobility behavior is the most effective way to influence the choice of transport mode.	Greene et al. (2024): A meta-analysis in the tourism sector evaluated measures aimed at changing behavior, with the goal of promoting sustainable travel behavior. Various nudging measures were examined, including penalties as a form of negative feedback. Compared to the other measures tested, penalties proved to be only partially effective.
Result: Feedback can be an effective nudging technique, but it should be constructive, because punishment as a form of negative feedback does not work to promote sustainable travel behavior.	

could promote sustainable tourism mobility in certain cases. However, this article also shows that nudging has its limitations (see Table 6), thereby raising the question as to how these boundaries can be overcome.

Conclusion

This article critically examined the potentials and limitations of nudging to steer travelers towards more sustainable tourism mobility and thus complements existing reviews that deal with the more general effectiveness of nudging (Hummel & Maedche, 2019; Mertens et al., 2022). Nudging can provide impulses for sustainable tourism mobility if certain rules for the effective design and distribution of nudges are observed (Nowak, 2025). At the same time, if certain rules are not followed, the effectiveness of nudges is limited.

Isolated nudges alone do not appear to be sufficient to achieve long-term behavioral changes but can rather be used as short-term controlling or supporting measures (Ixmeier & Kranz, 2024). Thus, the effectiveness of nudges can be particularly developed by using them as a supplement to political, regulatory, or economic measures. Nevertheless, clear recommendations for action can be derived from these heterogenic findings:

- For policymakers, the relevance of developing

hybrid approaches that combine nudging with regulatory, infrastructural, and economic measures is becoming clear.

- For companies and destinations, the development of and investment in behavior-based nudges, such as defaults and choice architecture, is particularly promising for influencing the behavior of their guests and customers.
- For research, more studies are needed on the combination of nudges or with structural measures, as well as on the long-term effects of nudging. In addition, of course, there is also a need for publication options where failed nudging interventions can be reported (Beermann et al., 2024; Maier et al., 2022).

Looking at studies that have been able to document effective nudging, one can also see the great potential that can be assumed in a stronger technological focus. Those “tech-nudges” offer the potential to follow the rules identified in this article; for example, they can be used in a context-sensitive and target group-specific way, various nudges can be easily and cost-effectively bundled, and cognitive nudges as well as default nudges can be easily combined. For example, apps for route planning can highlight sustainable travel options by default, thus avoiding choice overload and facilitating decision-making by providing a suitable selection of travel options. In this

Table 6: Limitations and Success Potentials for Nudging

Limitations of Nudging	Success Potentials to Overcome Those Limitations
Nudges that work in one context do not necessarily have the same effect in other contexts (Gravert & Olsson Collentine, 2021; Steffen et al., 2024).	Nudges must be designed and tested in context, e.g. explicitly tailored to the circumstances of a destination.
Nudges do not work equally well for all recipients, nor according to a scattergun approach. Not all target groups have the same needs, the same potential for change, or can be reached through the same channels (Anagnostopoulou et al., 2018; Arnott et al., 2014; Riggs, 2017; Sticher et al., 2024).	Nudges must be targeted at groups who exhibit a deficit in relation to the desired behavior, who can be encouraged to change their behavior by the nudges, and who can change their behavior. Therefore, involving target groups in the development of appropriate measures can be a useful approach. When implementing nudges, the channels through which the target groups can be reached must also be considered. In particular, the targeted integration of technologies such as AI can enable a target group-specific and individualized approach (Luger-Bazinger et al., 2023).
Decision information, e.g., cognitive nudges, alone can be effective but often prove insufficient (Dolnicar & Demeter, 2024).	The combination of different intervention categories can be more effective than a single category. These categories are decision information (increase availability and relevance of information, e.g. cognitive nudges), decision structure (alter utility of choice option through their arrangement or format, e.g., defaults, framing), and decision assistance (facilitating self-regulation, e.g., feedback) (Münscher et al., 2016). For example, cognitive nudges could be strengthened by social reinforcement or, for example, communities, whereby the purely cognitively processed nudges are supplemented by a real experience (e.g., sustainable travel challenges at destinations).
Evidence shows that ‘more is not always more’, i.e., nudges can also inhibit each other under certain circumstances (Riggs, 2017).	Nudges must be well coordinated, and the effectiveness of these combinations must be evaluated at best to exclude mutual inhibition and identify the most effective combinations possible.
Despite their potential effectiveness, nudges often do not address structural changes or habitualization (Yachin et al., 2024). Without appropriate complementary measures, behavior can easily fall back into old patterns.	Nudging measures should be complemented by appropriate infrastructural, political, regulatory, and economic measures. In addition, nudges should be accompanied by approaches that foster intrinsic motivation and personal identification with sustainable behavior, for instance through value-based communication, social engagement, or gamification. Research emphasizes that lasting behavioral change requires not only external prompts, but also the support of individuals’ autonomy, competence, and relatedness to foster intrinsic motivation (Deci & Ryan, 2000). Thus, interventions that build individuals’ skills and competencies, so-called “boosts,” are essential for enabling sustainable and self-determined behavior (Grüne-Yanoff & Hertwig, 2016).

context, the potential of the targeted integration of artificial intelligence, which, for example, helps to display personalized route suggestions or generate automated push notifications about carbon savings, should not be underestimated. Another promising approach in behavioral change can also be found in educational strategies. Herzog and Hertwig (2017, 2025) public officials have shown a growing interest in using evidence from the behavioral sciences to promote policy goals. Much of the discussion of behaviorally informed approaches has focused on ‘nudges’; that is, non-fiscal and non-regulatory interventions that steer (nudge suggest a technique called “boosting,” which aims to enhance individuals’ knowledge and competences, enabling them to make informed decisions. By fostering critical thinking and decision-making abilities, this tactic promotes long-term behavioral change that is both autonomous and sustainable. If a nudge is not enough, an additional educational approach such as boosting can thus potentially help make an impact across the spectrum.

THE AUTHORS

Jörn Basel holds a PhD from ESCP Business School Berlin and is Professor of Business Psychology at Lucerne School of Business. His research has mostly an applied focus and concentrates on behavioral inter-ventions in various areas with a particular emphasis on trust dynamics. In addition to his academic work, he serves as an advisor for The Negotiation Studio, a consulting firm specializing in negotiation training and coaching.

Christian Weibel earned a PhD in Psychology from the University of Bern and serves as Professor of Business Psychology at Lucerne School of Business. His research explores how people make consumer decisions, with a particular emphasis on promoting sustainable behavior and understanding the behavioral mechanisms behind pricing. Complementing his academic work, he brings several years of professional experience in both the industry and the financial sector.

Malin Wenzke is a research associate for behavioral economics at Lucerne School of Business and holds an MSc in Business Psychology from the Stuttgart Business School. Her research focuses on behavioral interventions to strengthen sustainable tourism

behavior. She previously worked as a consultant at FehrAdvice & Partners, where she gained extensive practical experience in the application of behavioral economics.

REFERENCES

- Abrahamse, W., Steg, L., Vlek, C., & Rothengatter, T. (2005). A review of intervention studies aimed at household energy conservation. *Journal of Environmental Psychology*, 25(3), 273–291. <https://doi.org/10.1016/j.jenvp.2005.08.002>
- Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50(2), 179–211. [https://doi.org/10.1016/0749-5978\(91\)90020-T](https://doi.org/10.1016/0749-5978(91)90020-T)
- Anagnostopoulou, E., Bothos, E., Magoutas, B., Schrammel, J., & Mentzas, G. (2018). Persuasive technologies for sustainable mobility: State of the art and emerging trends. *Sustainability*, 10(7), 2128. <https://doi.org/10.3390/su10072128>
- Aravind, A., Mishra, S., & Meservy, M. (2024). Nudging towards sustainable urban mobility: Exploring behavioral interventions for promoting public transit. *Transportation Research Part D: Transport and Environment*, 129, 104130. <https://doi.org/10.1016/j.trd.2024.104130>
- Arnott, B., Rehackova, L., Errington, L., Sniehotta, F. F., Roberts, J., & Araujo-Soares, V. (2014). Efficacy of behavioural interventions for transport behaviour change: Systematic review, meta-analysis and intervention coding. *International Journal of Behavioral Nutrition and Physical Activity*, 11(1). <https://doi.org/10.1186/s12966-014-0133-9>
- Bagozzi, R. P., Gopinath, M., & Nyer, P. U. (1999). The Role of Emotions in Marketing. *Journal of the Academy of Marketing Science*, 27(2), 184–206. <https://doi.org/10.1177/0092070399272005>
- Bamberg, S. (2013a). Applying the stage model of self-regulated behavioral change in a car use reduction intervention. *Journal of Environmental Psychology*, 33, 68–75. <https://doi.org/10.1016/j.jenvp.2012.10.001>
- Bamberg, S. (2013b). Changing environmentally harmful behaviors: A stage model of self-regulated behavioral change. *Journal of*

- Environmental Psychology*, 34, 151–159. <https://doi.org/10.1016/j.jenvp.2013.01.002>
- Bamberg, S., Hunecke, M., & Blöbaum, A. (2007). Social context, personal norms and the use of public transportation: Two field studies. *Journal of Environmental Psychology*, 27(3), 190–203. <https://doi.org/10.1016/j.jenvp.2007.04.001>
- Bamberg, S., & Möser, G. (2007). Twenty years after Hines, Hungerford, and Tomera: A new meta-analysis of psycho-social determinants of pro-environmental behaviour. *Journal of Environmental Psychology*, 27(1), 14–25. <https://doi.org/10.1016/j.jenvp.2006.12.002>
- Basel, J., Ohnmacht, T., Zimmermann, L., Schöb, D., Wyss, L., & Ries, T. (2023). *More than just a nudge: Kundenzentrierte Förderung nachhaltigen Mobilitätsverhaltens im Outdoor-Tourismus*. Lucerne University of Applied Sciences and Arts. <https://www.hslu.ch/de-ch/hochschule-luzern/forschung/projekte/detail/?pid=5905>
- Baumeister, R. F., & Leary, M. R. (1997). Writing narrative literature reviews. *Review of General Psychology*, 1(3), 311–320. <https://doi.org/10.1037/1089-2680.1.3.311>
- Beermann, V., Enkmann, J. M., Maier, M., & Bartoš, F. (2024). *How effective are digital green nudges? A publication bias-adjusted meta-analysis*. Proceedings of the International Conference on Information Systems (ICIS). https://aisel.aisnet.org/icis2024/lit_review/lit_review/5/
- Bicchieri, C., & Dimant, E. (2022). Nudging with care: The risks and benefits of social information. *Public Choice*, 191(3–4), 443–464. <https://doi.org/10.1007/s1127-019-00684-6>
- Deci, E. L., & Ryan, R. M. (2000). The “what” and “why” of goal pursuits: Human needs and the self-determination of behavior. *Psychological Inquiry*, 11(4), 227–268. https://doi.org/10.1207/S15327965PLI1104_01
- Dolnicar, S., & Demeter, C. (2024). Why targeting attitudes often fails to elicit sustainable tourist behaviour. *International Journal of Contemporary Hospitality Management*, 36(3), 730–742. <https://doi.org/10.1108/IJCHM-07-2022-0828>
- Goldstein, N. J., Cialdini, R. B., & Griskevicius, V. (2008). A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *Journal of Consumer Research*, 35(3), 472–482. <https://doi.org/10.1086/586910>
- Gollwitzer, P. M. (1999). Implementation intentions: Strong effects of simple plans. *American Psychologist*, 54(7), 493–503. <https://doi.org/10.1037/0003-066X.54.7.493>
- Gravert, C., & Olsson Collentine, L. (2021). When nudges aren’t enough: Norms, incentives and habit formation in public transport usage. *Journal of Economic Behavior & Organization*, 190, 1–14. <https://doi.org/10.1016/j.jebo.2021.07.012>
- Greene, D., Demeter, C., & Dolnicar, S. (2024). The comparative effectiveness of interventions aimed at making tourists behave in more environmentally sustainable ways: A meta-analysis. *Journal of Travel Research*, 63(5), 1239–1255. <https://doi.org/10.1177/00472875231183701>
- Grüne-Yanoff, T., & Hertwig, R. (2016). Nudge versus boost: How coherent are policy and theory? *Minds and Machines*, 26(1–2), 149–183. <https://doi.org/10.1007/s11023-015-9367-9>
- Hertwig, R. (2017). When to consider boosting: Some rules for policy-makers. *Behavioural Public Policy*, 1(2), 143–161. <https://doi.org/10.1017/bpp.2016.14>
- Herzog, S. M., & Hertwig, R. (2025). Boosting: Empowering citizens with behavioral science. *Annual Review of Psychology*, 76(1), 851–881. <https://doi.org/10.1146/annurev-psych-020924-124753>
- Hummel, D., & Maedche, A. (2019). How effective is nudging? A quantitative review on the effect sizes and limits of empirical nudging studies. *Journal of Behavioral and Experimental Economics*, 80, 47–58. <https://doi.org/10.1016/j.socec.2019.03.005>
- Ixmeyer, A., & Kranz, J. (2024). *The effectiveness of digital interventions to promote pro-environmental behaviour: A meta-analysis*. Proceedings of the International Conference on Information Systems (ICIS). https://aisel.aisnet.org/ecis2024/track17_greenis/track17_greenis/34/
- Luger-Bazinger, C., Geser, G., & Hornung-Prähauser, V. (2023). Digital behavioural interventions for sustainable mobility: A review of behaviour change techniques in mobile apps. *Behavioral Economics Guide 2023*, 68–75
- Maier, M., Bartoš, F., Stanley, T. D., Shanks, D. R., Harris, A. J. L., & Wagenmakers, E.-J. (2022). No evidence for nudging after adjusting for publication bias. *Proceedings of the National Academy*

- of Sciences, 119(31). <https://doi.org/10.1073/pnas.2200300119>
- McCalley, L. T., & Midden, C. J. H. (2002). Energy conservation through product-integrated feedback: The roles of goal-setting and social orientation. *Journal of Economic Psychology*, 23(5), 589–603. [https://doi.org/10.1016/S0167-4870\(02\)00119-8](https://doi.org/10.1016/S0167-4870(02)00119-8)
- Mertens, S., Herberz, M., Hahnel, U. J. J., & Brosch, T. (2022). The effectiveness of nudging: A meta-analysis of choice architecture interventions across behavioral domains. *Proceedings of the National Academy of Sciences*, 119(1). <https://doi.org/10.1073/pnas.2107346118>
- Münscher, R., Vetter, M., & Scheuerle, T. (2016). A review and taxonomy of choice architecture techniques. *Journal of Behavioral Decision Making*, 29(5), 511–524. <https://doi.org/10.1002/bdm.1897>
- Nowak, M. (2025). *Pro-sustainable consumer behaviour in tourism and hospitality: Drivers, barriers, and effective behavioural intervention design* [Dissertation]. Mid Sweden University.
- Ohnmacht, T., Schaffner, D., Weibel, C., & Schad, H. (2017). Rethinking social psychology and intervention design: A model of energy savings and human behavior. *Energy Research & Social Science*, 26, 40–53. <https://doi.org/10.1016/j.erss.2017.01.017>
- Riggs, W. (2017). Painting the fence: Social norms as economic incentives to non-automotive travel behavior. *Travel Behaviour and Society*, 7, 26–33. <https://doi.org/10.1016/j.tbs.2016.11.004>
- Semenescu, A., Gavreliuc, A., & Sârbescu, P. (2020). 30 Years of soft interventions to reduce car use: A systematic review and meta-analysis. *Transportation Research Part D: Transport and Environment*, 85, 102397. <https://doi.org/10.1016/j.trd.2020.102397>
- Snyder, H. (2019). Literature review as a research methodology: An overview and guidelines. *Journal of Business Research*, 104, 333–339. <https://doi.org/10.1016/j.jbusres.2019.07.039>
- Song, H., Wu, H., & Zhang, H. (2024). Can nudging affect tourists' low-carbon footprint travel choices? *International Journal of Contemporary Hospitality Management*, 36(5), 1534–1556. <https://doi.org/10.1108/IJCHM-09-2022-1175>
- Souza-Neto, V., Marques, O., Mayer, V. F., & Lohmann, G. (2023). Lowering the harm of tourist activities: A systematic literature review on nudges. *Journal of Sustainable Tourism*, 31(9), 2173–2194. <https://doi.org/10.1080/09669582.2022.2036170>
- Steffen, J., Hook, H., & Witlox, F. (2024). Improving interest in public, active, and shared travel modes through nudging interventions. *Transportation Research Part F: Traffic Psychology and Behaviour*, 103, 353–367. <https://doi.org/10.1016/j.trf.2024.04.020>
- Stern, P. C. (2000). New environmental theories: Toward a coherent theory of environmentally significant behavior. *Journal of Social Issues*, 56(3), 407–424. <https://doi.org/10.1111/0022-4537.00175>
- Sticher, S., Wallimann, H., & Balthasar, N. (2024). How (not) to incentivize sustainable mobility? Lessons from a Swiss mobility competition (Version 1). arXiv. <https://doi.org/10.48550/ARXIV.2409.11142>
- Tranfield, D., Denyer, D., & Smart, P. (2003). Towards a methodology for developing evidence-informed management knowledge by means of systematic review. *British Journal of Management*, 14(3), 207–222. <https://doi.org/10.1111/1467-8551.00375>
- White, K., Habib, R., & Hardisty, D. J. (2019). How to SHIFT consumer behaviors to be more sustainable: A literature review and guiding framework. *Journal of Marketing*, 83(3), 22–49. <https://doi.org/10.1177/0022242919825649>
- Yachin, J. M., Margaryan, L., Lexhagen, M., & Ioannides, D. (2024). Nudge plus in tourism: Reflexive behaviours and reflective attitudes. *Journal of Sustainable Tourism*, 1–18. <https://doi.org/10.1080/09669582.2024.2436907>
- Zaneva, M., & Dumbalska, T. (2020). Green nudges: Applying behavioural economics to the fight against climate change. *PsyPag Quarterly*, 1(116), 27–31. <https://doi.org/10.53841/bp-spag.2020.1.116.27>

The Behaviorally-Informed Insurance Company: Insights From the Allianz Global Center for Behavioral Economics

SORAYA CATOJA AND ANN-KATHRIN CREDE¹

Allianz, Global Center for Behavioral Economics

Applying behavioral economics (BE) in business has long been recognized as a powerful lever for optimizing decision-making, enhancing customer engagement, and improving operational efficiency. This article shares practical insights from the Allianz Global Center for Behavioral Economics on how BE insights can be effectively embedded within a large insurance organization and integrated into day-to-day operations. Drawing on a recent project aimed at increasing digital adoption in roadside assistance, we highlight the value of evidence-based designs and the impact of behaviorally-informed solutions. Based on pilots conducted across 13 countries, our findings show that presenting the digital channel as the default option, while clearly communicating its benefits and ease of use, leads to a significant increase in customers adopting the digital path rather than opting for a phone call.

Allianz Global Center for Behavioral Economics

Our History

Allianz, a global insurance provider, launched the Allianz Global Center for Behavioral Economics (GCBE) in 2021 to bring behavioral expertise in-house. Since then, the center has successfully completed more than 50 projects across different business lines and geographies. At the heart of the GCBE lies an international team of behavioral science experts dedicated to supporting Allianz entities worldwide in tackling challenges relating to human behavior change. These challenges range from increasing honest claim reporting to increasing the uptake of efficient, digital service channels. Our mission is to apply BE insights where they deliver financial impact.

What began as a strategic in-house initiative has now evolved beyond a traditional cost center, transforming into a profit-generating function.² As we move forward, our mission extends beyond Allianz, bringing evidence-based behavioral insights to the broader market.

How We Make It Work

As many behavioral science practitioners are keenly aware, incorporating BE into business processes is not just about running interesting experiments, it is also about delivering real value. By acting as an in-house consultancy, we work side by side with teams across functions and geographies to solve problems with practical significance. We believe that three factors have helped us in this process.

First, we operate in a unique setting: a global structure with local understanding. And because we are part of Allianz, we are often seen as peers—close enough to understand local realities, yet independent enough to challenge the status quo and bring change. Our centralized team supports Allianz entities across the globe, working closely with local teams to ensure behavioral solutions fit specific market needs. This setup allows us to scale successful interventions while respecting regulatory nuances, cultural norms, and operational constraints, which in turn helps us deliver change that is both scalable and sensitive, without assuming a one-size-fits-all approach. Since we are part of the business, we also have easier access to

1 Corresponding author: ann-kathrin.crede@allianz-suisse.ch

2 Recognizing this potential, Allianz has taken the next step—expanding our services to external clients with the launch of Decyde, our new BE consulting arm.

operational data and internal workflows, which helps us move quicker from insight to implementation.

Second, we try to blend scientific rigor with operational pragmatism. We use randomized controlled trials (RCTs) wherever feasible, but we are also mindful of the pace of business. When full RCTs in which we separate each individual modification are not possible, we shift gears, running sequential pilots, loaded treatments, or pre-post comparisons to ensure we still learn, adapt, and improve. This mindset allows us to keep moving while keeping our standards high.

Third, we translate science into strategy. Our team not only knows behavioral science theory, but is also experienced in applying it in environments with commercial targets, internal politics, and resource constraints. We work closely with stakeholders and make sure behavioral insights align with real business goals. And because we speak multiple languages, we are able to connect with partners in their native tongue, thereby building trust and making collaboration easier across markets.

Ultimately, we believe that this way of working can lead to solutions that are faster, cheaper, and more human-centered than traditional business interventions that often rely on gut feelings. However, BE is only effective if it can be applied in ways that businesses care about. To illustrate how this works in practice, the following sections highlight one of our use cases that shows how small, evidence-based changes can drive measurable results across markets. It is a case that many businesses can relate to: encouraging customers to move from traditional call-based services to faster, self-service digital channels.

Why Wait? The Challenge of Steering Customers Toward Faster Digital Solutions

Insurance transactions are increasingly moving into the digital realm, offering customers self-service options for everything from filing claims to updating policies. These tools promise greater convenience, speed, and flexibility, but many customers still prefer to pick up the phone (Blut et al., 2016; Fernandes & Oliveira, 2021). Why is that? BE offers valuable insights into the challenge of encouraging the adoption of self-service technologies. Many customers remain loyal to traditional channels, often due to limited

awareness, an unclear understanding of the benefits, or a natural preference for familiar routines. While developing robust digital tools is essential, the more complex task lies in driving the behavioral shift needed for their widespread use. To realize fully the potential of digital transformation, we must look beyond technical implementation and proactively support customers in adopting new ways of engaging with our insurance services.

The Use Case

In one of our large-scale projects, we examined the following use case more closely. Imagine you are on the road and your car suddenly breaks down. You need immediate assistance, so you call your insurance provider for roadside help. An automated voice answers, clearly presenting two options: stay on the line to speak with a call center agent, or switch to a faster digital solution via an SMS link and web app. In a project involving Allianz entities across 13 countries, we found that despite the digital channel offering significantly shorter wait times, most customers still chose to stay on the phone. Our challenge was to redesign the automated voice messaging by using behavioral insights to encourage more customers to opt for the digital solution.

Behavioral Audit: Identifying What Helps and Hinders Digital Adoption

During the behavioral audit phase, we explored two key questions: What prevents customers from choosing the digital path (“detractors”), and what would encourage them to do so (“enablers”) ? Based on an in-depth literature review, we identified behavioral drivers that appear to operate independently of the specific context (see Table 1). Given the wide range of countries involved in the project, from European markets like Germany, the Netherlands, and Poland to Australia, India, and Thailand, it was essential to develop solutions that fit each market’s cultural and customer-specific characteristics. To achieve this goal, we collaborated closely with local teams to identify enablers and detractors that were highly relevant to their respective context. Through individual behavioral audit workshops in each country, we gathered valuable insights into local customer behavior. To illustrate with a few examples, in Poland, widespread SMS phishing scams and police warnings

Table 1: Behavioral Drivers for the Adoption of Self-Service Technology

„Enablers“: These Factors Motivate or Facilitate Customers to Choose the Digital Path	
Cognitive Ease	When digital tools are intuitive and easy to use, customers are more likely to engage. This is supported by a clear mental model of what the tool is for and how it works, not just what information it presents. Ease of access and use, along with process fluency, reduce mental effort and increase comfort, making the digital journey feel seamless.
Usefulness / Benefits	Digital solutions that visibly save time and effort encourage adoption. Faster gratification and the flexibility of 24/7 support—accessible at the customer's own pace—enhance perceived value. How the benefits are framed also matters: highlighting both individual and collective gains can further strengthen motivation.
Trust	Trust plays a central role in digital adoption. Signals like data protection assurances and transparent communication about how the process works reduce ambiguity. In some cases, even subtle cues exemplifying a „human touch,“ through design or language, can activate the affect heuristic, making the experience feel more reliable and empathetic.
Social Proof	Customers are more likely to try something new if they know others like them are doing the same. Seeing that a digital tool is actively used by other customers (beyond curated testimonials) reinforces its legitimacy and helps overcome hesitation.
„Detractors“: These Factors Hinder or Discourage Customers From Adopting Digital Solutions	
Status Quo Bias	Customers tend to stick with familiar channels, even when better options are available, because of a preference for the known. This is influenced by omission bias, whereby customers feel more comfortable doing nothing (staying on the phone) than taking an action that introduces uncertainty (switching to digital).
Risk + Loss Aversion	Customers perceive the potential downsides of using digital channels, such as inferior service, missed entitlements, or security concerns, as more significant than the possible gains. The anticipated regret of making the „wrong choice“ can weigh more heavily than actual experience, leading to hesitation or avoidance.
Algorithm Aversion	Many customers feel uneasy about automated decision-making, especially in emotionally or financially sensitive situations. They may perceive algorithms as less trustworthy, less transparent, or unable to account for individual circumstances, thus prompting a preference for human interaction.
Perceived Risk of Failure	There is a fear that digital tools might not work as expected, thus causing delays, errors, or unresolved issues. In high-stress situations, this perceived risk of failure makes human contact feel like the safer, more reliable option.
Perceived Fairness	Some customers suspect that digital channels are primarily designed to reduce costs for the company, potentially at their expense. When a digital service feels less personal or lower in quality, it can trigger a sense of unfairness, especially if the customer believes the company is benefiting more than they are.

Sources: de Andrés-Sánchez & Gené-Albesa (2023), de Andrés-Sánchez & Gené-Albesa (2024), Faulkner et al. (2019), John & Blume (2017).

made customers cautious about clicking on links. This required us to emphasize clearly that the SMS was secure and genuinely sent by Allianz. In India, a general lack of trust in institutions meant that a web app confirming the arrival of a tow truck felt less reassuring than a call center agent personally promising assistance. As a result, we needed to ensure that the information provided by the app was perceived as trustworthy and reassuring. To give you one final example, we discovered that customers in Thailand were required to enter their vehicle ID to request a tow truck—an ID many did not know by heart. As a result, we had to inform customers proactively about the need to look up their vehicle ID in advance, in order

to prevent dropouts during this step of the process. Understanding these market-specific nuances was crucial for designing effective solutions.

From Audit to Action: Designing Behaviorally Informed Solutions

We translated the insights from the behavioral audit directly into the redesign phase, with the goal of aligning key behavioral drivers with targeted, evidence-based solutions. In dedicated redesign workshops, we co-created new versions of the automated voice messages that customers hear at the beginning of the phone call, embedding core BE concepts (see Table 2). For each country, we developed

Table 2: Key Behavioral Concepts Guiding the Redesign

Concept	Description	Application to Use Case
Default	People tend to go with the option that requires the least effort, often the pre-selected or default choice (Johnson & Goldstein, 2003).	In the automated voice message, the digital option was presented as the standard choice, making it feel like the natural next step while still leaving the phone option available.
Ease of Use	When an option feels simple, intuitive, and requires minimal effort to engage with, people are more likely to follow through (Davis, 1989).	The message emphasized how easy and straightforward it was to access and use the digital tool, lowering the perceived effort required to engage with the digital channel.
Tangible Benefits	Making benefits concrete helps individuals visualize how it fits into their own lives, which then increases personal relevance and motivation (Trope & Liberman, 2010).	The message highlighted practical benefits of the digital channel, such as faster assistance and the absence of wait times, to make its value feel tangible and personally relevant.
Trust	People are more likely to trust when they feel confident that a process is safe, legitimate, and managed by a reliable source (Berg et al., 1995).	The automated voice message reinforced trust by clearly identifying the sender, explaining the next steps, and signaling the legitimacy and security of the digital channel.
Social Proof	People are more likely to follow a behavior when they believe that others—especially those like them—are already doing it (Schultz et al., 2007).	The message referenced the growing number of customers choosing the digital option, hence helping reduce hesitation by showing that the behavior is familiar and widely accepted.
Uncertainty Avoidance	People are less likely to try something new when outcomes feel unclear or risky and tend to stick with familiar behaviors instead (Ellsberg, 1961).	The message reduced uncertainty by clearly explaining what the customer could expect from the digital channel, helping to build confidence in the process.

and iterated on multiple versions tailored to local preferences and behavioral patterns, which were then tested through controlled pilots. The version that proved most effective, measured by its success in encouraging the highest percentage of customers to choose the digital option, was ultimately rolled out in that market.

Testing Redesigned Voice Messages in a European Market

To make this more tangible, let us take a closer look at the automated voice messages used in one of the participating European countries. For data privacy reasons, we cannot share the exact scripts, but we can present the rephrased content and underlying ideas.

- The control message (Message 0) informed customers of the two available options: staying on the line or switching to the digital channel. In this version, the default was implicitly set to staying on the phone, which was framed as the standard way to proceed.

We then developed two alternative versions, each incorporating key BE concepts:

- Message 1 focused on tangible benefits, emphasizing that choosing the digital option could significantly reduce wait times. The digital channel was framed as the natural and expected way to continue (default on digital).
- Message 2 highlighted ease of use, assuring customers that requesting assistance via the web app would be quick and simple. Like Message 1, it also presented the digital tool as the default path forward.

To evaluate the effectiveness of the redesigned voice recordings, we conducted a simultaneous A/B test. Customers calling for roadside assistance were randomly assigned to one of three automated voice messages, and their behavior, i.e., whether they remained on the line or switched to the digital option, was systematically tracked.

So, which message was most effective in nudging customers toward the digital channel?

As illustrated in Figure 1, both message 1 and message 2 significantly increased the likelihood of customers opting for the digital solution by more than 20 percentage points compared to the control group ($p < .001$, based on OLS regressions). However, there was no statistically significant difference in




Figure 1: Share of customers selecting the digital option to request roadside assistance in a European market.

effectiveness between either message.

Cultural Differences

A key lesson from this project is the importance of market-specific testing. What resonated in one country did not always translate to success in another. For example, we tested the concept of social proof by emphasizing that an increasing number of customers were opting for the digital channel. While this message positively influenced customer behavior in one market, it had the opposite effect in another, resulting in a decline in digital adoption. As a concrete example, Figure 2 shows the results of testing the social proof message in one DACH (German-speaking) country versus one Anglo-Saxon country.




Figure 2: Change in percentage points of customers choosing the digital option under the social proof treatment vs. baseline.

In the DACH country, the message led to an increase in digital adoption of 18.8 pp (65% vs. 46.2%, $p < .001$, based on OLS regressions), whereas in the Anglo-Saxon country, it triggered a decrease of 18.4 pp (31.7%

vs. 50,1%, $p < .001$, based on OLS regressions). One hypothesis for this stark contrast is cultural differences in the domain of individualism versus collectivism. Customers in the DACH region, which tends to be more collectivist, may feel encouraged by the behavior of others, interpreting it as social validation. In contrast, customers in the more individualistic Anglo-Saxon context may perceive such messaging as intrusive or misaligned with their preference for independent decision-making, particularly when seeking help in a personal situation (Hofstede, 2001). These findings highlight that cultural differences pose a significant barrier to simply scaling behaviorally-informed solutions globally.

Key Outcomes and Learnings

Overall, our behaviorally-informed redesigns successfully increased digital roadside assistance requests across all 13 countries, delivering a faster and more efficient service for customers while generating significant cost savings—in total for all markets, nearly 1 million euros per year—by reducing call center volume. Looking ahead, further research is essential to understand the long-term effects of encouraging digital channel use. Important open questions remain: Do customers continue to choose the digital option in future interactions, or do they eventually revert to calling? How does this shift influence customer satisfaction and trust over time? Investigating whether these behavior changes are sustained is critical to ensuring that the observed impact is not only immediate, but also enduring.

As behavioral science continues to evolve, its practical application offers a valuable lens for improving decision-making and outcomes across a wide range of settings. Our use case on roadside assistance shows how behavioral science can unlock impacts for both customers and the business. By redesigning just one automated voice message, based on local insights and behavioral evidence, we nudged thousands of customers toward faster digital help, thereby reducing wait times and lowering call center volume. And while this case focused on one specific moment in the customer journey, the behavioral patterns behind it, namely, hesitation to act, sticking with the default, or struggling with unclear communication, are not unique to insurance. They appear across sectors, from banking, to logistics, e-commerce, healthcare, and

the public sector. Recognizing these patterns allows organizations across industries to design simpler, more effective experiences that better serve both people and business goals.

THE AUTHORS

Soraya Catoja is a senior behavioral economist at the Global Center for Behavioral Economics. Soraya designs and tests behavioral interventions across multiple markets. Her work focuses on optimizing customer experience, reducing operational inefficiencies, and minimizing fraud, especially in the context of insurance claims. With a background in both consulting and academia, she combines scientific rigor with business relevance.

Ann-Kathrin Crede is an engagement manager at the Global Center for Behavioral Economics at Allianz, where she has worked for the past three years. She holds a PhD in Behavioral Economics from the University of Bern, Switzerland. Before joining Allianz, she worked in consulting, advising public institutions on the application of behavioral insights to policy design. Her work is driven by a strong passion for bridging the gap between academia and practice by translating scientific research into actionable solutions for everyday business and life.

REFERENCES

- Berg, J., Dickhaut, J., & McCabe, K. (1995). Trust, reciprocity, and social history. *Games and Economic Behavior*, 10(1), 122–142. <https://doi.org/10.1006/game.1995.1027>
- Blut, M., Wang, C., & Schoefer, K. (2016). Factors influencing the acceptance of self-service technologies: A meta-analysis. *Journal of Service Research*, 19(4), 396–416. <https://doi.org/10.1177/1094670516662352>
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319–340. <https://doi.org/10.2307/249008>
- de Andrés-Sánchez, J., & Gené-Albesa, J. (2024). Not with the bot! The relevance of trust to explain the acceptance of chatbots by insurance customers. *Humanities and Social Sciences Communications*, 11(1), 1–12. <https://doi.org/10.1057/s41599-024-02621-5>

- de Andrés-Sánchez, J., & Gené-Albesa, J. (2023). Explaining policyholders' chatbot acceptance with an unified technology acceptance and use of technology-based model. *Journal of Theoretical and Applied Electronic Commerce Research*, 18(3), 1217–1237. <https://doi.org/10.3390/jtaer18030062>
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics*, 75(4), 643–669. <https://doi.org/10.2307/1884324>
- Faulkner, N., Jorgensen, B., & Koufariotis, G. (2019). Can behavioural interventions increase citizens' use of e-government? Evidence from a quasi-experimental trial. *Government Information Quarterly*, 36(1), 61–68. <https://doi.org/10.1016/j.giq.2018.10.009>
- Fernandes, T., & Oliveira, E. (2021). Understanding consumers' acceptance of automated technologies in service encounters: Drivers of digital voice assistants adoption. *Journal of Business Research*, 122, 180–191. <https://doi.org/10.1016/j.jbusres.2020.08.058>
- Hofstede, G. (2001). *Culture's consequences: Comparing values, behaviors, institutions, and organizations across nations* (2nd ed.). Thousand Oaks, CA: Sage Publications.
- John, P., & Blume, T. (2017). Nudges that promote channel shift: A randomized evaluation of messages to encourage citizens to renew benefits online. *Policy & Internet*, 9(2), 168–183. <https://doi.org/10.1002/poi3.148>
- Johnson, E. J., & Goldstein, D. (2003). Do defaults save lives? *Science*, 302(5649), 1338–1339. <https://doi.org/10.1126/science.1091721>
- Schultz, P. W., Nolan, J. M., Cialdini, R. B., Goldstein, N. J., & Griskevicius, V. (2007). The constructive, destructive, and reconstructive power of social norms. *Psychological Science*, 18(5), 429–434. <https://doi.org/10.1111/j.1467-9280.2007.01917.x>
- Trope, Y., & Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological Review*, 117(2), 440–463. <https://psycnet.apa.org/doi/10.1037/a0018963>

Preventive Screening and Habit Formation in the Discovery-Vitality Shared-Value Health Promotion and Disease Prevention Program

DEEPAK PATEL¹, NANXI LI, CARLA POOL, NATASHA SILVERMAN,
SANDRA LEHMANN AND MOSIMA MABUNDA

Vitality, South Africa

BRETT ROWLAND AND EMILIO GUERRA

Discovery Health

Healthcare systems and insurers worldwide are facing the dual pressures of aging populations and the rising prevalence and costs of chronic diseases. This creates significant strain on the sustainability of healthcare systems. The Discovery-Vitality shared-value model mitigates this strain by not only funding healthcare, but also making members healthier and helping them manage health risks. Members of the Discovery Health Medical Scheme (DHMS), who are also members of the Vitality program, have significantly better screening rates for chronic diseases and cancers, better clinical risk measures, and incur lower healthcare costs than members not on Vitality. Using Discovery and Vitality's rich data, we have developed a composite measure called the "habit index" to measure the management of a member's chronic condition. The habit index is a precursor to the Personal Health Pathways program, which provides personalized recommendations to all DHMS members, guiding each person on the most effective path to better health.

Introduction

It has long been recognized that trends in health and illness vary over time and across communities and continents (Jones et al., 2012). Globally, non-communicable chronic diseases (NCDs), such as heart diseases, stroke, hypertension, cancers, chronic lung diseases, and diabetes, dominate the medical landscape and accounted for approximately 74% of all deaths in 2019, while communicable diseases (grouped as infectious diseases, maternal and perinatal deaths, and deaths from injuries), dropped from 32% in 2000 to 18% in 2019 (World Health Organization, 2024a).

The COVID-19 pandemic temporarily disrupted this trend. In 2021, deaths from communicable diseases rose to 28%, while NCD-related deaths dropped to 65.3% (World Health Organization, 2024a). Despite the COVID-19 jolt, however, the global trend of an increase in NCDs is likely to resume as the pandemic

abates. This is in part because rates of screening for NCDs, including cancers, dropped significantly during the pandemic, resulting in a lower incidence (Cancino et al., 2020).

South Africa faces a high burden of diseases (Groenewald et al., 2024). In 2018, NCDs accounted for 58.9% of deaths, while communicable diseases, including HIV/AIDS and TB, were responsible for 29.1%, and injury and violence for 12% of deaths (Stats SA, 2025). Those with access to private healthcare in South Africa have a much lower prevalence of infectious diseases, including HIV, and a greater prevalence of chronic disease (Cairncross & Govuzela, 2019; World Health Organization, 2021).

What Is Responsible for the Rise in NCDs?

The rise in chronic diseases over the last half century is due, in part, to a reduction in early death

¹ Corresponding author: deepakp@discovery.co.za.

from infections, as well as perinatal and maternal illnesses. The resulting increase in life expectancy (aging) is associated with an increased incidence of NCDs. However, changes in lifestyle behaviors such as unhealthy eating, including excess calorie intake, physical inactivity, increased sedentary behavior, and excessive tobacco and alcohol use, play an important role in this regard (World Health Organization, 2024b).

Changing health behavior to prevent and manage chronic conditions presents several behavioral and structural challenges. Behavioral economics has shed considerable light on how individuals make systematic errors in health decisions because of ‘bounded rationality’ and cognitive bias (Simon et al., 1955). The following selection of biases contribute to unhealthy lifestyles and chronic diseases.

Present bias and discounting the future: many risk factors for chronic diseases, such as obesity, poor diet, and inactivity, begin in childhood and adolescence, when there is a higher rate of discounting the future and the tendency to favor immediate rewards over long-term benefits (present bias) (Loewenstein et al., 2013; Green et al., 1994).

Self-enhancement bias and over-optimistic bias: refers to the tendency for people to take a favorable view of themselves and overestimate their abilities and the likelihood of positive outcomes. Many smokers, for instance, acknowledge that smoking is a health risk, but they believe that they—somehow—will be spared from its deadly consequences (Patton et al., 2022; Peretti-Watel et al., 2007; Sedikides and Gregg, 2008).

Overweighting small probabilities and underweighting high probabilities: refers to the disproportionately high importance individuals attach to things that occur relatively uncommonly, such as dying in a plane crash, and the relative neglect of things that occur commonly, such as succumbing to heart disease, diabetes, or cancer (Kahneman & Tversky, 1979). Related to underweighting of small probabilities is the *peanut effect*, which refers to the tendency to assign little weight to actions that have small, immediate effects, such as smoking a cigarette, which has imperceptible proximate ill-effects but will have a cumulative effect if the behavior is maintained over a long period (Loewenstein et al., 2013).

For funders of healthcare, the ever-increasing rise in the prevalence of chronic diseases and the

associated increased volumes of patients on chronic medication, compounded by the escalating costs of new treatments and technologies, contributes to unsustainably high healthcare costs (OECD, 2015; Muka et al., 2015).

The Discovery Vitality Program

To address the rising prevalence of and costs related to NCDs, the Discovery Health Medical Scheme (DHMS), the largest private health plan in South Africa, has offered the Vitality behavior-change program to its members. In South Africa, the program has grown organically over almost three decades and is offered to DHMS and Discovery Life insurance plan members at a nominal monthly cost of Rands 559 (+/- 31 US dollars) for a family of three or more.

In 2023, we published a paper in the BE Guide (Patel et al., 2023) outlining the Discovery-Vitality model of behavior-change and how it incorporates health promotion and disease prevention for members of the DHMS and Discovery Life insurance plans. The Discovery-Vitality shared-value model is a distinctive model that not only funds healthcare but also seeks to improve the lifespan and health-span (years lived in good health) of insured members (Porter & Kramer, 2011; Porter et al., 2014). In the 2023 BE Guide paper (Patel et al., 2023) we quantified the cost saving for the DHMS from selection effects (healthy people selecting into the program) and from behavioral change on the program. For a broad overview of the health interventions and some of the rewards offered on the program, see Figure 1.

Health Promotion and Risk-Reduction in Vitality

A primary focus of Vitality is to address—from prevention to management—the ever-increasing burden of chronic diseases amongst members. Prevention is a broad term encompassing a wide range of interventions aimed at averting or delaying the occurrence of disease. In the context of NCDs and Vitality, this includes health-promoting activities such as staying physically active, eating healthily, maintaining a healthy weight, not smoking, and avoiding or limiting the intake of alcohol.

Promoting and enabling physical activity is a key feature of the program. Vitality offers several physical activity interventions, including steeply discounted access to a countrywide network of gyms, membership




Figure 1: The Vitality program is structured around four foundational pillars—Know Your Health, Improve Your Health, Reduce Your Risk, and Get Rewarded, which collectively aim to promote proactive health management by encouraging individuals to understand their health status, adopt healthier behaviors, and receive rewards for engagement.

of a running and cycling club called “Team Vitality,” substantial discounts on wearable devices, and sports gear². Vitality “points,” which are awarded for physical activity participation, are logged onto a smartphone platform called Vitality Active Rewards (VAR)³. We have previously described VAR (Patel et al., 2018) and how behavioral tools are embedded therein (Patel et al 2023). Our published research has also shown positive associations between physical activity engagement and improved healthcare costs (Patel et al., 2011), cancer outcomes (Mabena et al., 2025), and COVID-19 severity (Steenkamp et al., 2022).

Another key Vitality health-promoting intervention is the HealthyFood Benefit,

which offers rebates of 25% (on completion of Vitality Age and a Vitality Health Check—see details below) on purchases of thousands of healthy food items at two large supermarket chains⁴. The qualifying foods, which are labeled with a “V” logo in the stores, include fruits, vegetables, wholegrains, legumes, fish, skinless chicken, nuts and seeds, and fat-free dairy. The rebate is awarded in the app as “Discovery Miles,” which can be spent as a currency at a range of services and stores or can be monetized. We have previously reported on the positive association between the HealthyFood rebate and purchases of healthy foods (Sturm et al., 2013; An et al., 2013).

Vitality also offers discounted access to several risk-reduction programs: The HealthyWeight program is an online, interactive, dietician-supported weight-loss program especially for people who are overweight⁵, and smokers who wish to quit have discounted access to two in-person programs: Allen Carr and GoSmokeFree, a nurse-led program available at partnering pharmacy chains⁶.

What This Paper Offers

In the current paper we discuss a continuum of Vitality’s existing and newer interventions directed at identifying individuals at risk of chronic disease and managing it. More specifically, we discuss:

1. The Vitality Age online assessment

2. The Vitality Health Check
3. Screening for cancers
4. New measures of clinical habits and personalized pathways, to prevent and manage chronic diseases.

We report on outcomes associated with participation in the program. Furthermore, we compare Vitality members with non-Vitality members, we look at costs related to clinical outcomes, and we report projected lifespan and health-span improvements associated with changing habits with personalization.

The Vitality Age Assessment

Vitality Age is a purpose-built, evidence-based health risk algorithm administered via an online assessment. The assessment incorporates self-reported health behaviors, such as smoking, diet, and physical activity, and measured clinical risk factors, such as blood glucose, cholesterol, weight status, and blood pressure, to determine an individual’s risk, expressed as a “Vitality Age”. In essence, one’s Vitality Age is their actual age adjusted by the difference between population and individual life expectancy.

The Vitality Age assessment is both a health-risk assessment and a behavior-change tool. Individuals whose Vitality Age is lower than their actual age are expected to maintain their “younger,” healthier status by continuing to engage in health-promoting activity. Conversely, individuals with a Vitality Age higher than their actual age are expected to be motivated by loss aversion to engage in health-promoting and preventive activity (Loewenstein et al., 2013).

Figure 2 shows the activation of the HealthyFood benefit, and Figure 3 shows the activation of the Vitality Active Rewards benefit by the frequency of Vitality Age completion over the most recent three-year period. The more frequently members complete the Vitality Age assessment, the greater the association with activating health-promoting benefits on the program. Our analysis reveals that activating the HealthyFood benefit and the Gym benefit was higher amongst people with an age gap of

2 Fitness interventions: <https://www.discovery.co.za/vitality/get-fit>

3 Vitality Active Rewards: <https://www.discovery.co.za/vitality/active-rewards>

4 HealthyFood Benefit: <https://www.discovery.co.za/vitality/healthy-eating>

5 HealthyWeight: <https://www.discovery.co.za/vitality/healthyweight>

6 Smoking cessation programs: <https://www.discovery.co.za/vitality/smoking-cessation>




Figure 2: Activation of the HealthyFood Benefit by the frequency of Vitality Age assessment completion between 2022 and 2024. The more frequently members complete the Vitality Age assessment, the more likely they are to activate the HealthyFood Benefit.




Figure 3: Activation of Vitality Active Rewards by the frequency of Vitality Age assessment completion between 2022 and 2024. The more frequently members complete the Vitality Age assessment, the more likely they are to activate Vitality Active Rewards.

up to three years compared to people whose Vitality Age was equal to or lower than their actual age.

The Vitality Health Check

The Vitality Health Check (VHC) is an in-person, adult health assessment (commonly done by a nurse practitioner at a pharmacy, or on a company wellness day). It entails measuring weight status (body mass index and waist circumference), blood pressure, blood glucose, and finger-prick cholesterol, as well as making a non-smoker declaration. In addition, as part of a VHC, members are offered HIV testing, after appropriate counselling.

The VHC is intended to identify and appropriately refer members with out-of-range metrics to primary care physicians (PCPs) for further diagnostic testing and management. Members with chronic conditions get their metrics measured at the PCPs and not at a pharmacy. This is to ensure holistic management of their chronic condition and for them to get appropriate laboratory tests instead of point-of-care tests.

Currently, more than 350,000 VHCs are completed every year. The number of members completing the VHC declined significantly during the COVID-19 pandemic but has recovered and is now exceeding pre-COVID-19 levels.

The VHC allows for stratification of the Discovery-Vitality population and referrals for diagnosis and, where required, recruitment into various Discovery care and health coaching programs. Approximately 38% of all Diabetes Care Program, 17% of Diabetes Prevention Program, and 16% of Cardio Care program enrollments occur within 20 days of completing the VHC (DHMS data).

Nudging Members to Do a Vitality Age Assessment and Vitality Health Check

Gateway to Health-Promoting Benefits and Rewards

Completing a Vitality Age assessment and the VHC allows members to obtain a 25% rebate on purchases of healthy food items via the HealthyFood benefit. Moreover, completing these tasks unlocks discounts on numerous other health-related benefits, such as ActiveGear (sports gear, equipment, and wearable devices), HealthyCare (dental, eye, and sun care, stocked at partnering pharmacies), HealthyDining (healthier meals at partnering restaurants), HealthyBaby (baby care products), and discounts on local and international flights. Primarily, completing these assessments reinforces health-promoting behaviors.




Figure 4: Vitality Health Check results for Vitality compared to non-Vitality members. Vitality members have more readings in range (except for cholesterol) than members with no Vitality.

Easy Access

The Vitality Age assessment is easily accessible online, and the cost of the VHC is paid from the risk-pool of the DHMS, meaning that members incur no out-of-pocket costs at the point of care.

Rewarding with Vitality Points and Status

Members earn Vitality Points for completing the assessments. Furthermore, for the VHC, additional points are allocated for parameters that fall within the normal range. Individuals who have all parameters within the normal range can earn 22,500 Vitality points, which equates to 45% of the points required annually to reach the highest Vitality status (Diamond) for a single member. Members who maintain a Diamond status can get up to 50% of their Discovery Life premiums back. Our research shows that members who achieve Diamond status are inclined to continue a high level of engagement in the program to retain their status—motivated, perhaps, by loss aversion. Policy lapse rates are also lower amongst members with higher statuses.

The Impact of Vitality on VHC Outcomes and Costs

Figure 4 shows that DHMS Vitality members have more VHC readings in range (except for cholesterol)

than members of DHMS with no Vitality.

BMI: Body mass index; **BP:** Blood pressure; **CHOL:** Cholesterol; **GLUC:** Glucose

Moreover, Figure 5 illustrates that after adjusting for gender, age, and plan type, healthcare costs per life per month (PLPM) are considerably lower the greater the number of VHC measures “in range”.

Screening for Cancer

In addition to the VHC, Vitality awards points for other recommended screenings.

Figures 6 to 8 offer comparisons of longitudinal screening rates for the common points-earning screening tests between Vitality and Non-Vitality members for the years 2015 to 2024. Members on Vitality consistently and significantly ($p < 0.05$) have a higher screening rate than non-Vitality members.

The number of members completing general screening (offered by the VHC) and condition-specific screening for diabetes, breast cancer, cervical cancer, and prostate cancer all declined significantly during the Covid-19 pandemic, creating a screening deficit. Screening for these conditions dropped by 21% on average relative to pre-COVID-19 levels on the DHMS. Rates picked up in 2023 and 2024, and




Figure 5: Risk-adjusted costs per life per month (PLPM) by number of Vitality Health Check measures in range, baselined to 0 in range. After adjusting for gender, age, and plan type, healthcare costs are considerably lower the greater the number of VHC measures “in range”.




Figure 6: Eligible members completing a mammogram per 1,000 lives. Members on Vitality consistently have a higher screening rate compared to non-Vitality members.




Figure 7: Eligible members completing cervical cancer screening per 1,000 lives. Members on Vitality consistently have a higher screening rate compared to non-Vitality members.



Figure 8: Eligible members completing a prostate screening per 1,000 lives. Members on Vitality consistently have a higher screening rate compared to non-Vitality members.




Figure 9: Estimated percentage increase in use of screening tests after joining Vitality. After joining Vitality, members of the Discovery Health Medical Scheme increased their screening rates.

they are expected to continue increasing.

In a previous study (Mehrotra et al., 2014), we determined that DHMS members who joined Vitality increased their use of screening services, as shown in Figure 9.

The Habit Index and Personal Health Pathways

A habit is a specific, ingrained behavior that becomes automatic through repetition. It is performed regularly, and it is often initiated by an external cue. Habits are typically formed through consistent practice and can be beneficial or detrimental (Verplanken & Orbell, 2022).

Using extensive behavioral and clinical data, Discovery developed two habit measures: the Physical Activity Habit Index and the Clinical Habit Index. The Clinical Habit Index (CHI) is a composite measure of a member's overall management of their condition on a consistent basis.

Guided by clinical guidelines, each clinical action has a recommended frequency, referred to as a “cycle,” to maintain ongoing health and prevent complications. A person with uncomplicated type 2 diabetes, for instance, should collect their medication monthly, visit their health practitioner and do a HbA1c biannually, visit their ophthalmologist and podiatrist, and do a test for kidney function and a cholesterol test annually. Our research shows that completing recommended actions in three consecutive cycles predicts an 80% chance of future adherence—an indication that the member has likely formed a habit of performing that action (See Appendix 1 for the CHI formula). This index (ranging from 0-1) summarizes the overall habits displayed in managing the condition, with a score of 0 – 0.33 indicating poor management, a score up to 0.66 a medium habit, and a score greater than 0.66 reflecting a strong habit, and consistent completion of all clinically relevant actions.




Figure 10: Normal Cumulative Distribution Function (CDF) of the Clinical Habit Index (CHI) for people with diabetes, stratified by Vitality engagement. At least 55% of high-engaged Vitality members with diabetes have a strong CHI.

Clinical Habit Index of Vitality and Non-Vitality Members With Diabetes

Figure 10 shows the Habit Index of members registered for diabetes. The cohort was split by “not on Vitality,” “low-engaged on Vitality,” and “high-engaged on Vitality.” The X axis is marked at the strong habit level. The y-axis represents the cumulative probability, ranging from 0 to 1 of the proportion of members less than or equal to the corresponding habit.

For members with diabetes, roughly 55% of high-engaged Vitality members have at least a strong CHI, indicating they are considerably engaged in the management of their condition. When we compare this to non-members and low-engaged Vitality members, only about 35% have at least a strong CHI, a difference of about 20%. This distinction persists for people with hypertension and with hyperlipidemia, which suggests a strong correlation between our most engaged Vitality members with chronic conditions and their consistency of care. Our analysis also establishes that members who identify their diabetes risk at an early stage, and manage this risk and their illness effectively, can increase their health span by 8 years and add 5 years to their lifespan.

Personal Health Pathways and Behavior-Change

The Habit Index is the basis for the recent introduction of Personal Health Pathways (PHPs)—a behavior-change program that guides all adult DHMS members to adopt healthier habits and rewards them along the way. Key to behavior change is personalization, namely delivering personalized health and lifestyle recommendations, tailored messaging and individualized rewards, based on an individual’s specific health status, prior engagement, and preferences. Depending on the member’s action (or inaction), PHPs refresh and offer the next best action (NBA), building from simple to more complex behavior-change in a staged manner. The program leverages 80 million Discovery and Vitality life-years data, combined with advanced risk segmentation, machine-learning, and data science capabilities, to guide members along the most effective paths to complete their personalized health action. In parallel with machine-learning, Vitality conducts randomized controlled testing of messages, incentives and rewards, in different cohorts, to feed into the

PHP model. Moreover, the Vitality Age assessment is currently being reworked to make it a more dynamic, personalized, and engaging measure, responsive to changes in lifestyle behaviors such as increasing (or decreasing) physical activity and changes in clinical parameters. A dynamic Vitality Age will enhance the behavior-change features of PHP. The program is available on the Discovery Health app and through an intuitive WhatsApp journey, making it easier for members to form healthy habits.

Conclusion

By offering a range of health-promoting and risk-reduction interventions, improving screening uptake, incentivizing preventive care, and guiding members through personalized care pathways, Vitality contributes to better health outcomes and cost savings for individuals and the healthcare system. In this paper we have:

- Recorded how members who regularly complete the Vitality Age risk assessment are more likely to activate other health-promoting activities on the program.
- Determined that Vitality members have better VHC health parameters compared to non-Vitality members. The costs of healthcare reduce when more parameters are in range.
- Shown significant differences in cancer screening rates between Vitality members compared to non-Vitality members.
- Discussed a new composite measure of clinical adherence for chronic conditions, called the Clinical Habit Index. We have also described a new personalized, iterative machine-learning program called Personal Health Pathways, which guides all Discovery Health members on how they can form healthier habits.

The Vitality shared-value model is available in the UK as Vitality Health and Vitality Life, and it is available in partnership with insurers in more than 40 countries, covering more than 42 million lives. Vitality Health International is offered in seven other countries in Africa. Many interventions in the South Africa Vitality program, such as Vitality Active Rewards, Vitality Age, and the VHC, are already integrated into the global businesses. New innovations such as the CHI and PHP, as well as learning from their implementation, will be adopted by international

Vitality businesses.

The Discovery-Vitality shared-value model effectively integrates behavioral science, incentive design, risk stratification, and clinical habit formation to shift healthcare upstream and toward prevention, personalized care, and sustainable habit formation. It is an equitable, positive-sum game model that benefits members, healthcare funders, and society alike.

THE AUTHORS

Deepak Patel is a principal clinical specialist and head of research at Vitality South Africa. He obtained his undergraduate medical degree and his postgraduate pediatric qualifications from the University of the Witwatersrand, South Africa. He holds an Honors Degree in Development Studies from the University of Witwatersrand and a Master's in Sports Medicine from the University of Cape Town. He is an honorary lecturer in the Division of Sports & Exercise Medicine, Faculty of Health Sciences, University of the Witwatersrand.

Brett Rowland is a senior data scientist at Discovery Health, South Africa. He obtained his undergraduate degree in Actuarial Science and a Master's in Mathematical Statistics from the University of Pretoria. He is also a Technical Member of the Actuarial Society of South Africa.

Emilio Guerra is a data scientist in the Group Data Science Lab at Discovery South Africa. He obtained his undergraduate medical degree at the University of Pretoria and holds a Master's in Health Data Science from University College London.

Nanxi Li is a Technical Member of the Actuarial Society of South Africa and is currently an actuarial analyst at Vitality South Africa. He obtained a Bachelor of Business Science degree from the University of Cape Town.

Carla Pool is a registered dietitian for Vitality South Africa. She obtained both her undergraduate degree in Human Life Sciences and an Honor's Degree in Physiology from the University of Stellenbosch, as well as an Honor's Degree in Nutrition and Dietetics and a Master's in Public Health from the University of Cape Town.

Natasha Silverman is an actuarial manager at Discovery Vitality South Africa. She obtained her

actuarial science degree and Honor's from the University of the Witwatersrand, South Africa.

Sandra Lehmann has worked across the Discovery Group, both locally and internationally. Since 2020, she has been the Chief Actuary of Vitality South Africa, heading up the company's data and advanced analytics initiatives. Her team develops data-driven solutions, provides actuarial insights and guidance, leads data science efforts, oversees risk management, and conducts research analytics.

Mosima Mabunda is the Chief Medical Officer and Head of Vitality Wellness, South Africa. She obtained her undergraduate medical degree from the University of Cape Town and a Master of Business Administration from the University of Oxford.

REFERENCES

- An, R., Patel, D., Segal, D., & Sturm, R. (2013). Eating better for less: A national discount program for healthy food purchases in South Africa. *American Journal of Health Behavior*, 37(1), 56–61.
- Cairncross, C., & Govuzela, M. (2019). *Prevalence of chronic diseases in the population covered by medical schemes in South Africa*. Council for Medical Schemes. <https://www.medicalschemes.com/-/files/Research%20Briefs/Prevalence-of-chronic-diseases-in-the-medical-schemes-population.pdf>
- Cancino, R. S., Su, Z., Mesa, R., Tomlinson, G. E., & Wang, J. (2020). The impact of COVID-19 on cancer screening: Challenges and opportunities. *JMIR Cancer*, 6(2), e21697. <https://doi.org/10.2196/21697>
- Green, L., Fry, A. F., & Myerson, J. (1994). Discounting of delayed rewards: A life-span comparison. *Psychological Science*, 5(1), 33–36. <https://www.jstor.org/stable/40062338>
- Groenewald, P., Nannan, N., Joubert, J. D., Glass, T., Cheyip, M., Maqungo, M., Funani, N., Zinyakatira, N., Awotiwon, O., Nojilana, B., Kallis, N., Laubscher, R., Bezuidenhout, F., Clark, S. J., Kabudula, C., Martin, L., Kahn, K., Price, J., ... Bradshaw, D. (2024). *South African National Cause-of-Death Validation Project: Agreement and corrected cause specific profiles based on*

- data linkage. South African Medical Research Council. <https://www.samrc.ac.za>
- Jones, D. S., Podolsky, S. H., & Greene, J. A. (2012). The burden of disease and the changing task of medicine. *New England Journal of Medicine*, 366(25), 2333–2338. <https://doi.org/10.1056/NEJMp1113569>
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–292. <https://www.jstor.org/stable/1914185>
- Loewenstein, G., John, L., & Volpp, K. G. (2013). Using decision errors to help people help themselves. In E. Shafir (Ed.), *The behavioral foundations of public policy* (pp. 361–379). Princeton University Press.
- Mabena, N., Rugbeer, N., Lehmann, S., Torres, G., Patel, D., Mabunda, M., Greyling, M., Thornton, J. S., Choi, Y. H., Stranges, S., & Patricios, J. S. (2025). Association between recorded physical activity and cancer progression or mortality in individuals diagnosed with cancer in South Africa. *British Journal of Sports Medicine*. <https://doi.org/10.1136/bjsports-2024-108813>
- Mehrotra, A., An, R., Patel, D. N., & Sturm, R. (2014). Impact of a patient incentive program on receipt of preventive care. *American Journal of Managed Care*, 20(6), 494–501. <https://pmc.ncbi.nlm.nih.gov/articles/PMC4167869/>
- Muka, T., Imo, D., Jaspers, L., Colpani, V., Chaker, L., van der Lee, S. J., Mendis, S., Chowdhury, R., Bramer, W. M., Falla, A., & Pazoki, R. (2015). The global impact of non-communicable diseases on healthcare spending and national income: A systematic review. *European Journal of Epidemiology*, 30(4), 251–277. <https://doi.org/10.1007/s10654-014-9984-2>
- National Health Service (NHS) Digital. (2023, February 16). *NHS Breast Screening Programme, England 2021–22*. <https://digital.nhs.uk/data-and-information/publications/statistical/breast-screening-programme/england---2021-22>
- Organisation for Economic Co-operation and Development (OECD). (2015). *Fiscal sustainability of health systems: Bridging health and finance perspectives*. OECD Publishing. <https://doi.org/10.1787/9789264233386-en>
- Patel, D., Lambert, E. V., Da Silva, R., Greyling, M., Kolbe-Alexander, T., Noach, A., Conradie, J., Nossel, C., Borresen, J., & Gaziano, T. (2011). Participation in fitness-related activities of an incentive-based health promotion program and hospital costs: A retrospective longitudinal study. *American Journal of Health Promotion*, 25(5), 341–348. <https://doi.org/10.4278/ajhp.100603-QUAN-172>
- Patel, D., Moche, L., Singh, K., Joseph, C., Lehmann, S., & Mabunda, M. (2023). Nudging toward good health: Leveraging behavioural science in the shared-value insurance model. In A. Samson (Ed.), *The Behavioral Economics Guide 2023* (pp. 94–104).
- Patel, D. N., Nossel, C., Patricios, J., & Maboreke, J. (2018). Bright spots, physical activity investments that work: Vitality Active Rewards—a smartphone app that incentivises programme members to be physically active. *British Journal of Sports Medicine*, 52(23), 1494–1496. <https://doi.org/10.1136/bjsports-2018-099271>
- Patton, S. R., Cushing, C. C., & Lansing, A. H. (2022). Applying behavioral economics theories to interventions for persons with diabetes. *Current Diabetes Reports*, 22(5), 219–226. <https://doi.org/10.1007/s11892-022-01487-2> [Correction published in 2023, *Current Diabetes Reports*, 23(7), 173.]
- Peng, X., Wan, L., Yu, B., & Zhang, J. (2025). The link between adherence to antihypertensive medications and mortality rates in patients with hypertension: A systematic review and meta-analysis of cohort studies. *BMC Cardiovascular Disorders*, 25(1), 145. <https://doi.org/10.1186/s12872-025-04538-6>
- Peretti-Watel, P., Constance, J., Guilbert, P., Gautier, A., Beck, F., & Moatti, J. P. (2007). Smoking too few cigarettes to be at risk? Smokers' perceptions of risk and risk denial, a French survey. *Tobacco Control*, 16(5), 351–356. <https://doi.org/10.1136/tc.2007.020362>
- Porter, M. E., & Kramer, M. R. (2011). Creating shared value. *Harvard Business Review*, January–February. <https://hbr.org/2011/01/the-big-idea-creating-shared-value>
- Porter, M., Kramer, M., & Sesia, A. (2014). *Discovery Limited* (Case Study 715–423). Harvard Business

- School. <https://hbr.org/product/discovery-limited/715423-PDF-ENG>
- Sedikides, C., & Gregg, A. P. (2008). Self-enhancement: Food for thought. *Perspectives on Psychological Science*, 3(2), 102–116. <https://doi.org/10.1111/j.1745-6916.2008.00068.x>
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1), 99–118. <https://www.jstor.org/stable/1884852>
- Statistics South Africa. (2025). *Mortality and causes of death in South Africa, 2021: Findings from death notification* (Statistical Release P0309.3). <https://www.statssa.gov.za/publications/P03093/P030932021.pdf>
- Steenkamp, L., Saggars, R. T., Bandini, R., Stranges, S., Choi, Y. H., Thornton, J. S., Hendrie, S., Patel, D., Rabinowitz, S., & Patricios, J. (2022). Small steps, strong shield: Directly measured, moderate physical activity in 65,361 adults is associated with significant protective effects from severe COVID-19 outcomes. *British Journal of Sports Medicine*, 56(10), 568–577. <https://doi.org/10.1136/bjsports-2021-105159>
- Sturm, R., An, R., Segal, D., & Patel, D. (2013). A cash-back rebate program for healthy food purchases in South Africa: Results from scanner data. *American Journal of Preventive Medicine*, 44(6), 567–572. <https://doi.org/10.1016/j.amepre.2013.02.011>
- Verplanken, B., & Orbell, S. (2022). Attitudes, habits, and behavior change. *Annual Review of Psychology*, 73, 327–352. <https://doi.org/10.1146/annurev-psych-020821-011744>
- World Health Organization. (2021). *World health statistics 2021: Monitoring health for the SDGs, sustainable development goals*. World Health Organization. <https://apps.who.int/iris/bitstream/handle/10665/342703/9789240027053-eng.pdf>
- World Health Organization. (2024a). *World health statistics 2024: Monitoring health for the SDGs, sustainable development goals*. Geneva: World Health Organization.
- World Health Organization. (2024b, December 23). Noncommunicable diseases. <https://www.who.int/news-room/fact-sheets/detail/noncommunicable-diseases>

APPENDIX

Clinical Habit Index Formula

$$\frac{1}{N} \sum_{i=1}^N \frac{w_i \times \text{completed cycles}}{\min(\text{cycle exposure}, 3)^*} = w_1 \left(\frac{1}{3}\right) + w_2 \left(\frac{2}{3}\right) + w_3 \left(\frac{1}{3}\right) + w_4 \left(\frac{3}{3}\right) + w_5 \left(\frac{2}{3}\right) + w_6 \left(\frac{3}{3}\right)$$

N : Total number of clinically relevant actions available for a member.

w_i : Normalized value-weight of the i -th clinical action captures its relative importance in shaping the overall CHI. This weight is determined by models that evaluate the economic impact of each action, assigning higher weights to actions that generate greater economic value. By doing so, the CHI is more significantly influenced by high-value actions.

From Points to Profits: Designing the Optimal Loyalty Programme

HENRY STOTT, JERRY LUUKKONEN¹ AND MIRIAM MAKIN

Dectech

Loyalty programmes have a long history dating back to 1851 or earlier. Today, the global loyalty programme market is valued at approximately \$200 billion. While their common purpose is to incentivise repeat purchasing, some loyalty schemes are considerably more successful than others. This report draws on data from our own Behaviourlab experiments to understand which features optimise or hinder loyalty scheme performance, and what their business outcomes are. A well-designed loyalty scheme, incorporating our top-performing features, could attract 30% more customers over a five-year period. However, our research also highlights that the same features can perform very differently across sectors. In this report, we provide practical recommendations on how you can tailor a loyalty programme to your particular market.

Executive Summary

Loyalty programmes have a long history dating back to at least 1851. Yet, some are considerably more successful than others. In this report, we discuss the different commercial objectives sought from these schemes, how you can optimise their design, and what outcomes you should expect. The main conclusions are:

Reach: A typical loyalty scheme will reach between 20% and 90% of a retailer's base. In part, the higher penetration rates are down to longevity. Schemes sign up around 17% of the base at launch and grow at +2.1% per year thereafter.

Effects: Loyalty programmes target a combination of commercial benefits. Behavioural – an increased share of wallet. Conative – a relationship that drives cross-sales. Attitudinal – brand liking that generates intangible value.

Acquisition: Using Behaviourlab, our large-scale online randomised controlled trial approach, we find that sign-up incentives and exclusive access work best for acquisition. Moreover, a better design will add +30% to your base over 5 years.

Participation: US consumers join 14.8 schemes on average, and yet they actively use only half. Given an active loyalty programme member spends x4.6 more

than a non-member, the value of activating sign-ups can fund compelling rewards.

Activation: Features such as personalised offers and partner firm discounts help increase participation. A well-designed scheme can triple participation rates. The features that drive acquisitions and participation are different.

Optimisation: There are some general rules around loyalty programme design, albeit the details vary by category. Free delivery works best where online sales are greater. Personalised offers work best for higher-ticket products.

Bake-off: Nectar has the best design for maximising behavioural loyalty, in that they include the highest-impact features and communicate them well with their members.

Introduction

Loyalty programmes have been around for nearly two centuries. Their purpose is to incentivise repeat purchasing through the accumulation of points that can be traded in for discounts or other rewards and can come in many different forms (Kim et al., 2021). Contrasting loyalty programmes with fiat currencies and asset-backed securities, they have five defining characteristics. First, loyalty programmes

¹ Corresponding author: j.luukkonen@dectech.co.uk

are operated by non-sovereign entities, typically private companies. Second, there is no underlying asset, like grain or gold, that is used either directly or to underpin the value of a promissory note. Third, a loyalty programme does not accrue interest, unlike, say, a private currency. Fourth, a loyalty programme accumulates when you spend money at – or undertake some activity to the commercial benefit of – the issuer(s). Fifth, any accumulation can ultimately only be exchanged for goods or services available from the issuer(s).

Based on these definitions, the first loyalty programme probably started in 1851 when Benjamin Babbit sold his eponymous Best Soap with “trade marks” that could be cut from the soap packaging and exchanged for a company lithograph. There then followed a century of similar programmes by US retailers, culminating in American Airlines’ Aadvantage frequent flyer scheme in 1981. Meanwhile, in the UK, Babbit’s idea eventually manifested in the form of Green Shield stamps (1958), which was the OG Clubcard. Today, the global loyalty programme market is worth \$200Bn, according to Beroe, a procurement platform provider, and schemes are offered by a range of brands all competing for customer loyalty.

The UK’s largest loyalty programme is Tesco Clubcard. Launched in 1995, Clubcard has 22m members – about 80% of UK households. Tesco is also the most adopted scheme, with 90% of their shoppers signed up (see Figure 1). Likewise, other schemes that are over 25 years old, namely Boots and Sainsbury’s, enjoy 70% customer adoption. Conversely, Asda, Morrisons, Lidl, and M&S are relative newbies with commensurately lower membership rates. Fitting a line to this trend, about 17% of your customers sign up in the launch phase, and then you add 2.1% per year thereafter. By this measure, the 3-year-old Asda offering, at 49% adoption, is remarkably successful.

Nevertheless, loyalty programmes aren’t just about sign-ups; this is only half the battle. A typical US consumer joins 14.8 schemes but is only active in half, so these schemes need to earn that loyalty with a better product or service. Accordingly, when designing a scheme, what features drive sign-up? Likewise, after people have joined, how do you activate and nurture their participation?

This paper sets out to answer what combination of features best drives both sign-ups and customer participation. First, we consider existing findings and literature on the effectiveness of loyalty schemes.




Figure 1: Largest loyalty schemes. Source: Dectech fieldwork, May 2024 (n = 2,115 UK loyalty scheme members). Household loyalty scheme membership was measured by asking respondents to indicate which of 21 loyalty schemes they belonged to. Customer adoption rates are defined as the proportion of loyalty scheme membership among customers who have shopped at a store at least once in the past year.

While our study is largely exploratory in nature, the existing research guided us through some key hypotheses, which we briefly discuss. Second, we examine the impact different scheme features have on sign-up. Specifically, we do this through a loyalty scheme sign-up task we designed, as this made examination much cleaner than attempting to disentangle the effect of brand and other factors on real-world scheme sign-ups. Third, we examine how loyalty scheme features impact purchasing behaviour. Here, we primarily focus on reported purchases with real-world loyalty schemes, as these are more reliable than asking people to imagine their purchase propensity. Finally, we provide recommendations on what to keep in mind when designing or optimising a loyalty scheme.

Commercial Impact

Before we address acquisition and activation, you need to ask why you want a loyalty programme in the first place. What commercial benefits are you targeting? What will be the return on investment? The typical answer to these questions is that you want to increase the share of wallet and thereby sales. A

less common, but possibly more honest, answer is that your main competitor operates one, and you feel obliged to match them.

Figure 2 supports the thesis that having a scheme can drive growth. Sainsbury's and Tesco have well-established and widely used schemes, and they are growing. Lidl has a scheme and is growing faster than Aldi, the latter of which has explicitly said that loyalty programmes are prohibitively expensive.

Deconstructing these growth ambitions, the research literature tends to coalesce around three different types of loyalty, each of which generates higher enterprise value in different ways:

Behavioural: Consistently repurchasing from a brand, and thereby driving up short-term sales and gross margin, after allowing for cannibalisation (Oliver, 1999).

Conative: Identifying as a customer and so generating higher cross-sales propensity and greater intention to purchase new products (Gremler et al., 2020).

Attitudinal: Liking a brand, which then generates wider value, such as brand trust and positive word of mouth (Wirtz et al., 2007).




Figure 2: Post-pandemic market share trends. Source: Market share data are Kantar Worldpanel's for 12 weeks ending (Kantar, 2025). Growth rates use a linear trend whereby the gradient is divided by the intercept to estimate the annual growth rate during this post-pandemic period.

Nonetheless, and despite loyalty programmes having been in their current form for decades across the world, their effectiveness is still being debated. Some firms see them as transformational, while others have had poor experiences. After all, a firm's growth is due to many factors, and it can be difficult to say what role a loyalty scheme plays. As a result, there's a parallel debate in the academic literature. And so it was that in 2021, Professor Alex Belli, now at the University of Melbourne, undertook a meta-analysis of 110 papers spanning 30 years, measuring 429 effect sizes to extract the main messages therein (Belli et al., 2022).

Unsurprisingly, given publication bias (which, in fairness, he and his co-authors were explicitly conscious of and attempted to minimise through searching for unpublished papers), Belli found strong evidence that loyalty programmes increase customer loyalty. More interestingly, they also identified how those effect sizes vary with by scheme design. These results are illustrated in Figure 3. The most impactful feature is whether the scheme is invitation-only and, relatedly, if it includes access to exclusive events. Such exclusivity may result in consumers valuing the scheme and events more due to mimetic dominance (Imas & Madarász, 2020), whereby accessing

something others cannot access allows for some form of social status (Henderson et al., 2011). However, Figure 3 also reveals that it is better to offer rewards that use one's own products (direct); furthermore, people don't like attention, because despite perhaps being expected to increase a customer's sense of belonging (Yim et al., 2008), it may also simply result in excessive personalised communications (Kim et al., 2021). So, hold the anniversary gifts, especially products you don't typically offer. As we did not test invitation-only schemes, we therefore expected the following:

- *H1. Features promoting exclusivity through an early product access offering will be the most effective in increasing customer loyalty.*
- *H2. Features promoting special attention, such as through birthday rewards, will reduce customer loyalty.*

Belli's study also suggests that despite all programmes having monetary impacts, discounts and joining fees may not actually influence customer loyalty, while savings through loyalty points do. This may be due to discounts creating a focus on costs which, while perhaps effective in isolation, will also increase cross-price elasticities, thereby pushing customers to search for better discounts




Figure 3: Loyalty scheme success drivers. Source: Meta research by Alex Belli et al. whilst at The University of Adelaide, based on 110 papers. Performance is measured using the Pearson Correlation of the published effect sizes, typically the difference on some outcome of interest (e.g. annual sales, product holdings or brand liking) between loyalty programme members and a non-member control group. Filled bars are statistically significant at $p = 0.01$.

elsewhere (Henderson et al., 2011). Points, meanwhile, may help build loyalty through increasing switching costs and associated loss aversion (Zhang et al., 2000), through providing enjoyment via collecting points in anticipation of rewards (Taylor & Neslin, 2005), due to a sense of accomplishment (Leenheer et al., 2007), or due to valuing points differently than money (Stourm et al., 2015). Finally, while customers obviously dislike costs, joining fees introduce switching costs to those who have joined, through sunk costs, that may counterbalance this regarding loyalty (Bhattacharya et al., 1995). Naturally, such fees would also be expected to deter some from joining in the first place. We therefore expect that:

- *H3. Discounts provided by a loyalty scheme will not increase customer loyalty.*
- *H4. Savings through loyalty points will increase customer loyalty.*
- *H5. The inclusion of an enrolment fee will not reduce customer loyalty.*
- *H6. The inclusion of a small enrolment fee (£1 per month) will reduce sign-ups less than larger enrolment fees (£3 or £5 per month).*

Belli and his co-authors also found that loyalty schemes have stronger effects in industries associated with lower purchase frequencies, as less frequent but

higher-value purchases are more likely to require trust. Thus, we also expect that:

- *H7. Loyalty schemes will promote more trust in retail stores with lower purchase frequency (electronics) than in retail stores with higher purchase frequency (clothing and groceries).*

Finally, while not something we examined, it is worth noting that loyalty programmes also confer indirect benefits, because the exhaustive data from a loyalty programme also allow companies to learn about their customers. Dunnhumby, the firm Tesco purchased to monetise Clubcard data, now generates £326m in sales. However, and more importantly, it also allows Tesco to market more effectively and tailor the customer shopping experience. Moreover, unlike in-store promotions, these personalised offers are impossible for competitors to monitor and counteract.

Acquiring Members

Having determined the business case and operational objectives, you then have to design the scheme. As noted, this means solving two problems – sign-up and participation. This section addresses the former. To explore this issue, we ran an experiment and assessed how to optimise sign-up. The experiment involved asking paid participants to undertake a




Figure 4: Acquisition impact. Source: Dectech fieldwork, May 2024 (n = 2,115 UK loyalty scheme members). The acquisition impact of features is based on exponentiated betas from a logistic regression model predicting overall sign-ups to the loyalty scheme in our experiment.

lifelike process, in this case a loyalty programme sign-up, in a randomised controlled trial in which scheme cost, sector, and features varied (see Appendix).

By far the biggest sign-up driver was cost. Unlike expected, participants were not particularly sensitive to the size of the monthly subscription fee (i.e. £1, £3, or £5), just whether there was one. Having a fee reduced the sign-up rate by a full two-thirds. It turns out that the zero price effect is very real, and people will value something free beyond what standard economic theory may predict (Shampanier et al., 2007). So, if you can afford to make your scheme free, you will get three times as many members. Beyond this it didn't matter how many features were foregrounded in the sign-up process, and certainly more isn't better. The key driver

is which features are foregrounded.

Overall, we tested 22 scheme features. The five features with the best and worst acquisition effects are shown in Figure 4. Customers are most attracted by free shipping, which will x1.27 your sign-up rate, again highlighting the power of free. Similarly, a £10 incentive payment, or advanced notice of product launches, boost acquisition by around x1.15. You need a hook, like Amazon Prime being famous for free delivery. Conversely, offering a monthly prize draw, access to content, or advanced notice of promotions (as opposed to being offered a personalised promotion) reduces sign-up propensity, possibly due to some customers perceiving these as bothersome (Kim et al., 2021).

Table 1: Sector Variation

	Overall	Grocery	Clothing	Electronics
Free Annual Delivery	x1.27	x0.89	x1.93	x1.33
Charity Donations	x1.09	x0.87	x1.40	x1.11
Personalised Offers	x0.96	x0.77	x0.62	x1.82
Promotion Notifications	x0.91	x0.72	x1.61	x0.78
Occasional Purchase Rewards	x0.82	x1.19	x0.94	x0.84

Source: Dectech fieldwork, May 2024 (n = 2,115 UK loyalty scheme members). The acquisition impact of features is based on exponentiated betas from a logistic regression model predicting overall sign-ups to the loyalty scheme in our experiment, as well as separate logistic regression models for the three sectors we tested.

The three best features will increase your acquisition rate x2.7 relative to the three worst. Over five years, using the +2.1% base rate, a well-designed scheme will attract 30% more customers – the difference between stalling at 20% post-launch or extending your reach to 50% of the base. But don't just ape Figure 4. Table 1 illustrates how the highest-impact features vary by category. Personalised offers work best in higher price point markets such as electronics. Free shipping is best for clothing and electronics, which are three times more likely to be purchased online compared to grocery. Test your design so that it is tailored to your sector.

Driving Sales

The second question revolves around how to activate members once they've signed up. As

noted, the meaning of “activate” depends on your commercial objectives. Maybe you are trying to generate brand-liking, and activation means greater engagement. Maybe you want more cross-sales, and so activation is defined by direct marketing response. Maybe you want higher sales, in which case activation should be measured by footfall or repeat purchasing. Our research shows that loyalty programme members spend 4.6 times more than non-members, so despite some reverse causation caveats (higher spenders being more likely to join), there are gains to be made here.

Figure 5 assesses this question based on real-world purchasing behaviour. The vertical axis indicates how much a given loyalty programme feature is associated with increased purchasing. Where the scheme offers member prices and people are aware of this, they




Figure 5: Sales impact. Source: Dectech fieldwork, May 2024 (N = 2,115 UK loyalty scheme members). The horizontal axis is the percentage of respondents who thought a loyalty programme offered the feature shown. The vertical axis is the betas from a linear model predicting the number of purchases made per year across all programmes between those who believed the loyalty scheme they were asked about included a given feature and those who didn't think it included it.

make 7.6 more purchases per year. That's nearly triple the 4.7 purchases an average consumer makes across these brands. There are, of course, again reverse causation caveats, as higher-frequency purchasers with higher spend may be more likely to join these schemes. However, the scheme features' impacts from the randomised controlled trial on stated purchase propensity (conative loyalty) and perceptions (attitudinal loyalty) were similar to the behavioural loyalty measure in Figure 5. Meanwhile, other aspects of the loyalty scheme had less of an impact, as the size of the subscription fee in the experiment did not influence purchase intentions among those who signed up despite reducing perceptions when raised to £3, and the retail sector only influenced purchase intentions (being highest for the grocery store).

Crucially, the features in Figure 5 that drive participation differ from those that drive sign-up. For example, member stories and free competitions aren't good for acquisition, but they do indirectly increase sales by converting a loyalty programme into a community. Schemes are not just about discounts. Some people may also simply identify with the scheme, and so it may act as a signifier in a semiotic sense

(Henderson et al., 2011). Conversely, offering joining rewards and providing advanced notice of events or product launches do not help with activation. They are the incentives for sign-up, following which, post-acquisition, they have done their job and should be set aside.

Meanwhile, the horizontal axis in Figure 5 is the average awareness of each feature across schemes. A low score means it's either not offered or members aren't aware it's being offered. Loyalty points and member prices are commonly offered and widely expected. The features that are less common or underpromoted that drive higher purchasing are in the top-left corner. For example, relatively few schemes offer discounts or rewards for other brands, possibly because partner offers don't drive acquisition. Figure 5 demonstrates that this benefit can be used to drive participation and is currently either uncommon or not widely understood. This creates the possibility of using this feature to both differentiate and enhance your scheme.

Finally, given we know which features drive purchasing and which schemes are known by members to have those features, the overall effectiveness of




Figure 6: Top 10 programmes. Source: Dectech fieldwork, May 2024 (N = 2,115 UK loyalty scheme members). The estimated impact of the featured loyalty schemes is based solely on the modelled sales impact of the features multiplied by the proportion of customers who believe the loyalty schemes have specific features.

different loyalty programmes can be calculated. The winners of this activation design bake-off are in Figure 6. Nectar (Sainsbury's and Argos) wins, beating Tesco Clubcard on the basis of being better known for personalised offers and having brand partners. Similarly, Co-op is known for member prices, and Boots is strong on its loyalty points offering. If your scheme is not on this list, call us, and we can tell you why.

Recommendations

Our research details how loyalty programmes can provide extraordinary commercial uplift and wider intangible benefits for retailers, if well-designed. Part of the challenge is that they must solve two problems, namely acquisition and participation. Based on our findings, there are six main recommendations on how to construct a winning scheme:

Align to win: Loyalty programmes can deliver many benefits, from a larger share of wallet to brand advocacy, the value of which will vary by sector. Clearly, it is difficult to design a loyalty programme when commercial objectives are poorly articulated, inadequately understood, or not shared across the organisation.

Nail the numbers: These commercial objectives need to be measured, forecasted, and tracked using the relevant metrics. For example, our work suggests

a stretch target of 25% base sign-up at launch and +5% per year thereafter. Such forecasts also support iterating the design and costing the programme. What size is the optimal discount?

The power of free: The zero price effect is well-documented. People respond disproportionately to free. Don't charge for your loyalty programme unless it's the only way to make the numbers stack up. And in that scenario, consider making the scheme free but exclusive, contingent on some minimum spend.

Invest in acquisition: Your loyalty programme will need a sign-up incentive. The question is, which is the most cost-effective incentive? In our experiment, a £10 voucher was 1.15 times better than a £10 gift (unless it's a Rubik's Cube and it's 1981). Alternatively, consider advanced notice or exclusive access offers.

Classic for a reason: For activation, you need to go with the script that customers already know and can navigate – member pricing, personalised offers, and loyalty points. That said, your scheme also needs to be famous for something. Beyond these basics, you need to invent a social media-friendly feature.

Tailor, tailor, tailor: The optimal strategy varies by category. You have to test your design iteratively. This might be post-launch A-B tests or pre-launch lab trials, like those used herein. But know this: focus groups and attitudinal surveys aren't enough, as evidenced by the graveyard of failed schemes.

THE AUTHORS

Henry Stott is a co-founder and managing director of Decision Technology. He has a PhD in decision-making and is a Certified Member of the MRS. At Decision Technology, Henry helps businesses and policymakers understand and manage customer decision-making. He was previously a director of Oliver Wyman, a financial services strategy consultancy, where he helped found their highly respected Risk Practice, developing many of their proprietary techniques in credit risk measurement and business valuation.

Jerry Luukkonen is a principal at Dectech. He has a PhD in behavioural science from the University of Warwick, and his areas of expertise include designing, conducting and analysing surveys and experiments in order to deliver actionable insights to senior business stakeholders.

Miriam Makin is a senior associate at Dectech. She holds an MSc in Behavioural Science from The London School of Economics and Political Science (LSE). She is passionate about applying behavioural science in her work with financial service providers to improve customer outcomes. She was a finalist for the MRS Young Researcher of the Year in 2024.

REFERENCES

- Belli, A., O'Rourke, A.-M., Carrillat, F. A., Pupovac, L., Melnyk, V., & Napolova, E. (2021). 40 years of loyalty programs: How effective are they? Generalizations from a meta-analysis. *Journal of the Academy of Marketing Science*, 50(1), 147–173. <https://doi.org/10.1007/s11747-021-00804-z>
- Bhattacharya, C. B., Rao, H., & Glynn, M. A. (1995). Understanding the bond of identification: An investigation of its correlates among art museum members. *Journal of Marketing*, 59(4), 46–57. <https://doi.org/10.1177/002224299505900404>
- Gremler, D. D., Van Vaerenbergh, Y., Brüggén, E. C., & Gwinner, K. P. (2020). Understanding and managing customer relational benefits in services. *Journal of the Academy of Marketing Science*, 48(3), 565–583. <https://doi.org/10.1007/s11747-019-00701-6>
- Henderson, C. M., Beck, J. T., & Palmatier, R. W. (2011). Review of the theoretical underpinnings of loyalty programs. *Journal of Consumer Psychology*, 21(3), 256–276. <https://doi.org/10.1016/j.jcps.2011.02.007>
- Imas, A., & Madarász, K. (2020). *Mimetic dominance and the economics of exclusion: Private goods in public context*. (Working Paper No. #8435). <https://doi.org/10.2139/ssrn.3657971>
- Kantar. (2025). *Grocery Market Share - Kantar*. Kantarworldpanel.com. <https://www.kantarworldpanel.com/grocery-market-share/great-britain/range/02.02.20/26.01.25>
- Kim, J. J., Steinhoff, L., & Palmatier, R. W. (2021). An emerging theory of loyalty program dynamics. *Journal of the Academy of Marketing Science*, 49(1), 71–95. <https://doi.org/10.1007/s11747-020-00719-1>
- Leenheer, J., Van Heerde, H. J., Bijmolt, T. H. A., & Smidts, A. (2007). Do loyalty programs really enhance behavioral loyalty? An empirical analysis accounting for self-selecting members. *International Journal of Research in Marketing*, 24(1), 31–47. <https://doi.org/10.1016/j.ijres-mar.2006.10.005>
- Oliver, R. L. (1999). Whence consumer loyalty? *Journal of Marketing*, 63(4), 33–44. <https://doi.org/10.1177/00222429990634s105>
- Shampanier, K., Mazar, N., & Ariely, D. (2007). Zero as a special price: The true value of free products. *Marketing Science*, 26(6), 742–757. <https://doi.org/10.1287/mksc.1060.0254>
- Stourm, V., Bradlow, E. T., & Fader, P. S. (2015). Stockpiling points in linear loyalty programs. *Journal of Marketing Research*, 52(2), 253–267. <https://doi.org/10.1509/jmr.12.0354>
- Taylor, G. A., & Neslin, S. A. (2005). The current and future sales impact of a retail frequency reward program. *Journal of Retailing*, 81(4), 293–305. <https://doi.org/10.1016/j.jretai.2004.11.004>
- Wirtz, J., Mattila, A. S., & Lwin, M. O. (2007). How effective are loyalty reward programs in driving share of wallet? *Journal of Service Research*, 9(4), 327–334. <https://doi.org/10.1177/1094670506295853>
- Yim, C. K., Tse, D. K., & Chan, K. W. (2008). Strengthening customer loyalty through in-

timacy and passion: Roles of customer–firm affection and customer–staff relationships in services. *Journal of Marketing Research*, 45(6), 741–756. <https://doi.org/10.1509/jmkr.45.6.741>

Zhang, Z. J., Krishna, A., & Dhar, S. K. (2000). The optimal choice of promotional vehicles: Front-loaded or rear-loaded incentives? *Management Science*, 46(3), 348–362. <https://ssrn.com/abstract=2552242>

APPENDIX: METHODOLOGY

Sampling

The primary research undertaken for this report was conducted online in May 2024. Respondents were a nationally representative sample of 2,115 UK consumers aged 18 and over who were a member of at least one loyalty scheme and were responsible within their household for purchasing products for which they were later shown loyalty schemes for in the experiment (groceries, clothes, and/or consumer electronics).

Behaviourlab

Behaviourlab is our bespoke online test platform that uses a randomised controlled trial to address key commercial questions more accurately. The method follows modern academic standards of eliciting consumer preferences and forecasting their behaviours. This research involved putting participants through a realistic simulation of a loyalty scheme

sign-up journey for a fictional retail brand (see Figure 7 for an example). Each participant was asked to read the details of the loyalty scheme they were presented with and to decide whether to sign up or not. Participants were shown a loyalty scheme for one of three types of retail stores: a grocery store, a clothing store, or a consumer electronics store. To proceed, participants had to either complete the sign-up process or choose to not to do so by pressing the “I’m not interested” button on the loyalty scheme information screen. The loyalty scheme shown to participants was randomised in a number of ways (see Table 2). Specifically, in addition to randomising the store, the loyalty scheme would either include no sign-up fee (shown to half of respondents) or include a £1, £3, or £5 per month sign-up fee (each shown with equal probability to the other half of respondents). In addition, the loyalty scheme shown would include two, three or four features, each chosen randomly from a separate category (see Table 3).

Table 2: Summary of Experiment Conditions

	Element 1	Element 2	Element 3	Element 4
Retail Store	Groceries	Clothing	Consumer Electronics	
Sign-Up Fee	Free	£1 per month	£3 per month	£5 per month
Number of Features	2	3	4	




Figure 7: Example loyalty scheme sign-up journey.

Table 3: Features Included in the Loyalty Schemes

Feature Category	Features	Feature Title and Description Shown
Savings	Discounts	Discounts on all purchases: Enjoy 2% off all our products. On top of any promo sale.
	Personalised offers (selected by retailer)	Offers just for you: Enjoy 10% off three products we think you will love (max item value of £20). On top of any promo sale.
	Personalised offers (selected by customer)	Discounts on your favourites: Enjoy 10% off three products of your choice any time you buy them (max item value of £20). On top of any promo sale.
	Member prices	Member-only prices: Enjoy member-only prices on specific products.
	Loyalty points	Collect and spend points: Earn 1 loyalty point for every £1 you spend. Use your points for any purchase. 100 points are worth £2.
Joining Reward	£10 voucher at sign-up	Receive a free welcome gift: Enjoy a £10 K&H voucher from us when you join as a member.
	£10 gift at sign-up	Receive a free welcome gift: Choose a free welcome gift worth £10 from a range of products when you join.
Gifts	Occasional purchase rewards	Rewards on purchases: Occasionally, you'll receive special surprises & rewards when you purchase.
	Prize draw	Exclusive prize draws: Have a chance to win in our free monthly member-only prize draws.
	Birthday gift	Birthday gift: Receive a special gift from K&H on your birthday.
Referral Rewards	£10 on referral	Refer a friend for rewards: Get rewarded with a £10 K&H voucher when you get a family member or friend to join Super Perks.
	£10 to both on referral	Refer a friend for rewards: Get a family member or friend to join Super Perks, and you both get rewarded a £5 K&H voucher.
Charity Donation	Charity donation (chosen by customer)	Charity donations: Every time you shop, we'll donate 10p to a charity of your choice from K&H's charity partners.
	Charity donation (chosen by retailer)	Charity donations: Every time you shop, we'll donate 10p to a K&H chosen charity.
Free Benefit	Free hot drink vouchers	Free hot drink with purchase: Get a free tea or coffee at our stores after every purchase.
	Free annual delivery	Free annual delivery pass: Get free priority delivery and click + collect on all of your orders.

Feature Category	Features	Feature Title and Description Shown
Member Notifications and Content	Advanced notice of product launches	Receive timely notifications: Be first to hear about new product launches or special in-store events.
	Promotions notifications	Receive timely notifications: Be first to hear about promotions or special prices on our products.
	Exclusive content	See exclusive content: Access articles, blogs, and videos from K&H.
Partner Offers	Food box delivery	Exclusive partner offers: Get access to special offers and discounts from our food box delivery partner brands.
	Footwear	Exclusive partner offers: Get access to special offers and discounts from our footwear partner brands.
	Consumer electronics	Exclusive partner offers: Get access to special offers and discounts from our electronics partner brands.
Behavioural Rewards	Points for tasks	Engage and get rewarded: Collect points by visiting our website, leaving products reviews, and engaging with our posts on social media.

After the loyalty scheme sign-up journey, participants who had signed up indicated their likelihood of doing so in reality on an 11-point Likert Scale ranging from “Extremely Unlikely” to “Extremely Likely”. All participants were also asked to rate the likelihood of purchasing from the fictional brand on the same 11-point Likert Scale and to rate agreement with five perception statements regarding the loyalty scheme on a 7-point Likert Scale ranging from “Strongly Disagree” to “Strongly Agree”. The perception statements related to value for money, sufficient detail, clarity of the benefits, suitability for the person’s needs, and innovativeness of the benefits. Principal component analysis was conducted with these perception ratings to find a hybrid measure of Customer Satisfaction. This score was approximately normally distributed ($N(0,1)$), and it was used as an indication of satisfaction with the loyalty schemes.

We also asked participants a series of questions about two retail brands they had purchased a product from in the last 12 months – one they had a loyalty membership with and, where possible, one they did not have a loyalty membership with. A minority (5.6% of participants), who only purchased from brands they had signed up with, were asked about a second brand with which they had also signed up. Questions

about the brands included purchase frequency and types of items purchased in the last 12 months, spend and number of items purchased during the most recent purchase, and likelihood of recommending the brand to friends or colleagues. For brands whose loyalty schemes the participants had signed up to, we also included questions about whether they had used their loyalty card/app during the most recent purchase, what features they believed the loyalty scheme included, and whether a series of events related to the loyalty scheme had occurred in the last 12 months (e.g. receiving a discount, attending a brand event, or getting an unexpected reward). For brands participants had not signed up with, we also included a question about why they had not done so.

Modelling

The analysis involved statistically modelling whether the type of retail store, sign-up fee amount, number of features, and specific features included affected the sign-up rate (binary logistic regression model), sign-up likelihood (ordinal regression model), purchase likelihood from the fictional brand (ordinal regression model), and perceptions of the loyalty scheme (linear model) shown in the experiment, although we chose to focus on sign-ups and purchases

in this report.

We also statistically modelled whether respondents had signed up to the loyalty schemes of the real-world brands they had been asked about (binary logistic regression model), their likelihood of recommending the brands (ordinal regression model), number of purchases in the last 12 months from the brands (linear regression model), and the perceptions of the real-world loyalty schemes (linear regression model). As we did not ask participants directly about their perceptions of the real-world loyalty schemes, we instead applied the real-world loyalty scheme data to a version of the model we used to predict the effect of loyalty scheme features in the experiment on

perceptions. Specifically, where possible, we matched the real-world loyalty scheme features participants believed their loyalty scheme had with those tested in the experiment, assuming effect sizes would be similar on perceptions.

The purpose of modelling is in part to control for the impact of other information (such as a consumer's age) and thereby isolate and estimate the impact of different benefits on the dependent variables. The set of controlling factors included existing loyalty scheme usage, current spend, current purchase frequency, and demographics. Modelling also allows us to identify any statistically significant effects and avoid reporting insights that are simply noise.

Can a Simple Message Make Customers More Honest? Insights From the Insurance Industry

GERHARD FEHR¹ AND LUCAS AMHERD¹

Fraudulent insurance claims represent a significant financial burden for insurers, and they are often influenced by how individuals rationalize their reporting decisions rather than being a clear intent to deceive. This study examines how behavioral science-based communication strategies can encourage truthful claims reporting through subtle, cost-effective interventions. In two countries, a randomized controlled trial tested four variations of Interactive Voice Response (IVR) messages, each leveraging different behavioral mechanisms, such as emphasizing social expectations, reinforcing moral responsibility, or fostering trust between customers and the insurer. The findings indicate that messages promoting a sense of shared ethical responsibility were the most effective in reducing dishonest reporting, while those that focused on monitoring and deterrence had mixed or no effects, sometimes leading to increased dishonesty. These insights offer insurers a scalable, customer-friendly strategy to enhance honesty in claims reporting while maintaining trust and engagement.

Introduction

The Insurance Group, one of Europe's leading insurers, manages a premium volume of several billion euros and serves individuals and businesses across the continent. Insurance fraud remains a significant financial challenge, with fraudulent claims accounting for up to 10% of total payouts in Switzerland, for example (SVV, 2017). Public perception further complicates this issue; for instance, nearly a quarter of Swiss consumers consider claim exaggeration as a socially acceptable act rather than as fraud. Conventional fraud prevention strategies often rely on detection and deterrence, but such measures can be costly, difficult to enforce, and risk damaging customer relationships. This raises a key question: Can behavioral science be used to promote honesty in claims reporting in a way that is both effective and customer-friendly?

Research in behavioral economics suggests that dishonesty is rarely the result of outright fraud but is often driven by psychological justifications. Customers may perceive the insurer as a distant entity rather than individuals being affected, or assume that others exaggerate claims as well. Understanding these behavioral drivers is critical to designing effective interventions that encourage honesty without harming customer trust.

Background

Understanding Honesty: Strategic Importance for the Insurance

Honesty, in our context, can be defined as the commitment to truthfulness, particularly in situations where there is an incentive to deceive. In terms of the insurance, promoting honesty among customers is crucial for reducing fraudulent claims, which directly impact both financial outcomes and customer trust. In insurance settings, dishonesty—such as exaggerating damages or misrepresenting details—poses a significant risk and leads to financial losses, increased premiums, and a damaged reputation. Therefore, the insurance seeks to mitigate these risks by optimizing customer communication strategies that promote honest behavior during claims reporting. The challenge lies in influencing customer behavior at critical touchpoints, without creating a negative customer experience.

Honesty can be broadly categorized into two types:

1. **Stated honesty:** Where individuals report their intentions or beliefs, which may be influenced by social desirability or perceived expectations.
2. **Revealed honesty:** Where honesty is inferred from actual behavior in decision-making

¹ Corresponding authors: gerhard.fehr@fehradvice.com and lucas.amherd@fehradvice.com

scenarios where there are tangible incentives to be dishonest.

Research indicates that dishonesty is not purely a cost-benefit decision; for example, people deceive only to the extent that they maintain a positive self-image (Mazar et al., 2008). This suggests that subtle behavioral interventions could enhance honesty by making ethical behavior more salient.

Identifying Intervention Points

To reduce dishonesty in claims reporting, FehrAdvice & Partners and the insurance team identified key moments where behavioral interventions could be most effective. Customers have systematic incentives to misreport claims, creating a Principal-Agent problem whereby financial motives often override honesty. The focus was placed on developing behavioral nudges that could be integrated at critical decision points to reinforce ethical behavior before claims were submitted. The Interactive Voice Response (IVR) system was ultimately selected for experimental testing due to its scalability, cost-effectiveness, and ability to influence customer behavior at a crucial pre-decision moment in the claims process.

Behavioral Mechanisms Behind IVR Treatments

The behavioral mechanisms tested in this study were co-developed with the insurance company, thereby ensuring their relevance to real-world claims reporting. Rather than applying generic behavioral principles, behavioral economics tools were used to identify psychological drivers most likely to influence honesty in this specific context. Each IVR treatment represents a valid behavioral hypothesis, grounded in existing research but requiring empirical validation within this insurer's customer base. While concepts like social norms, reciprocity, and identity priming have been effective in other settings, their impact is context-dependent. This is why experimentation is essential—behavioral science is not about applying standardized nudges but about measuring their effectiveness in practice. The following section details select behavioral mechanisms integrated into the IVR treatments.

1. Reciprocity – honesty as a mutual expectation: Customers tend to reciprocate trust. When the IVR message expresses trust in their honesty, they feel obliged to act ethically. Studies suggest that when

people feel trusted, they are more likely to uphold that trust (Fehr & Gächter, 2000).

2. Conscientiousness – appealing to moral responsibility: People have an internal drive to act according to their values. Reminding customers that their honesty is important reinforces their moral obligation to report truthfully.

3. Fairness – emphasizing just outcomes: Highlighting the fairness aspect appeals to the customer's intrinsic value thereof. When treated fairly, reciprocity encourages people to cooperate and to punish unfair treatment, thus sustaining cooperative behavior in certain economic interactions (Fehr & Gächter, 2000). When fairness is emphasized, they are more likely to behave fairly themselves, aligning with the insurance's goal of ensuring equitable treatment for all policyholders.

4. Social comparison – “others are honest, so you should be too”: Research has shown that people tend to follow what they believe others are doing (descriptive norm). When this matches what is socially approved (injunctive norm), it strengthens the message and encourages positive behavior change (Cialdini, 2003). In the Community treatment, social comparison is used to suggest that being part of the insurance means being part of a community of honest customers. This plays on the psychological tendency of individuals to conform to group norms and expectations and as a result promotes honest behavior. This approach is further supported by the finding that people imitate observed behavior, and when they see others acting ethically, they are more likely to follow suit (Keizer et al., 2008).

5. Moral appeal and inequity aversion – avoiding the discomfort of gaining unfairly: Treatments leveraging moral appeal and inequity aversion emphasize the moral high ground of honesty and the discomfort of being perceived as dishonest compared to others. In the context of the insurance, this means that customers who consider themselves part of a fair community may experience cognitive dissonance if they act dishonestly. This aligns with self-concept maintenance theory, which suggests that people balance their desire for financial gain with their need to maintain a positive self-view, often cheating only to the extent that allows them to avoid updating their moral self-concept (Mazar et al., 2008). Research by Gneezy (2005) supports this notion by showing that

people are less likely to lie when their dishonesty results in greater harm to others, suggesting that emphasizing the consequences of dishonest claims can reduce fraudulent behavior. Furthermore, inequity aversion suggests that customers may avoid dishonest behavior to avoid feeling that they are gaining an unfair advantage over others who are honest, therefore aligning their actions with social norms and fairness (Fehr & Schmidt, 1999).

6. Feedback and cooperation – creating a trust-based relationship: Messages that express appreciation for direct contact and honesty create a feedback loop that encourages cooperation. Research shows that when people perceive a trusting relationship, they are more likely to act in pro-social ways, reciprocating trust with increased effort and cooperation, such as being honest (Falk & Kosfeld, 2006).

7. Identity priming – aligning actions with self-perception: Messages that tap into the customer's identity are designed to align their actions with their self-concept as a trustworthy and honest individual. This is founded in identity economics research, in that when people see themselves as part of a group with strong ethical standards, they are more likely to act in accordance with those standards (Akerlof & Kranton, 2000).

The insights gained from this experiment are critical for the insurance, as they decide on the best communication strategies to reduce fraudulent claims and strengthen customer trust. By understanding the specific behavioral drivers that each message targets, the insurance can craft more effective and tailored customer communication strategies.

Designing Effective IVR Messages to Encourage Honesty

Multiple IVR message prototypes were designed to encourage honest behavior in claims reporting. From these, four were selected for testing in an online experiment to evaluate their effectiveness. Each treatment was crafted based on specific psychological principles rooted in behavioral economics, in order to influence customer decisions subtly. Table 1 is an overview of each treatment, along with an explanation of the behavioral mechanisms driving their design. Due to proprietary agreements with the insurer, the exact IVR message content cannot be disclosed.

For all treatments, a female human voice delivered the IVR message, maintaining consistency across the interventions. After the message was delivered, participants had to click a button to proceed to the honesty decision point, thus ensuring that the method of presentation did not differ between treatments. The control group followed the same protocol, albeit with the insurer's original message, hence maintaining comparability while preserving consistency in message delivery.

Measuring Honesty: An Experimental Approach

This project employed a randomized controlled trial to examine how IVR message framing influences honesty in insurance claims reporting. Participants were placed in a realistic decision-making scenario, where they simulated purchasing a car, managing a budget, and filing an insurance claim after an accident. Before reporting the damage, they reviewed their insurance policy terms, ensuring they were aware of the financial implications of their decision. The protocol was exactly the same in both online experiments in Italy and Switzerland, ensuring comparability of results across both countries.

To test honesty, participants called a fictitious customer service hotline and were randomly assigned to hear one of four behaviorally designed IVR messages or a control message. They then chose between an honest claim (paying a CHF 1,000 deductible fee) or a dishonest claim (avoiding the deductible fee by falsely blaming another driver). Their choice served as a revealed honesty measure and helped assess whether financial incentives outweighed ethical considerations. Participants also rated the IVR messages on clarity, transparency, and intrusiveness, providing additional insights into message effectiveness.

The experiment was not explicitly framed around honesty, so participants were unaware that honesty was relevant to their responses or that their answers could be verified. This design choice mirrors real-world conditions in which individuals may not be fully aware of the extent to which their claims are scrutinized. However, it is worth noting that actual insurance fraud scenarios often involve uncertainty about claim verification, which may influence perceived risks and behavior.

Table 1: Tested IVR Treatments

IVR Treatment	Behavioral Mechanisms	Message Explanation	Example Phrases	Message Duration
1. Community	Reciprocity, Conscientiousness, Fairness, Social Comparison, Moral Appeal, Inequity Aversion	Reinforces a sense of belonging and shared responsibility among policyholders. Emphasizes mutual trust and fairness, positioning honesty as a collective norm.	“The home for a million customers”, “We rely on your honest information”	10 seconds
2. Evidence	Reciprocity, Conscientiousness, Fairness, Social Comparison, Moral Appeal, Inequity Aversion	Builds on the Community treatment by adding an element of transparency and accountability. Highlights the importance of accurate claims reporting, stronger focus on monitoring.	“Recorded for evidence purposes”	12 seconds
3. Appreciation	Reciprocity, Feedback, Cooperation, Identity	Acknowledges and values direct customer engagement. Establishes a positive relational dynamic, emphasizing trust between the insurer and the policyholder.	“Appreciate that you are contacting us”, “We are here for you”	8 seconds
4. Gratitude	Reciprocity, Feedback, Cooperation, Identity	Expresses appreciation for customer honesty and trust, fostering a cooperative relationship. Frames honesty as an integral part of receiving good service.	“Thank you for your honest information”	6 seconds
5. Control Group (No Treatment)	None	No behavioral nudge applied. Serves as a baseline to measure the impact of IVR interventions.	None	2 seconds

Note: This overview summarizes the tested IVR treatments, their underlying behavioral mechanisms, intended message effects, and message duration on audio. The control group received no intervention, serving instead as a baseline.

Bias Mitigation Strategies in the Experiment Design

To ensure the validity of the findings, the experiment incorporated multiple bias mitigation strategies in order to minimize distortions in measured honesty rates and ensure external validity.

Ensuring causal attribution: A randomized controlled design ensured that participants were equally likely

to be exposed to any IVR treatment or the control condition. This minimized selection bias, allowing observed differences in honesty rates to be attributed directly to IVR messaging, rather than individual characteristics or pre-existing tendencies.

Social desirability bias: Participants were not informed that honesty was being measured. Instead, they believed they were participating in a general

study on insurance experiences. This strategy prevented them from modifying their behavior to appear more honest, ensured genuine decision-making in the claims reporting scenario, and aligned with research indicating that concealing the true purpose of an experiment reduces self-presentation bias (Goffman, 1956). By removing explicit honesty cues, the experiment captured real-world decision-making dynamics.

Incentivizing realistic behavior: The experiment was designed to ensure incentive compatibility, replicating real-world conditions where financial gain can motivate dishonesty. Participants were given a fictitious CHF 15,000 budget, which they managed throughout the study. When reporting an accident, they faced a realistic financial trade-off: choosing the dishonest claim (no deductible) resulted in a monetary benefit, while the honest claim (CHF 1,000 deductible) led to a financial loss.

To reinforce meaningful decision-making, participants could use any remaining fictitious money to purchase raffle tickets for a real prize, thus making their financial choices tangible rather than hypothetical. This structure ensured that observed dishonesty reflected real economic incentives, mirroring the pressures policyholders face when filing insurance claims.

Enhancing realism through contextual immersion:

The experiment was designed using realistic decision-making contexts that closely mirror the actual experiences of insurance customers, enhancing external validity. Through a storytelling approach, the experiment immersed participants in scenarios that felt genuine, such as buying a car and dealing with an accidental scratch. This method made the decision points more relatable and authentic, and it encouraged participants to respond as they would in real life, thereby enhancing the reliability of the data on honest and dishonest behavior.

All these strategies were critical in isolating the causal effects of different IVR messages, providing the insurance with robust, actionable insights for optimizing communication strategies, in order to reduce fraudulent claims and build customer trust.

Switzerland: Results of the Online Experiment

The experiment in Switzerland involved 2,987 insurance customers. The majority of participants were over 50 years old (1,516; 51%), with 1,173 (39%) aged 30 to 49, and 298 (10%) under 30. The sample was predominantly male, with 1,987 men (67%) and 1,000 women (33%). While none of the differences between treatments and the control group was statistically significant at standard levels using t-tests,




Figure 1: The measured proportion of dishonest responses refers to the number of dishonest claims out of the total responses in the online experiment conducted in Switzerland, using a fictitious comprehensive insurance case after exposure to the respective IVR message. Direct measurement of the decision immediately after listening to the IVR message and under real monetary incentives.

comparisons between the treatments themselves reveal meaningful differences, as shown in Figure 1.

The experiment uncovered differences in how IVR messages affected dishonesty rates among insurance customers in Switzerland. The Community treatment, stressing social norms and collective responsibility, reduced dishonesty the most to 7.2% of dishonest responses. By framing honesty as the expected norm among responsible policyholders, this intervention reframed dishonesty as a violation of trust within a shared community. Furthermore, emphasizing that fraudulent claims harm fellow policyholders rather than a distant corporate entity reinforced social accountability. Its statistically significant superiority over Evidence ($p = 0.097$) and Appreciation ($p = 0.019$) suggests community-oriented messages outperform monitoring or mere acknowledgment.

In contrast, the Evidence treatment increased dishonesty to 10.1%. This likely triggered psychological reactance, whereby participants responded negatively to perceived oversight. The Gratitude treatment, which thanked customers for their trust and honesty, reduced dishonesty to 7.9%. Positive reinforcement likely made the customers feel valued, in turn enhancing intrinsic motivation to act honestly. Compared to the control group (9.0%), both the Community and Gratitude treatments, although not statistically significant, reduced dishonesty without explicit deterrents. The difference between Gratitude and Appreciation ($p = 0.060$), however, indicates that expressing gratitude is more effective than simple acknowledgment.

The Appreciation treatment, i.e., acknowledging customers for contacting the insurance company, resulted in the highest dishonesty rate of 11.3%. While the hypothesis was that expressions of appreciation can strengthen social bonds and encourage cooperation, the findings in this context indicate that in isolation they are not enough to promote honesty. This counterintuitive finding suggests that generic expressions of appreciation, without linking to ethical accountability, may lower customers' perceived responsibility to report truthfully.

Beyond the content of the messages, the experiment determined that the perceived clarity and tone of the IVR message also influenced behavior. Messages that were rated as clear and transparent were more effective in reducing dishonesty, while

those perceived as intrusive were associated with higher dishonesty rates (both statistically significant with $p < 0.01$). This highlights the importance of crafting IVR messages that strike a balance between nudging ethical behavior and maintaining a positive customer experience.

FehrAdvice & Partners recommended that the insurance should implement the Community IVR message for broader rollout, based on the results of the experiment and the following insights:

- First, the concept of priming—small changes to the message's introduction—can significantly influence behavior, triggering honest actions by reminding customers of their membership of a trustworthy group.
- The focus on social identity and social proof, where customers are reminded that they are part of a community that values honesty, was particularly effective in promoting ethical behavior.
- Finally, they advised the insurance to ensure that the messages were clear and not perceived as intrusive, as overly confrontational or accusatory tones may backfire and reduce honesty.

Italy: Scaling the Experiment to a Larger Market

A second experiment in Italy examined whether behaviorally designed IVR messages could reduce dishonest claims in a market with lower baseline honesty rates than Switzerland. Prior research highlights significant national differences in ethical behavior, with Switzerland demonstrating higher civic honesty (Cohn et al., 2019). This study provided an opportunity to test whether similar IVR interventions could influence honesty in a setting where fraudulent behavior is more prevalent. The results confirmed this discrepancy, with 31% of participants in Italy misreporting their claims, consistent with previous findings on lower honesty norms. The experiment aimed to identify messaging strategies that could effectively counteract this tendency and encourage more truthful reporting.

The experiment in Italy involved 7,041 insurance customers, with 45% female and 55% male participants. The age distribution included 1,098 participants below 35 years (16%), 3,308 between 35 and 54 years (47%), and 2,635 over 55 years (37%).




Figure 2: Measured proportion of dishonest responses in the online experiment in Italy, using a fictitious comprehensive insurance case after exposure to the respective IVR message. Direct measurement of the decision immediately after listening to the IVR message and under real monetary incentives.

The experiment revealed differences in relation to the Swiss findings, reflecting broader differences in civic honesty, as illustrated in Figure 2. While dishonesty in the Italian control group was higher at 32.2%, the most effective intervention followed the same pattern: framing honesty as a collective norm. The Community treatment reduced dishonesty to 28.4%, confirming that reinforcing social responsibility is effective across contexts, albeit with a smaller relative impact than in Switzerland. Its effectiveness over the control group was also statistically significant ($p = 0.038$) when compared to Appreciation ($p = 0.060$) and Gratitude ($p = 0.067$).

The Evidence treatment had a statistically insignificant, dishonesty-lowering effect, whereas similar monitoring-based messaging in Switzerland led to a backlash. This suggests that oversight cues can sometimes deter dishonesty, but their effectiveness varies by cultural context. The Appreciation and Gratitude treatments also failed to influence behavior in either country, resulting in dishonesty levels similar to the control group. None of these effects was statistically significant, which reinforces the idea that generic gratitude, without direct linkage to ethical accountability, does not effectively enhance honest behavior.

The results demonstrate that behavioral nudges can still drive meaningful reductions in fraudulent claims. Moreover, these findings highlight the need for context-sensitive interventions, ensuring messages align with local honesty norms to maximize their impact.

Key Takeaways

Even small behavioral adjustments can lead to significant changes in customer decision-making. This study highlights how subtle shifts in communication—such as framing honesty as a shared responsibility—can meaningfully reduce fraudulent claims, without harming customer relationships. The success of these interventions not only underscores the power of behavioral economics tools as a foundation for designing effective measures, but also reinforces a crucial principle: what works depends on the context.

Developing the right interventions requires a structured, collaborative approach. The hypotheses tested in this study were not arbitrary but were developed jointly using behavioral science frameworks to identify the psychological mechanisms most likely to influence honesty. However, having a strong hypothesis is not enough—only through

rigorous experimentation can we determine what actually works in a specific environment.

This is why principled experimental design is essential. Ensuring incentive compatibility and testing interventions in realistic decision-making scenarios allows us to measure actual behavior rather than relying on assumptions. Behavioral economics is more than just applying social proof or nudging customers in a generic way—it requires precise, context-specific testing to create interventions that truly drive impact.

The findings from this study not only provide actionable insights, but also emphasize the importance of continuous learning. What worked in this experiment offers a strong foundation for broader implementation, yet further experimentation across different settings, customer segments, and communication channels is key to refining and optimizing behavioral interventions in the insurance industry.

THE AUTHORS

Gerhard Fehr is an applied behavioral economist, Behavioral Designer, and Delegate of the Board of Directors at FehrAdvice & Partners, which he co-founded with Ernst Fehr to bring behavioral economics insights into practice for businesses and policymakers. Renowned for his dedication to #Experimentability as a transformative management skill, Gerhard led FehrAdvice to become the leading behavioral economics consultancy in the German-speaking world. Gerhard completed his Executive Education in Behavioral Economics at Harvard Business School, enhancing his expertise in designing evidence-based solutions. His mission is to establish FehrAdvice & Partners as the European market leader in digitalization and #Experimentability, promoting an experimental culture where quick, small-scale failures accelerate learning and innovation.

Lucas Amherd is a consultant and member of the Behavioral Insights Team at FehrAdvice & Partners, contributing to cutting-edge projects by applying behavioral economics principles to business and policy challenges. He has a Master's in Economics at the University of Zurich, having studied at Sciences Po and the University of Sydney. His expertise spans behavioral sciences, data science, causal inference, and business development. Passionate about innovation

and evidence-based solutions, Lucas continually seeks to expand his knowledge and skills to drive impactful change through behavioral insights.

REFERENCES

- Akerlof, G. A., & Kranton, R. E. (2000). Economics and identity. *The Quarterly Journal of Economics*, 115(3), 715–753. <https://doi.org/10.1162/003355300554881>
- Brehm, J. W. (1966). *A theory of psychological reactance*. Academic Press.
- Cialdini, R. B. (2003). Crafting normative messages to protect the environment. *Current Directions in Psychological Science*, 12(4), 105–109. <https://doi.org/10.1111/1467-8721.01242>
- Cohn, A., Maréchal, M. A., Tannenbaum, D., & Zünd, C. L. (2019). Civic honesty around the globe. *Science*, 365(6448), 70–73. <https://doi.org/10.1126/science.aau8712>
- Falk, A., & Kosfeld, M. (2006). The hidden costs of control. *American Economic Review*, 96(5), 1611–1630. <https://doi.org/10.1257/aer.96.5.1611>
- Fehr, E., & Gächter, S. (2000). Fairness and retaliation: The economics of reciprocity. *Journal of Economic Perspectives*, 14(3), 159–182. <https://doi.org/10.1257/jep.14.3.159>
- Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3), 817–868. <https://doi.org/10.1162/003355399556151>
- Gneezy, U. (2005). Deception: The role of consequences. *American Economic Review*, 95(1), 384–394. <https://doi.org/10.1257/0002828053828662>
- Goffman, E. (2023). *The presentation of self in everyday life*. In *Social theory re-wired* (pp. 450–459). Routledge.
- Keizer, K., Lindenberg, S., & Steg, L. (2008). The spreading of disorder. *Science*, 322(5908), 1681–1685. <https://doi.org/10.1126/science.1161405>
- Mazar, N., Amir, O., & Ariely, D. (2008). The dishonesty of honest people: A theory of self-concept maintenance. *Journal of Marketing Research*, 45(6), 633–644. <https://doi.org/10.1509/jmkr.45.6.633>
- Schweizerischer Versicherungsverband. (2017). *Versicherungsbetrug: Zahlen und fakten* (Studie

der GfK im Auftrag des SVV). Schweizerischer
Versicherungsverband (SVV). [https://www.
svv.ch/sites/default/files/2017-11/SVV%20](https://www.svv.ch/sites/default/files/2017-11/SVV%20)

Zusammenfassung%20der%20Ergebnisse%20
der%20GfK%20Studie%20zum%20
Versicherungsmissbrauch%202017.pdf

What Lies Behind Our Behavior? An Approach to Understanding Human Decision-Making

OLYMPIA ARANGO CASTRO, MIGUEL ANTONIO RODRÍGUEZ RAMÍREZ,
JUAN DE RUS GUTIÉRREZ AND EMMA BERNARDO SAMPEDRO¹

Neovantas

This article presents a behavioral analysis framework that moves beyond traditional data-driven approaches by focusing on psychological mechanisms—specifically, motivation and affect—as key drivers of decision-making. Drawing on both observational and experimental data, we illustrate how these latent cognitive factors can be identified and modeled through linguistic cues in natural language. We examine two empirical cases: one explores intrinsic and extrinsic motivations among Airbnb hosts, using topic modeling and multi-label classification; the other integrates the COM-B behavioral model into large language models (LLMs) to distinguish behavioral inaction sources in digital health coaching. Additionally, we propose a practical application using sentiment analysis and SHAP (SHapley Additive exPlanations) to infer motivational signals in real-world language data. By combining behavioral theory with computational tools, this approach facilitates a more nuanced interpretation of human behavior and sets the foundation for future research in AI-powered intention modeling. Ethical considerations and methodological limitations are also discussed.

The Use of Data in Economics

Economics and other disciplines have advanced in the empirical study of social sciences (Einav & Levin, 2014), including researchers' use of databases combined with statistical tools to investigate social and individual phenomena. The type of data used by scientists is divided into two categories: observational and research (Carthey, 2003; Brosnan, 2008; Baillie & Higgins, 2024).

Observational data refers to data that occurs without human intervention, where causal relationships are hard to identify. One example is medical data, in terms of, say, how many patients are diagnosed with an illness (Carthey, 2003). Another example is customer data, such as types and frequency of purchase, or budget spent by a group of consumers. This data is collected in a real-world environment, meaning that there are no experimental conditions or treatments applied to the subjects providing it (Baillie & Higgins, 2024).

Research data refers to information collected through various scientific methods, including experimentation, meaning that it is produced, analyzed, and

shared by experts using scientific tools (Gomez-Diaz & Recio, 2022). This type of data is generated to answer a very specific question and may be collected in a controlled environment, such as a laboratory, or online, and it can be used to quantify a specific impact or effect.

Both data categories can help us study and describe human behavior (Altmann, 1974; Paolisso & Hames, 2010), but economics has traditionally relied on models centered on final decisions to explain behavior (Radu & Radu, 2014). For this reason, the development of behavioral economics has helped to also incorporate psychological factors such as intentions, emotions, and cognitive processes in the study of human decision-making (Albarracín et al., 2018).

Motivations are defined by the American Psychological Association as 'the impetus that gives purpose or direction to behavior and operates in humans at a conscious or unconscious level' (APA, 2018). Apart from being an important influence on decision-making, they are also conditioned by emotions (Ng, 2023); for example, happiness can

¹ Corresponding author: ebernardo@neovantas.com

boost the motivation to pursue a behavior, and fear can disincentivize it. However, several authors have highlighted the difficulties and subjectivity in motivation research (Girard & Cohn, 2016; Souza et al., 2021; Beauchamp, 2024).

The aim of this article is to move beyond traditional data-driven approaches by illustrating the analysis of emotions or affective states as a means to uncover motivations in human decision-making. In order to achieve this goal, we present two empirical analyses and a third example with a more practical approach to understanding economic and human behavior by using sentiment analysis. As a result, we illustrate the identification of affect, defined as an ‘encompassing term used to describe the topics of emotion, feelings, and moods together’ (Fox, 2008), in a dataset to describe and understand motivation and, consequently, human behavior.

Identifying Motivations in Data: Research Cases

In this section, we present two empirical cases that exemplify how psychological factors can be modeled and interpreted from data. These cases do not merely describe behavior, but also seek to uncover the internal drivers behind observable outcomes, using rigorous computational and methodological tools.

The first approach is provided by Chung et al. (2021), who study behavioral patterns on Airbnb, a global online platform connecting individuals looking to rent out their homes (hosts) to people seeking accommodation (guests). The object in this research is to understand the motivations behind individuals deciding to become Airbnb hosts, and how they are reflected on the platform.

To uncover them, the authors analyze open-ended responses to the question: ‘Why did you start hosting?’ The key variables of interest are the types of motivations (intrinsic vs. extrinsic, where intrinsic motivations are driven by internal rewards such as enjoyment, purpose, or self-fulfillment, and extrinsic motivations are driven by external rewards like money and recognition) and observable host behaviors on the platform, such as the number of pictures uploaded, length of property description, and guest satisfaction ratings.

The methodology combines two main techniques. First, they used Poisson factorization for topic

modeling, a probabilistic tool that identifies latent patterns in sparse datasets such as textual responses, without requiring labeled data. This method allowed the researchers to extract underlying motivational themes, such as meeting new people or earning supplemental income, from the hosts’ narratives. Second, they applied a multi-label classification approach, where responses were categorized simultaneously into multiple motivational labels (e.g., financial motives and social motives), recognizing that motivations are often multidimensional rather than mutually exclusive.

Their findings reveal a strong relationship between intrinsic motivations and higher engagement behaviors. Hosts motivated by intrinsic reasons tended to upload more photos, write longer and more detailed descriptions, and achieved higher guest satisfaction scores. In contrast, those driven primarily by extrinsic (financial) motives exhibited less engagement in these dimensions.

This study demonstrates a crucial behavioral insight: motivations leave measurable traces in actions. Even when motivations are not directly observable, patterns in user behavior, such as the effort put into property presentation, can serve as proxies for deeper motivational states. By employing appropriate analytical techniques, it is possible to bridge the gap between what individuals do and why they do it, moving beyond superficial interpretations of behavior.

A complementary approach to identifying motivations is presented by Hegde et al. (2024), who integrate the COM-B model (Capability, Opportunity, Motivation: Behavior) into the design of AI-based activity coaching systems.

The COM-B model (Michie et al., 2021) explains behavior as a function of three elements, namely, Capability (physical/psychological ability), Opportunity (external facilitators/barriers), and Motivation (internal processes driving action). This structure explains that the three elements must be present for a behavior to occur; a deficit in any element can lead to intended behaviors failing to materialize.

Building on this foundation, the authors embedded the COM-B framework into large language models (LLMs) used for digital health coaching. Through user interactions, such as refusal to engage, incomplete responses, or hesitation markers, the system learned

to distinguish whether inaction was driven by insufficient knowledge, environmental barriers, or low motivational state. This distinction is fundamental: two users may behave the same way (e.g., skipping a planned activity), but their underlying motivation, and therefore the optimal interventions, are entirely different.

The methodological strength of this case lies in its fusion of cognitive-behavioral theory with adaptive systems. Rather than treating all inactivity as disengagement, the COM-B-enhanced LLM interprets signals within context and adjusts its response accordingly, offering reminders, providing information, or prompting reflection based on the inferred source of inaction. In essence, it translates behavior into intention-based reasoning, thus producing a more personalized and ethically aligned mode of engagement.

Together with the Airbnb study, this case demonstrates that the challenge of identifying motivations in data is not impossible. Whether through the analysis of rich textual data or the structured modeling of behavioral components, it is possible to move beyond superficial interpretations of action toward a deeper, more nuanced understanding of human decision-making. In the next section, we propose an approach for applying these insights systematically in data-rich environments.

A Practical Approach to Identifying Motivations

In this section, the authors present a practical approach for understanding economic and human behavior, using sentiment analysis (DellaVigna & Pope, 2021). In addition, they highlight psychological markers, specifically affect, in real public opinions expressed in relation to an investment process, using Natural Language Processing (NLP) as a tool for analysis and SHAP as a visualization technique.

NLP as a Gateway to Behavioral Inference

NLP has emerged as a powerful tool for extracting behavioral insights from observational data. While traditional structured datasets provide observable actions (e.g., number of transactions, product returns), observational data, such as reviews, survey responses, and chatbot conversations, captures the internal narratives that drive behavior (Feuerriegel

et al., 2025).

Among the various NLP techniques, sentiment analysis plays a pivotal role in the early detection of affect. Sentiment not only reflects the affective state, but it also often serves as a proxy for future behavior and decision-making orientation. Positive or negative affect embedded in language can reveal levels of commitment, risk appetite, dissatisfaction, or openness to change (Pang & Lee, 2008), all of which are drivers of motivation.

In practice, sentiment analysis models classify textual expressions along a polarity axis (positive, neutral, negative) or more detailed affective spectra (e.g., joy, trust, anger, fear). When applied carefully, these affective signals can be interpreted as cognitive precursors of action. For instance, a customer's statement expressing frustration with service quality could predict a future churn decision, while enthusiastic expressions about a brand could signal an intention to repurchase (Mohammad, 2016).

From Detection to Explanation: Interpreting Motivations Through Model Explainability, Using SHAP-Based Interpretability

In the context of behavioral analysis, one of the greatest methodological challenges lies not only in making accurate predictions, but also in understanding how those predictions are made. When working with complex models, such as affective classifiers, interpretability becomes a key requirement for drawing behavioral conclusions (Mohammad, 2016). This is especially true when the goal is not just to detect affective tone, but also to infer underlying motivations.

SHAP (SHapley Additive exPlanations), proposed by Lundberg and Lee (2017), offers a unified and theoretically grounded framework to interpret model outputs. Based on cooperative game theory, it assigns each feature in an input (such as a word in a sentence) a contribution value that quantifies its impact on the model's prediction. This makes SHAP particularly useful for opening the black box of natural language models, revealing how individual linguistic elements influence classification outcomes.

In the context of sentiment analysis, SHAP has proven to be a powerful tool for (1) visualizing the influence of specific words or phrases on sentiment polarity (e.g., Ribeiro et al., 2016; Lundberg et al.,

2018) and (2) identifying misleading correlations or overfitting in model behavior by highlighting disproportionately weighted terms (Zhang et al., 2021).

This is particularly relevant when working with a behavioral approach to data, because affects are rarely isolated, and they often co-occur and influence one another. SHAP allows us to trace back from prediction to language, enabling the analyst to identify patterns consistent with known cognitive biases or motivational signals. In this sense, the framework is a behavioral inference tool that helps reduce the gap between computational outputs and psychological interpretation.

Illustrative Application: Sentiment Model With SHAP-Based Intention Decomposition

To demonstrate how SHAP can provide a granular understanding of the psychological structure of motivations within language, we designed illustrative examples using real commentary. These examples focus on contexts of investment and purchasing decision-making, where affective language plays a central role in shaping expectations and risk perception.

The underlying sentences were processed using a pre-trained DistilBERT sentiment classifier provided by the Hugging Face library. DistilBERT (Sanh et al., 2019) is a distilled version of BERT (Bidirectional Encoder Representations from Transformers), optimized to retain over 95% of BERT's language understanding capabilities while being 60% faster and lighter. This makes it especially suitable for sentiment analysis tasks where inference speed and computational efficiency are crucial. Hugging Face, as an open-source ecosystem for NLP, provides access to robust, benchmarked transformer models like DistilBERT that are fine-tuned for classification tasks, including those involving subjective affect (Liu, 2012).

The SHAP values were computed using the `shap.Explainer` function on the model's output logits. This allowed for token-level decomposition of predicted sentiment, revealing the marginal contribution of each word to the model's final decision. To aid interpretation, two complementary forms of visualization were employed: (1) text plots, which show word-level attribution along the sentence flow, using contextual coloring, and (2) bar plots, which offer a clearer comparison of magnitude across tokens.

As Liu (2012) notes in his comprehensive review on sentiment analysis, visual interpretability is essential in behavioral applications, particularly when affect is tied to decision-relevant constructs such as confidence, fear, or uncertainty. Therefore, both formats were used to capture not only the directional affect, but also its linguistic structure and weight.

The first statement was extracted from a user comment on Yahoo Finance, reflecting investor sentiment towards Tesla Inc. during a period of market volatility:

'Just because I'm down money doesn't mean this is a bad investment. The whole market, everything sucks right now. 2-5 years from now, I'm confident it will pay off'. (Source: Kaggle – Stock Tweets for Sentiment Analysis and Stock Price Prediction, 2021)

This comment reflects a real-world investor expression, rich in affective and temporal language, thereby offering an ideal case for sentiment deconstruction using SHAP.

The SHAP text plot (Figure 1) shows the model's predicted emotion is positive with high confidence (score: 0.967) despite multiple negative lexical cues. Terms such as "sucks" and "bad investment" carry negative connotations, reflected by their downward SHAP contributions in the bar plot (Figure 2). However, the prediction is outweighed by positive terms like "confident," "mean," and "pay off," which demonstrate strong upward SHAP values. This contrast confirms how certain lexemes, especially forward-looking optimistic terms, dominate the model's sentiment decision.

Behavioral Interpretation

This sentence encapsulates a textbook case of affect-driven investment reasoning. The subject acknowledges temporary losses ("down money") and macro pessimism ("everything sucks") but maintains an assertive belief in long-term gains. This aligns with the concept of *overconfidence* and *biased self-attribution*, as discussed in Daniel et al. (1998). Their theory suggests investors overweight their private judgments and reinterpret adverse outcomes as temporary noise, especially when market performance challenges their expectations.




Figure 1: SHAP text plots.




Figure 2: SHAP bar plots.

Furthermore, the strong SHAP value for “confident” reflects what Shleifer (2000) calls *affective forecasting* in asset markets, where expectations are influenced not only by data, but also by investors’ projected emotional payoff. The time frame mentioned (“2–5 years from now”) is indicative of a high willingness to wait, meaning that the investor has greater patience. A potential explanation for this behavior is the *endowment effect* (Konstantinidis et al., 2019). A cognitive bias that establishes a relationship between the ownership of an investment and the perception of its own value. People tend to evaluate more positively those financial assets they possess; this sense of property or emotional attachment makes it harder for the investors to sell them.

Lastly, the acknowledgment of broader market conditions (“the whole market, everything sucks right now”) suggests *external attribution*, whereby investors attribute personal investment losses to external factors, thereby preserving their belief in the investment’s quality. This coping mechanism helps maintain investment commitment during downturns (Hayward & Hambrick, 1997).

SHAP provides the necessary granularity to move from overall sentiment classification to a behaviorally meaningful interpretation of language. This form of modeling can be applied in contexts such as financial product recommendation, customer retention strategy, and risk communication, where understanding cognitive predispositions such as motivations is critical.

Model Limitations

Despite its strengths, the SHAP methodology has notable limitations when applied to behavioral analysis. First, while SHAP values offer word-level attribution, they are highly dependent on the underlying model’s representations and learned biases. If the sentiment classifier has internalized misleading associations, such as overemphasizing intensifiers or misinterpreting sarcasm, SHAP will reflect those distortions faithfully. Interpretability is only as reliable as the model’s alignment with real behavioral constructs. In addition, SHAP estimates each word’s contribution as if its effect could be analyzed separately from the rest of the sentence,

and it assumes a mostly linear relationship with the model’s output. However, in natural language, meaning often depends on the combination of words. This can lead to inaccurate attributions in complex or highly contextual sentences (Slack et al., 2020).

Second, SHAP does not inherently distinguish between genuine emotional expression and strategic or rhetorical language. A user may express strong emotions not as a reflection of internal motivation but to make a point, seek attention, or persuade. Moreover, because SHAP explains predictions at the word level, it may overlook important patterns in tone or meaning that emerge across a sentence or paragraph. From a behavioral science perspective, this increases the risk of overinterpreting surface-level sentiment as true underlying intention. SHAP should therefore be treated as a diagnostic tool, not a definitive indicator of intent, and used alongside contextual interpretation, complementary behavioral signals, and theory-driven validation.

While the SHAP-based sentiment analysis provides valuable insights into the motivational and affective structure behind investment-related language, a key limitation lies in the absence of behavioral validation. Future research could strengthen these findings by linking expressed sentiments with subsequent real-world investment actions in the given example. For instance, analyzing whether individuals that expressed optimism executed stock purchases thereafter would help to understand the alignment between stated motivation and realized behavior, thereby offering a more complete behavioral validation of the model.

Ethical Considerations and Biases

When analyzing human behavior, it is important to take into consideration ethical concerns and biases to prevent any type of harm or bad practice. With the framework we proposed, there are two relevant ethical concerns worth mentioning.

The first one refers to the fact that the distinction between the motivation and behavior might not want to be disclosed by the agent, for privacy issues. For example, in 2016, a dataset of more than 70,000 users² of a dating platform was published by a group of researchers. The data contained sensitive information

2 <https://www.vox.com/2016/5/12/11666116/70000-okcupid-users-data-release>

regarding sexual identity, ideology, fidelity, and politics, raising questions around privacy issues and intimacy. Also, dating life is a complex topic, and some people might use appearances and convenient attributes to signal certain characteristics, which can differ from their true nature. Precisely because of this, disentangling affect from true behavior can be tricky and must be handled with caution.

The second concern is the temporal validity of this type of analysis. Motivations, affect, as well as behavior change over time and evolve. This must be considered when assessing the real impact of the developed analysis, because it might not always be adequate or useful.

Conclusion

The aim of this article was to explain a behavioral analysis approach that goes beyond the scope of traditional data analysis. Instead of understanding behavior merely as the outcome of a decision, the authors consider motivations, shaped by cognitive states such as affect, as critical drivers in the decision-making process. Therefore, unlike conventional analytical approaches, this one incorporates sentiment analysis to uncover the latent psychological mechanisms embedded in natural language. By identifying motivations and affect via relevant linguistic cues in verbal expressions related to economic decisions, and by applying SHAP as a visualization technique, the authors provide a novel method for interpreting behavior through text.

This type of methodology can assist researchers and practitioners in behavioral science in navigating the analysis of affect in a more intuitive and transparent way. Understanding what shapes and conditions human behavior, even when it is not directly observable, represents a meaningful step toward unveiling latent decision processes.

Future research should explore how these methods can scale into AI-driven intention models, where large language models (LLMs) integrate behavioral priors to anticipate, classify, or even influence human affect in real time. Such models could open up new paths in applied behavioral science, thereby enabling personalized interventions, dynamic policy design, and deeper insights into the unobservable layers of cognition and affect that underpin economic and social behavior.

THE AUTHORS

Olympia Arango Castro holds a degree in Economics from Pompeu Fabra University and a Master in Economics and Psychology from the Paris School of Economics and Université Paris 1 Panthéon-Sorbonne. She has worked as a research assistant at academic institutions such as IESE Business School, focusing on topics like decision-making, consumer behavior, and leadership. Currently, she is a Behavioral Economics Analyst at Neovantas, where she supports projects related to behavioral analysis, experimental design, and training programs. Her professional career focuses on creating synergies between academic research and applied work, exploring how behavioral science can inform practical interventions.

Miguel Antonio Rodríguez Ramírez holds a degree in Actuarial Science from the Instituto Tecnológico Autónomo de México (ITAM), where he also specialized in Data Science and Machine Learning. He has worked in leading financial institutions in Mexico, focusing on data analytics, experimental design, and behavioral economics. His experience bridges quantitative analysis with behavioral insights to drive strategic decision-making. Currently, Miguel is a business specialist at Neovantas, where he leads behavioral data initiatives for the largest bank in Mexico. His professional work lies at the intersection of data science and behavioral economics, applying rigorous methodologies to generate actionable insights with real-world impacts.

Juan de Rus Gutiérrez is Managing Partner at Neovantas and Adjunct Professor of Behavioral Science and Marketing at IE Business School, Universidad Carlos III of Madrid, and the University of Barcelona. Juan holds an MSc in Behavioral Science with distinction from the London School of Economics & Political Science (LSE), and he also holds an MBA from Universidad Carlos III de Madrid. He is also a certified member of GAABS.

Emma Bernardo Sampedro is Head of Marketing at Neovantas. She holds a degree in Journalism from the Complutense University of Madrid (UCM) and a degree in Psychology from the International University of La Rioja (UNIR). She recently completed a Master's degree in Behavioral Data Science at the University of Barcelona and also holds a specialization in People-Centered Innovation from the H2i

Institute. With over a decade of experience in the bancassurance and health sectors, she combines a background in communication, psychology, and customer experience with a growing focus on behavioral economics.

REFERENCES

- Albarracín, D., Chan, M. P. S., & Jiang, D. (2018). Attitudes and attitude change: Social and personality considerations about specific and general patterns of behavior. In K. Deaux & M. Snyder (Eds.), *The Oxford handbook of personality and social psychology* (pp. 439–463). Oxford University Press.
- Altmann, J. (1974). Observational study of behavior: Sampling methods. *Behaviour*, 49(3–4), 227–267. <https://www.jstor.org/stable/4533591>
- Baillie, L., & Higgins, S. (2024). Observational data. In J. E. Edlund & A. L. Nichols (Eds.), *The Cambridge Handbook of Research Methods and Statistics for the Social and Behavioral Sciences: Volume 2: Performing Research* (pp. 665–685). Cambridge University Press.
- Beauchamp, A. J. (2024). *Observational methods in psychology research*. EBSCO Research Starters. <https://www.ebsco.com/research-starters/psychology/observational-methods-psychology-research>
- Brosnan, D. (2008). Research methods. In M. Poulter (Ed.), *The library and information professional's guide to the Internet* (pp. 27–44). Chandos Publishing. <https://doi.org/10.5772/intechopen.68749>
- Buisson, F. (2021). *Behavioral data analysis with R & Python: Customer-driven data for real business results*. O'Reilly Media.
- Carthey, J. (2003). The role of structured observational research in health care. *BMJ Quality & Safety*, 12(Suppl 2), ii13–ii16. https://doi.org/10.1136/qhc.12.suppl_2.ii13
- Chung, J. Y., Oh, S., & Choi, Y. (2021). *Uncovering consumer motivations in the sharing economy: A text-mining approach*. Columbia University.
- Daniel, K., Hirshleifer, D., & Subrahmanyam, A. (1998). A theory of overconfidence, self-attribution, and security market under- and over-reactions. *Journal of Finance*, 53(6), 1839–1885.
- DellaVigna, S., & Pope, D. (2021). What motivates effort? Evidence and expert forecasts. *Journal of Economic Behavior & Organization*, 185, 146–163. <https://doi.org/10.1093/restud/rdx033>
- Einav, L., & Levin, J. (2014). Economics in the age of big data. *Science*, 346(6210), 1243089. <https://doi.org/10.1126/science.1243089>
- Fox, E. (2008). *Emotion science: Cognitive and neuroscientific approaches to understanding human emotions*. Palgrave Macmillan.
- Feuerriegel, S., Maarouf, A., Bär, D., Geissler, D., Schweisthal, J., Pröllochs, N., Robertson, C. E., Rathje, S., Hartmann, J., Mohammad, S. M., Netzer, O., Siegel, A. A., Plank, B., & Van Bavel, J. J. (2025). Using natural language processing to analyse text data in behavioural science. *Nature Reviews Psychology*, 4, 96–111. <https://doi.org/10.1038/s44159-024-00392-z>
- Gomez-Diaz, T., & Recio, T. (2022). Research Software vs. Research Data I: Towards a Research Data definition in the Open Science context. *F1000Research*, 11, 118. <https://doi.org/10.12688/f1000research.78195.2>
- Girard, J. M., & Cohn, J. F. (2016). A primer on observational measurement. *Assessment*, 23(4), 404–413. <https://doi.org/10.1177/1073191116635807>
- Hayward, M. L. A., & Hambrick, D. C. (1997). Explaining the premiums paid for large acquisitions: Evidence of CEO hubris. *Administrative Science Quarterly*, 42(1), 103–127. <https://doi.org/10.2307/2393810>
- Hegde, N., Vardhan, M., Nathani, D., Rosenzweig, E., Speed, C., Karthikesalingam, A., et al. (2024). Infusing behavioral science into large language models for activity coaching. *PLOS Digital Health*, 3(4), e0000431. <https://doi.org/10.1371/journal.pdig.0000431>
- Konstantinidis, A. D., Katarachia, A., & Siskou, T. (2019). Disposition effect, endowment effect, attachment bias and investment advisers. *IOSR Journal of Business and Management*, 21(2), 16–21. <https://doi.org/10.9790/487X-2102041621>
- Ng, B. (2023). The neuroscience of emotion and intrinsic motivation. In W. L. D. Hung, A. Jamaludin, & A. A. Rahman (Eds.), *Applying the science of learning to education* (pp. pp 79–97). Springer. https://doi.org/10.1007/978-981-99-5378-3_4

- Liu, B. (2012). Sentiment analysis and opinion mining. *Synthesis Lectures on Human Language Technologies*, 5(1), 1–167.
- Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *NeurIPS*. <https://doi.org/10.48550/arXiv.1705.07874>
- Paolisso, M., & Hames, R. (2010). Methods for the systematic study of human behavior. *Field Methods*, 22(4), 1–13.
- Michie, S., van Stralen, M. M., & West, R. (2011). The behaviour change wheel: A new method for characterising and designing behaviour change interventions. *Implementation Science*, 6(42). <https://doi.org/10.1186/1748-5908-6-42>
- Mohammad, S. M. (2016). Sentiment analysis: Detecting valence, emotions, and other affectual states from text. In R. A. Calvo, S. K. D'Mello, J. Gratch, & A. Kappas (Eds.), *The Oxford handbook of affective computing* (pp. 201–237). Oxford Academic.
- Pallier, G., Wilkinson, R., Danthiir, V., Kleitman, S., Knezevic, G., Stankov, L., & Roberts, R. D. (2002). The role of individual differences in the accuracy of confidence judgments. *Journal of General Psychology*, 129(3), 257–299. <https://doi.org/10.1080/002213300209602099>
- Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, 2(1–2), 1–135. <http://dx.doi.org/10.1561/15000000011>
- Radu, C. M., & Radu, A. D. (2014). The contribution of behavioral economics in explaining the decisional process. *Procedia – Social and Behavioral Sciences*, 109, 1052–1059. [https://doi.org/10.1016/S2212-5671\(14\)00821-1](https://doi.org/10.1016/S2212-5671(14)00821-1)
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why should I trust you?”: Explaining the predictions of any classifier. *KDD '16: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135–1144. <https://doi.org/10.1145/2939672.2939778>
- Sanh, V., Debut, L., Chaumond, J., & Wolf, T. (2019). DistilBERT, a distilled version of BERT: Smaller, faster, cheaper and lighter. *arXiv Preprint*. <https://doi.org/10.48550/arXiv.1910.01108>
- Shaikh, A., Dandekar, R. A., Panat, S., & Dandekar, R. (2024). CBEval: A framework for evaluating and interpreting cognitive biases in LLMs. *Vizuara AI Labs*. <https://arxiv.org/abs/2412.03605>
- Shleifer, A. (2000). *Inefficient markets: An introduction to behavioral finance*. Oxford University Press.
- Slack, D., Hilgard, S., Jia, E., Singh, S., & Lakkaraju, H. (2020). Fooling LIME and SHAP: Adversarial attacks on post hoc explanation methods. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 180–186. <https://doi.org/10.1145/3375627.3375830>
- Slovic, P., Finucane, M. L., Peters, E., & MacGregor, D. G. (2002). The affect heuristic. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 397–420). Cambridge University Press.
- Souza, G. C., Meireles, E., Mira, V. L., & Leite, M. M. J. (2021). Academic motivation scale: Reliability and validity evidence among undergraduate nursing students. *Revista Latino-Americana de Enfermagem*, 29, e3420. <https://doi.org/10.1590/1518-8345.3848.3420>
- Yadollahi, A., Shahraki, A. G., & Zaiane, O. R. (2017). Current state of text sentiment analysis from opinion to emotion mining. *ACM Computing Surveys*, 50(2), 25. <https://doi.org/10.1145/3057270>
- Zhang, Z., Jin, H., & Zhou, L. (2021). Sentiment analysis: From traditional to deep learning models. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*.

The Integrated Ideation Model for Data-Use Culture (IIMDC): Tackling Behavioral Challenges Within Health Information Systems

AKINYI ODERA, MELAKU MULETA AND ARIAM HAILEMARIAM¹

ThinkPlace Ethiopia

The Integrated Ideation Model for Data-use Culture (IIMDC) offers a behavioral framework for strengthening how health workers collect, interpret, and use patient- and clinical-related data in decision-making in Ethiopia. The framework identified drivers like low awareness of data relevance, diffusion of responsibility, and limited perceived impact of decisions at the facility level. It guides targeted interventions to enhance salience, clarify roles, and strengthen leadership accountability, promoting routine data use. While prior efforts focused on technical capacity, behavioral drivers were often overlooked. Designed interventions targeted providers, health information technicians, review team members, and leaders, with the aim of improving ownership, perceived value, and accountability. By applying behavioral insights across the data journey—recording to decision-making—we show herein how tailored interventions can strengthen data-use culture and improve health outcomes.

Introduction

Behavioral Science and Data-Use Culture

Health workers, program managers, and policy-makers rely on routine information, such as health services utilization levels, health service provider capacity, availability of drugs and medical supplies, the prevalence of disease outbreaks, etc., to make decisions that improve health outcomes (World Health Organization, 2008).

However, in most health facilities in Ethiopia, the quality of this information and, consequently, its use during decision-making is lacking. Systemic challenges such as infrastructural constraints, duplicative tools, and staff shortages contribute to this issue. In addition, behavioral factors such as a low perceived value of health data, limited reinforcement of positive data management behaviors, and unintentional and inconsistent data management frequently undermine the accurate and timely generation of high-quality data and its resultant use during decision-making. These conditions impede the development of a robust data-use culture within the Ethiopian health sector.

Desired State

The Data Use Partnership (DUP) envisions a health system where data is consistently complete, timely, and accurate. Health workers record and submit data in alignment with required offline and online registration formats, ensuring all necessary variables are captured, and reporting occurs in real-time or immediately after service delivery, thus reducing the risk of missing or distorted information. A culture of accountability is cultivated whereby department heads and staff feel compelled to report data that genuinely reflects facility realities. Most importantly, data is actively utilized for decision-making, with department heads and performance monitoring teams consistently referring to centralized information to guide planning and improvement efforts.

The Purpose of the IIMDC Model

The IIMDC provides a structured approach to achieving this desired state. It enhances data utilization in health systems by integrating behavioral insights into the ideation process. It achieves this aim by drawing on three foundational frameworks: PRISM, (Aqil et al., 2009), COM-B (Michie et al., 2011),

1 Corresponding author: ariam.hailemariam@thinkplace.co.ke

and ideation theory (Health Communication Capacity Collaborative, 2015).

The PRISM framework highlights that technical, organizational, and behavioral determinants shape data use. Technical factors include data collection tools and IT systems, while organizational determinants encompass structures, resource availability, and institutional culture. Complementing this approach, the COM-B model provides a behavioral science perspective, identifying capability, opportunity, and motivation as essential drivers of behavior. It emphasizes that fostering a data-use culture requires enhancing knowledge (capability), creating an enabling environment (opportunity), and fostering intrinsic motivation to engage with data.

Ideation theory operationalizes these insights by structuring the ideation process across three key perspectives: Cognitive, social, and emotional enablers. Cognitive strategies focus on simplifying decision-making through nudges, framing techniques, and habit formation. Social strategies leverage peer influence, group commitments, and normative messaging to create collective momentum for behavior change. Emotional strategies, such as recognition, rewards, and storytelling, tap into intrinsic motivation to sustain engagement.

Case Study: Application of the IIMDC

The IIMDC guides two stages of a design journey:

1. **Problem Identification:** Understanding behavioral barriers to data quality and use.
2. **Solution Design:** Ideating and shaping behaviorally informed solutions and embedding them in the broader health system.

Problem Identification: The Seven Predominant Behavioral Drivers

In the DUP's problem identification phase, research findings related to this case study indicated that a core behavioral challenge was that health providers undervalued the importance of data in improving healthcare services, leading to negligence in documentation and reluctance to engage in data-driven decision-making.

Furthermore, insufficient training and skills in data management among healthcare workers hindered their ability to collect, process, and interpret health information effectively, thereby creating gaps in data utilization. Competing priorities and work overload further exacerbated the challenge, as they juggled multiple responsibilities, often leaving insufficient time and attention for data-related tasks. The absence of incentives also played a critical role in discouraging




Figure 1: The integrated ideation model for data-use culture.

data use, as health workers who received no tangible rewards or recognition for data-related activities tended to deprioritize them.

In addition to these individual-level barriers, leadership and supervisory gaps significantly affected data use in health facilities. Weak enforcement of data management protocols and limited managerial support resulted in poor adherence to data collection and utilization practices. Lastly, resource constraints, including limited staffing and inadequate infrastructure, posed significant obstacles to effective data management, restricting the capacity for proper data collection, storage, and analysis.

Through the application of the IIMDC during DUP research with healthcare workers in Ethiopia, the following predominant behavioral barriers were identified:

- **Visibility:** Poor data quality occurs when the impact of the data is not visible. Data-informed decisions are not habitually communicated to health workers who record and report patient information.
- **Self-efficacy:** Poor data quality occurs when there are fewer active data-use moments at the health facility level. Data from service providers is not habitually used for department- or facility-level decision-making.
- **Leadership and accountability:** Poor data quality and use occur when there is a lack of oversight on data recording activities. This demotivates the intentional and consistent generation of high-quality data.
- **Recognition and reinforcement:** Poor data quality occurs when positive data recording and reporting behaviors are not recognized. These behaviors are not often acknowledged formally or informally.
- **Limited skill capacity:** Poor data quality occurs when there is a lack of confidence in data recording, visualization, analysis, and reporting skills. Health workers are willing to record and report according to required standards but are not confident they can do it correctly.
- **Competing priorities:** Poor data quality occurs when health service provision is prioritized over data recording and reporting. This deprioritizes data-related tasks.

- **Reliability of indicators:** Poor data quality occurs when reporting indicators focus more on health facilities' performance than on health service delivery milestones. Service delivery milestones reflect "metrics that matter" more to the health workers who record, report, and use data.

Behavioral insights were generated using a qualitative ethnographic study design with in-depth interviews, key informant interviews, focus group discussions, and other HCD research data collection methods across six regions in Ethiopia, namely, Addis Ababa, Amhara, Oromia, Sidama, Tigray, and Somali, in close collaboration with local academic institutions. This research informed the development of the IIMDC framework by unpacking the various behavioral drivers and barriers and how they interplay to affect the data-use culture.

Other similar studies highlight that many health workers perceive data analysis as too complex, leading to avoidance. Low confidence, exacerbated by limited support structures, undermines motivation to engage with data, especially when tasks require technical interpretation (Mekonnen & Tesfaye, 2021). Additional studies also highlight that staff are less likely to internalize data as a norm without leaders modeling strong data-use behaviors or reinforcing accountability (Fetene et al., 2020). In addition, motivation wanes when data-use behaviors are not acknowledged, health workers respond to positive reinforcement, and its absence leads to disengagement (Abera et al., 2022), and data tasks are filtered out as non-essential, especially when systems are duplicative or infrastructure is weak (Tilahun et al., 2022).

Solution Design: Ideation, Prototyping, and Testing Behavioral Interventions

The IIMDC guided the design of various behavioral prototypes by targeting key barriers to data use. While several examples were tested and implemented across sites, this document highlights two prioritized interventions to illustrate the application of behavioral insights selected based on their relevance, scalability, and potential for impact, as well as to manage the volume of information presented. Two prototypes, used to nudge intentional leadership and improve the perceived value of data among health providers, are described below.

Strengthening Recognition: The Data Stars Spotlight

The Data Stars Spotlight is a recognition system designed to increase health workers' intrinsic motivation by instantly recognizing staff who perform desired data quality behaviors, such as reporting patient information immediately after health service delivery or at the end of the day. It is designed to 'catch people doing something right' and frequently reinforces this behavior in genuine, non-monetary ways. Any reward is more impactful when provided in a timely manner and when the recipient's salience (conscious awareness) is high (Gao et al., 2021). Through genuine and instant recognition, the Data Stars Spotlight provides the dopamine needed to condition the desired data quality behavior.

The Data Stars Spotlight was designed to address two critical gaps. First, there was an absence of negative incentives for the lack of performance of data-related duties, i.e., there was no difference between when data recording and reporting duties were diligently done and when they were not performed at all. No one would notice. This introduced the second gap, namely, a glaring absence of leadership in the health facilities' data management practices. Field research activities showed that the lack of intentional and consistent oversight strongly contributed to poor data quality and, as a result, the lack of its use in day-to-day operational decision-making and health service delivery. The Data Stars Spotlight therefore sets out to engage health facility and local government leadership to participate more actively in data oversight activities. It gives tips about the desired behaviors that leaders can look out for daily and weekly, and how to affirm and reinforce those behaviors genuinely. It also outlines a more formal recognition system, where 'Data Stars' are socially recognized on a monthly, quarterly, and annual basis on national platforms. This way, the Data Stars Spotlight addresses recognition and leadership engagement gaps.

The IIMDC positions organizational factors, including supervision, as influencing a health facility's data use culture. The model also highlights the need for non-directive strategies, including recognition. These behavioral strategies are used within the Data Stars Spotlight, as described below.

Social Recognition

The Data Star Spotlight uses informal recognition strategies embedded within the social interactions between health facility staff and leadership, earned based on the performance of desired data quality behaviors. Beyond health facility-level recognition, health workers progressively earn opportunities for acknowledgment among authorities at the national level and, as a result, career progression. The DUP's research activities indicated that while desirable, health workers want tighter monitoring and evaluation mechanisms, to reduce the chances of gaming the recognition system.

Autonomy

When people are more intrinsically motivated, they are more likely to put forth their best effort, even when there is no reward, when it is challenging, and when no one is looking (Cho & Perry 2012, Grant 2008; Ryan et al., 2008). Introducing a sense of autonomy is one way to increase intrinsic motivation. When people have more autonomy over their behavior, they enjoy it and want to do it more, but when a behavior is forced, for instance through punishment, there is a higher likelihood for rebellion to regain control (Ryan & Deci, 2000). The Data Stars Spotlight is designed to inspire rather than to prescribe how and when leaders can affirm—and therefore reinforce—their staff's performance in terms of data recording, reporting, and use behaviors. It also uses selection criteria co-created between the health facility staff and leadership. While introducing this sense of autonomy in the design of the selection criteria was desirable, the DUP research findings in Amhara indicated a need to balance its implementation with structure, such as pre-set goals and check-in periods, and evaluation tools, such as a checklist and individual charts. While Data Star strongly leverages autonomy as a component of intrinsic motivation, it is used alongside already existing extrinsic motivation strategies, such as certificates and mobile cards issued as rewards to health workers who diligently perform data-related duties.

Competence

When people feel like a behavior moves them up meaningfully, they are more committed to it, for

example, if the behavior promotes career development and increases commanded respect. Early studies showed that positive performance feedback enhanced intrinsic motivation, whereas negative feedback diminished it (Deci, 1975). The Data Stars Spotlight is designed to increase health workers' commanded respect through frequent and genuine acknowledgment of competence and the publication of the Data Star in national and regional media and authority forums such as the annual review meeting.

'I feel more confident using the board to explain changes in service numbers' (health worker, Sidama).

Monetary incentives are widely used to improve health system performance. However, even when they are well designed and work perfectly, they can be expensive to maintain (Gneezy & Rustichini, 2000), people may experience loss aversion if the incentive is removed (Jeffrey 2004), and they are likely to crowd out non-monetary or intrinsic incentives. The motivation crowding effect suggests that external intervention via monetary incentives or punishments may undermine and—under different identifiable conditions—strengthen intrinsic motivation (Frey & Jegen, 2001).

The DUP research indicated a need to adapt the Data Star Spotlight, considering that monetary rewards had already been popularized in some health facilities. One main constraint was the visible contribution of the monetary rewards to the health facility's improved infrastructure, environment, and health workers' living standards, and, separately, strong support from the leadership. This necessitated the inclusion of strategies to ensure that extrinsic and intrinsic motivation strategies could co-exist.

The Data Star Spotlight attempts to do this by crowding in intrinsic motivators, with the intention of potentially progressively crowding out extrinsic motivators. While crowding in is most likely to happen in areas where the person is not already doing the extrinsically-motivated behavior (Frey & Jegen, 2001), DUP research findings indicated an existing opportunity to leverage health workers' desire to save lives as a superior reason for the performance of desired behaviors, namely, accurate recording and reporting, with or without the associated monetary benefits. In some regions, where performance-based incentives

are widely used, some health workers expressed that they find value when data informs their health service delivery duties and, as a result, saves lives.

'[...] Information is helpful to the ministry and the country. I am mentally satisfied when I provide good information. I would be happy' (health worker, Oromia).

The research insights also indicated that the possibility of achieving this would need to be closely linked to the ability of health workers to see their data positively impacting health service delivery. Therefore, the Data Stars Spotlight needs to be supported by interventions to increase access to reported data.

Strengthening Data Visibility: Eyita Analytics Bot

The Eyita Telegram Analytics Bot is a mobile-based data analytics platform that allows health workers to access reported data via the Telegram app, on-demand. Previously, access to reported data was limited to leaders at the higher administrative levels and data clerks, but Eyita expands access to health service providers, creating a strong reason for data use at the point of health service delivery.

The IIMDC's ideational factors highlight the role of salience and descriptive norms in influencing data culture. These are considered in the Eyita design as follows.

Increasing Salience

Eyita increases data access to a broader stakeholder base, making it more present in day-to-day activities within the health facility. It de-mystifies its functionality. Previously, data was externalized, i.e., it was seen as a product only useful for stakeholders higher up the administration, and it was used to reward or punish rather than for supportive supervision. Therefore, reporting would be done in a way that minimized any potential negative effects, such as loss of funding and recognition, rather than demonstrating the needs of the health facility. Research insights showed that health workers would be more inclined to misreport if there was a risk of reduced funding. The research activities also indicated that they would likely generate high-quality data if they understood its functionality.

'Before, data stayed on paper. Now, we talk about it during meetings' (facility head, Oromia).

In the Sidama region, healthcare workers reported using *Eyita*-generated insights during weekly department-level discussions and to prepare reports. In the same region, health facility managers requested additional indicators used to track disease trends and post-level health data, to increase visibility on full primary Health Care Unit (PHCU) performance. *Eyita* also pushed the functionality of data as a benchmark or reference point.

In this regard, a performance management team member reported,

‘Eyita helped me easily see our performance, where we are as a facility against our plan[...].’

While *Eyita* showed great potential, two main limitations were notable. The first was related to internet access. The majority of health workers had access to the internet through data bundles, but using *Eyita* during daily case management activities would require an investment of about \$1 per day. Some who had already purchased data for social media usage didn’t mind using the same internet package for work-related purposes; however, others who were not frequently online were less willing to invest in data bundles at their own expense. Further research is needed to understand the feasibility of health facility-level investment to encourage usage.

The second notable limitation was low technology acceptance among older health workers. Research Insights revealed excitement and enthusiasm among younger workers, as compared to anxiety and fear among older health workers introduced to *Eyita*. In the Tigray region, some older health workers were not online often, and so they were unfamiliar with and afraid to touch the app. Further research will help understand the efficacy of younger peers as onboarding champions.

Improving Social Norms

Social norms are the unwritten, informal rules that dictate acceptable and appropriate behaviors within a group, community, or culture, guiding human interactions and expectations (Bicchieri, C. 2006). *Eyita* intends to normalize the use of data to inform routine decision-making in health facilities, which it achieves by:

- Increasing access to reported data: making it easy to refer to data during day-to-day activities related to health service provision and daily operations.
- Creating social proof: broadcasting use cases among peers to influence late adopters and create an in-group effect. This is done by communicating usage frequency through an inbuilt leaderboard.
- Increasing competence: DUP research in Amhara illustrated the intrinsic benefits of health worker-driven performance tracking:

‘After entering the data on visualization, we see how happy and confident users feel’ (Amhara researcher).

This also provided the positive reinforcement needed to ensure data-use behaviors are repeated in the future.

Indicative Results

Qualitative research methods were used to generate insights about the efficacy of the behavioral prototypes. Focus group discussions, one-on-one interviews, and observations were conducted in three design sites per region over one month. Due to a collaborative design process, critical signals of change emerged.

Health workers reported increased intrinsic motivation and confidence in performing data-related tasks, attributing this to timely and genuine recognition of their efforts. Leadership engagement in data oversight also improved, with facility heads actively using tips from the intervention to monitor and affirm desired behaviors.

In the Sidama region, there were signs of improved attitudes towards the recording of health information, as well as a sense of improved social support. Research Insights show that while health workers know their information recording roles and responsibilities, producing high-quality monthly reports is a shared responsibility and is valuable for the health facility. Due to this improved attitude, health workers initiated simple solutions, such as color-coded bins, to help separate registered and unregistered patient cards and files. Previously, these were kept together, but now they are separated, and the unregistered documents

are filled in right after health service delivery.

In the Amhara region, once seen as routine, weekly morning meetings are now seen as reflective spaces where staff review data trends and discuss progress. During a morning meeting in one health center, the leader led discussions and shared insights from prior training on data quality and registration. He emphasized that data quality is not just about reporting, but also a resource-pooling mechanism. This initiative indicates improved, intentional leadership—a building block of the IIMDC.

In the Somali region, previously dominated by informal systems, a sense of structure is emerging through the use of *Eyita* to record and track trends weekly. In a health center in Hodale, health workers tracked typhoid cases using an offline data collection system. They compared the current data to the annual plan and identified a spike in cases. An investigation revealed that polluted ponds, a key water source, contributed to the outbreak. Chlorine tablets were distributed at the health center to purify household drinking water, and as a result, typhoid cases decreased. Balancing this emerging formal structure with conversations helps maintain a feedback loop, and in turn it builds trust and encourages regular use. There is an emerging attitudinal change regarding perceived value and the practical use of information.

Initial results from the behavioral prototypes, such as increased motivation among health workers, improved leadership engagement, and more consistent data recording practices, suggest positive changes in data-use culture in Ethiopian health facilities. While a randomized controlled trial (RCT) is planned, to test the impact of these behaviors rigorously, early indications point to meaningful improvements in the health information system.

Conclusion

The Integrated Ideation Model for Data Culture (IIMDC) plays a primary role in the research and design phases involved in improving data-use behavior change. Its strength lies in its ability to diagnose behavioral barriers and shape prototypes through iterative and co-creation approaches before broader scaling efforts begin. This initiative contributes important learning to the emerging field of

applied behavioral science in health systems. Early outcomes highlight the value of using behavioral frameworks like IIMDC to design context-sensitive, people-centered interventions. As the prototypes evolve and undergo further evaluation, they will deepen the evidence base on what drives effective, sustained data use in resource-constrained settings.

THE AUTHORS

Akinyi Odera is a senior executive designer at ThinkPlace Global, where she has worked for the past seven years. She is the project director for Ethiopia's Data Use Partnership Project (Workstream 3) and has over 15 years of industry experience in strategic communication, human-centered design, behavior change, and systems thinking. Her most recent work has resulted in the incubation of multiple digital and offline prototypes designed to improve the perceived value of health data among frontline health workers in low-resource settings.

Melaku Muleta is deputy director and a health information systems specialist for the Data Use Partnership Project at ThinkPlace Ethiopia, where he leads initiatives that apply human-centred design and behavioural science to improve the performance and usability of health information systems. With extensive experience in public health program design, implementation, and evaluation, he has managed large-scale initiatives across primary healthcare, NTDs, HIV/AIDS, and WASH. He has published articles in international journals and presented research at multiple international conferences. Initially trained in clinical and community health at a degree level, Melaku holds a Master's degree in public health, and he is an internationally certified healthcare quality improvement advisor.

Ariam Hailemariam Asrat is a social and behavior change expert and consultant at ThinkPlace, where she is engaged in the development and testing of behavioral solutions for strengthening health systems in Ethiopia. Her experience in public health, behavioral science, project management, and lecturing in higher institution brings a strong focus on evidence-based design and implementation. Her expertise spans human-centered design and mixed method behavioral research. Ariam's recent work includes the

development of data-use culture interventions and scalable behavior change prototypes for the Data Use Partnership Project (workstream 3).

REFERENCES

- Abera, A., Shiferaw, T., Girma, A., Mohammed, R., Ahmed, M., Shalmeno, B., Weldesenbet, A. B., Tolera, A., & Dessie, Y. (2022). Health facilities performance monitoring team focused motivation interventions to improve the use of health information for better decision making: An implementation research study protocol. *Ethiopian Journal of Health Development*, 36(SI-1). <https://www.ejhd.org/index.php/ejhd/article/view/5374>
- Aqil, A., Lippeveld, T., & Hozumi, D. (2009). PRISM framework: A paradigm shift for designing, strengthening and evaluating routine health information systems. *Health Policy and Planning*, 24(3), 217–228. <https://doi.org/10.1093/heapol/czp010>
- Asulin, Y., Heller, Y., & Munichor, N. (2024). Comparing the effects of non-monetary incentives and monetary incentives on prosocial behavior. *European Economic Review*, 165, 104740. <https://doi.org/10.1016/j.euroecorev.2024.104740>
- Bicchieri, C. (2017). *Norms in the wild: How to diagnose, measure, and change social norms*. Oxford University Press.
- Cho, Y. J., & Perry, J. L. (2012). Intrinsic motivation and employee attitudes. *Review of Public Personnel Administration*, 32(4), 382–406. <https://doi.org/10.1177/0734371X11421495>
- Dannals, J. E., & Li, Y. (2024). A theoretical framework for social norm perception. *Research in Organizational Behavior*, 44, 100211. <https://doi.org/10.1016/j.riob.2024.100211>
- Fetene, N., Patel, A., Benyam, T., Ayde, A., Desai, M. M., Curry, L., & Linnander, E. (2020). Experiences of managerial accountability in Ethiopia's primary healthcare system: A qualitative study. *BMC Family Practice*, 21, 261. <https://doi.org/10.1186/s12875-020-01332-5>
- Frey, B. S., & Jegen, R. (2001). Motivation crowding theory. *Journal of Economic Surveys*, 15(5), 589–611. <https://doi.org/10.1111/1467-6419.00150>
- Gao, Z., Wang, H., Lu, C., Lu, T., Froudust-Walsh, S., Chen, M., Wang, J., Hu, J., & Sun, W. (2021). The neural basis of delayed gratification. *Science Advances*, 7(49), eabg6611. <https://doi.org/10.1126/sciadv.abg6611>
- Gneezy, U., & Rustichini, A. (2000). Pay enough – or don't pay at all. *The Quarterly Journal of Economics*, 115(3), 791–810. <https://doi.org/10.1162/003355300554917>
- Hummel, D., & Maedche, A. (2019). How effective is nudging? A quantitative review on the effect sizes and limits of empirical nudging studies. *Journal of Behavioral and Experimental Economics*, 80, 47–58. <https://doi.org/10.1016/j.socec.2019.03.005>
- Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.
- Krath, J., Schürmann, L., & Von Korflesch, H. F. (2021). Revealing the theoretical basis of gamification: A systematic review and analysis of theory in research on gamification, serious games, and game-based learning. *Computers in Human Behavior*, 125, 106963. <https://doi.org/10.1016/j.chb.2021.106963>
- Mekonnen, T., & Tesfaye, S. (2021). Exploring health professionals' competence in data analysis and visualization in Ethiopian healthcare settings: A pre-post intervention study. *Ethiopian Journal of Health Development*, 35(1), 30–38. <https://www.ejhd.org/index.php/ejhd/article/view/4604>
- Michie, S., van Stralen, M. M., & West, R. (2011). The behavior change wheel: A new method for characterizing and designing behavior change interventions. *Implementation Science*, 6, 42. <https://doi.org/10.1186/1748-5908-6-42>
- Ryan, R. M., & Deci, E. L. (2000). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American Psychologist*, 55, 68–78. <https://psycnet.apa.org/doi/10.1037/0003-066X.55.1.68>
- Tilahun, B., Areru, M., Taddese, A., Shiferaw, F., & Kachero, B. (2022). Applying behavioral science to health systems: Pathways and evidence gaps. *Global Health: Science and Practice*, 10(5), e2200253. <https://doi.org/10.9745/GHSP-D-22-00253>

- Tilahun, H., Abate, B., Belay, H., Gebeyehu, A., Ahmed, M., Simanesew, A., Ayele, W., Mohammedsanni, A., Knittel, B., & Wondarad, Y. (2022). Drivers and barriers to improved data quality and data-use practices: An interpretative qualitative study in Addis Ababa, Ethiopia. *Global Health: Science and Practice*, 10(Suppl 1), e2100689. <https://doi.org/10.9745/GHSP-D-21-00689>
- Waithaka, D., Cashin, C., & Barasa, E. (2022). Is performance-based financing a pathway to strategic purchasing in Sub-Saharan Africa? A synthesis of the evidence. *Health Systems and Reform*, 8(2), e2068231. <https://doi.org/10.1080/23288604.2022.2068231>
- World Health Organization. (2008). *Framework and standards for country health information systems* (2nd ed.). World Health Organization.
- Health Communication Capacity Collaborative. (2015). *Ideation: An HC3 research primer*. <https://www.healthcommcapacity.org/wp-content/uploads/2015/02/Ideation.pdf>

CEO Briefing: Behavioral Science in the Workplace

MINH HUA¹

This briefing introduces practical applications of behavioral science that can drive competitive advantage in commercial organizations. Drawing from field experience in multiple industries, as well as research, the author shows how improved management techniques has lead to real business value. The first case study shows how artificial intelligence can improve productivity and details some of the pitfalls that come with new technology. The second case study takes a deep dive into Amazon's hiring mechanism and looks at what may be in store for the future. The third case study illustrates how an investment bank got more value out of their bonus payout—without spending additional money. The article concludes with advice on transforming a company's culture to be more evidence-based and willing to experiment..

Introduction

'A long-standing question is whether differences in management practices across firms can explain differences in productivity', wrote Nicholas Bloom, a Stanford professor and former policy advisor to Her Majesty's Treasury.

This is an important question for CEOs because higher productivity creates an opportunity to pull on a host of strategic levers and then take the cost savings to boost profit margins, lower prices to increase market share, or divert resources to new ideas.

A team led by Bloom evaluated management practices inside 6,000 companies representing every region of the world. They concluded that a one-standard deviation improvement in scores on management practices (a composite measure in several areas, including operations, monitoring and people management) is associated with 38% higher productivity (Bloom & Reenen, 2007). Another team of researchers led by Bloom found in a randomized controlled study that free consulting on management practices improved productivity by 17% in the first year (Bloom et al., 2013). As a result, the researchers asked a crucial question: Why hadn't those companies adopted better management practices before?

One explanation is lack of know-how. A mechanism is needed to acquire or update knowledge systematically, ensure it is fit for purpose, and help an organization put that knowledge into practice—tantamount to getting smarter and working smarter.

Another explanation is natural skepticism, in that managers benefiting from the status quo may perceive little incentive to propagate change. So-called “best practices” can be polysemic to a workforce fearing a zero-sum game, so CEOs need a way to help their companies get smarter and work smarter without alienating the workforce. *Behavioral science* offers a way forward.

Behavioral science is the study of human perception, decision-making, and actions. The behavioral science mindset is about being evidence-based and willing to experiment, and it uses tools ranging from surveys to statistical analysis, to A/B testing to machine learning.² A timely example is a 2024 field study by Aidan Toner-Rodgers.

Case Study 1: Artificial Intelligence

Toner-Rodgers measured the impact of artificial intelligence on the R&D lab of a large U.S. company with 1,018 participating scientists. He found a 39% increase in patent filings and a 17% increase in product innovation from scientists aided by AI tools. Surprisingly, the bottom third of scientists did not show much improvement in their productivity, while the top scientists saw their output double.

Even more surprising was that 82% of all the scientists reported less satisfaction with their work due to decreased creativity and skill utilization. In particular, qualitative interviews found that nearly everyone involved felt insecure about whether AI

¹ <https://www.linkedin.com/in/mhua/>

² A company has a *competitive advantage* when it has something that enables it to outperform its competitors. One such example is a lower cost structure. That advantage is durable when a competitor cannot replicate or nullify it within three years.




Figure 1: An original comic strip for this article by Michelle Canete. Developed with ChatGPT within 10 minutes, using one prompt to generate the illustration and a second prompt for the text. Without AI, this development would have likely taken hours and incurred a financial cost.

would get the credit for new discoveries (Toner-Rodgers, 2024). Here, we have a good example of how a new change can improve productivity and, at the same time, how humans experience that change.

Publicly available insights, such as what Toner-Rodgers provided, can help a CEO look around corners, but it raises questions: “Do those insights hold true here?” and “How can we benefit best from

those insights?” This is where a behavioral science capability can (Step 1) validate new ideas and insights for true fit; (Step 2) generate new ideas to help a company outperform; (Step 3) test those ideas for efficacy and unintended consequences; and (Step 4) fine-tune them for maximum value. The table below summarizes what happened at this large U.S. company.

Table 1: Mapping of Value From One Field Experiment

Behavioral Science Steps	Evidence Based Discovery	Business Value
1. Validate ideas and insights for true fit.	AI applications improved productivity by enabling faster R&C cycle time.	39% increase in patent filings. 17% increase in product innovation.
2. Generate new ideas.	None.	None.
3. Tests for efficacy and unintended consequence.	A third of participants did not benefit. 82% felt less job satisfaction.	Opportunity for management to see around corners and make adjustments.
4. Fine tune for maximum value.	None.	None.

As you can see from the above chart, a single field experiment—conducted by an academic practicing a form of behavioral science—helped demonstrate the value of AI applications and provided management with insights into unintended consequences. It is also

evident that the work was limited in scope. A hand-off from Toner-Rodgers to an embedded behavioral science team could generate additional ideas or insights and fine-tune the overall change to maximize value. Advice to CEOs: Don’t leave value on the table.

Case Study 2: Hiring the Right Talent

Having the right talent can be the difference between a company's success, mediocrity, or failure.³ Consistently making better hiring decisions than your competitors creates an advantage in areas such as innovation, product development, sales, customer service, and operations.

Amazon is one company that has excelled at talent selection. Every hiring decision is made through the consensus of three to five interviewers trained in behavioral interviewing, bias resistance, and legal compliance. To allow for deeper probing, each interviewer is assigned specific areas to focus on. To avoid groupthink, interviewers are instructed not to discuss the candidate until the decision meeting and cannot see the votes of others until they have submitted their own vote (Amazon, 2021).

According to Amazon's written standards, only bar-raising talent can be hired. A trained bar-raiser outside the hiring manager's chain of command interviews all the final candidates and facilitates the decision meeting. The bar-raiser has veto authority over a decision to hire, which has happened in the hiring process.

Amazon defines a "bar-raising talent" as someone with the potential to be stronger than 50% of incumbents doing the same job. That standard is represented by the equation below

$$H_n = 1.5 * M_{n-1}$$

where H_n represents the next hire, 1.5 is the requirement of a 50% improvement in talent quality, and M_{n-1} is the median performance quality of all previous hires. Translating a concept into a symbolic equation is helpful because it encourages pressure testing of abstract ideas—especially those so intuitively attractive that we want them to be true. Meanwhile, the not-so-logarithmic equation below represents how incumbents at the company felt with each successive hire said to be 50% better than them:

: = (

Between 2004 and 2024, Amazon made job offers to over a million people. Such immense growth put a lot of pressure on the various talent pools from which the company hired. The equation below represents the likelihood of actually hiring bar-raising talent H_n :

$$H_n \propto \frac{D \cdot Ofn \cdot Uv}{P_{better} \cdot (Pt + Pg)}$$

This footnote explains the variables in the above equation.⁴ Thinking through the equation shows that hiring a bar-raising talent with each successive job opening *soon* becomes impractical and *eventually* becomes impossible.

Amazon's hiring mechanism started near the turn of the century and should be applauded. A quarter century later, it serves as a starting point rather than the final model to copy. How can hiring accuracy be improved? Are some people better at evaluating talent than others? Can machine learning augment human decision-makers?

The core mission of an interview is to predict which candidate would be the best hire, casually defined as the one most likely to perform job duties as expected and work well with others on the team. More than 60 years of research has shown that traditional interviews have weak predictive validity ($r = 0.20 - 0.38$), equivalent to being wrong about a third of the time (McDaniel et al., 1994). A one-third error rate on such a consequential decision is a lot of lost value.

Still, the predominant practice is to screen resumes, do a series of one-on-one interviews, and hope the hiring team chooses the best person. That traditional approach does pretty well in screening for apparent factors, but it does not provide consistent accuracy in discerning between, for example, the best candidate for the job and the best interviewer.

In a landmark study, McDaniel et al. (1994) analyzed 245 correlation coefficients across 86,311 individuals and concluded that structured—rather than

3 The "right talent" should be thought of in the context of the job requirements and operating culture, and not necessarily the person with the perfect resume or the best interview performance.

4 H_n increases in line with higher offer declines (D), more false negatives (Ofn), more significant inefficiency from false positives (Ofp), and an increased number of people taking themselves out of the talent pool by going to another employer, retirement, a career change, or other reasons (Uv). Conversely, H_n decreases when the proportion of suitable 'bar-raising' talent (P_{better}) is higher and when the overall talent pool size is larger, combining existing talent (Pt) and pool growth (Pg).

unstructured—interviews have better predictive validity. A structured interview uses a predetermined set of questions while an unstructured interview is a more free-flowing conversation where questions emerge spontaneously. Behavioral interviewing is the most well-known form of structured interview.

Behavioral interviewing was quickly adopted by the human resources industry as a promising solution—so much so that it became dogma in some recruiting circles. But then, in 2013, Oh et al. ran an updated meta-analysis that controlled for variables such as publication bias and overlapping confidence intervals. The researchers wrote, ‘It can be safely concluded that unstructured interviews may be as valid as structured interviews in most cases’.

So, which is the best path forward? Neither in isolation. Structured interviews provide better consistency, but not at the cost of being robotic or predictable. When interviews feel robotic, the hiring mechanism loses valuable benefits such as the chance to build rapport and receive unprepared responses.

One straightforward but underutilized technique is work sampling. Reviewing a work sample provides as much, if not more, predictive validity ($r = 0.33 - 0.54$) than interviews (Roth et al., 2005; Schmidt & Hunter, 1998). This makes intuitive sense: If you’re hiring someone to develop, for example, a cash flow

model, why not ask them to build a model based on publicly available data? It would seem that what a job candidate can do is more important than how well they interview.

Hiring techniques are the *how*; the people that make the decision are the *who*. In a literature review, De Kock et al. (2020) sought to identify the characteristics that make individuals better assessors. Their discoveries were surprising, in that general intelligence was the most consistent predictor of judgment accuracy, but the most intelligent people are usually less accurate than moderately intelligent people. Conventional wisdom says to let the smartest people make the decisions, but science says that this could be a mistake.

Advice to CEOs: Your hiring managers need help, and that help may come from machines.

In a well-designed field experiment, Bo Cowgill provided empirical evidence that algorithms are not inherently fair or unfair but are context-dependent. He argued that algorithms must be carefully measured against the human baseline they augment or replace. For example, in Cowgill’s double-blind study, an algorithm was able to predict better which job candidate would receive an offer (+14%), accept the offer (+18%), and be more productive once hired (0.2 to 0.4 standard deviation). Interestingly, the algo




Figure 2: A model of individual differences in judgment and rating accuracy. Adapted from “The profile of the ‘Good Judge’ in HRM: A systematic review and agenda for future research,” by De Kock et al. (2020).

selected more candidates at the margin—those who tend not to be picked by human assessors—hinting at less bias and foreshadowing an ability to improve inclusive hiring outcomes (Cowgill, 2019).

Case Study 3: Rewarding Top Performers

At a large investment bank, a recurring business problem was figuring out how to distribute their bonus pool in a manner that retained and motivated top performers. The prevailing management practice was to allocate a budget and instruct managers to stay within it, but this did little to optimize bonus allocation, left the managers feeling unsupported, and created attrition risk with top performers.

According to prospect theory, people evaluate potential outcomes relative to a reference point and experience that outcome as either a gain or a loss (Kahneman & Tversky, 1979). Consistent with prospect theory, I have observed people generally evaluate their bonus based on raw size ('Does this feel enough?'), how it compares to the previous year's payout ('Did this get better?'), and how it compares to what others received ('Is this fair?').

On top of prospect theory sits another layer of psychology: loss aversion. People will experience stronger emotional reactions to perceived losses than to gains, even when the gains and losses are equal in quantum. For example, one might experience 15 watts of positive emotions after winning a \$50 bet and 25 watts of negative emotions after losing a \$50 bet. For now, the watts are metaphorical and not a literal measurement of emotions—which science can't quite do yet but may one day with neuroimaging.

If a top performer is paid more than their reference point, they will feel a sense of gain, followed by emotions ranging from satisfaction to excitement. However, if paid less than their reference point, they will feel a loss followed by emotions ranging from disappointment to anger. The latter is more consequential because a loss evokes a higher emotional amplitude than a gain. Most managers intuitively try to minimize the number of people that will be upset and gravitate towards the $1/N$ heuristic, a tendency to distribute compensation evenly across the team (Gigerenzer et al., 2022). The problem is that top performers want to be paid commensurate with their results.

To exemplify this point, I worked with a small team that wanted to improve the process and convinced upper management to approve a pilot program. We listened to qualitative feedback on existing talent, analyzed performance scorecards, and reviewed talent movement history. We made two key findings. First, there were meaningful differences in performance amongst the bankers, even after controlling for variances in their portfolios. Second, in past years, the correlation between performance and bonus variance existed but was too small to move the needle on motivation.

To solve this problem, we designed and implemented a match-funding feature within the annual bonus program. The idea was simple: Reward managers with three dollars for every one dollar not given to a below-median performer. For example, if the bonus for a below-median performer were reduced from \$100,000 to \$80,000, the manager would receive \$60,000 back into their bonus pool to allocate to the rest of the team. The extra \$40,000 was from an initial hold-back of the top-level budget prior to allocation to various teams.

What were the results? Millions of dollars of additional pay differentiation (the specific number is confidential) was unlocked at no extra cost to shareholders. Managers liked the program because they preferred being incentivized to differentiate, whilst top performers were happy to be paid for their performance and therefore stayed.

Advice to CEOs: Behavioral science can help you make sure the people who spend limited money are spending it wisely.




Figure 3: Gemini's interpretation of loss aversion theory.

Transforming Your Organization

Asking people to be more evidence-based and reconsider their beliefs is fighting human nature. When entrenched habits and personal self-interest are present, the best way for a CEO to transform an organization is to flood the zone: Go all in, all at once, and overshoot.⁵

Jack Welch, the former Chairman and CEO of General Electric, pushed Six Sigma into the culture of what was then already the most admired company in the world. Six Sigma is a quality control methodology that requires people to think differently, be evidence-based, and use statistical tools. Welch decided the best way to move the 112-year-old conglomerate was to go all in, all at once, and overshoot.

For example, while most companies would roll out quality control training only to teams that absolutely needed it, Welch declared that every white-collar employee must be Six Sigma-certified, which meant passing an exam and doing a quality control project. Those who did not comply could lose their job. Six Sigma certification was hard-coded into the agenda of GE's talent reviews for over a decade.

The Six Sigma certification hierarchy is Green Belt, Black Belt, and Master Black Belt. Achieving Black or Master Black Belt was considered a career accelerator. Welch constantly boasted about Six Sigma in his annual letter to shareholders and cited \$12 billion worth of savings over five years.

Welch set expectations higher than necessary, talked about the initiative repeatedly, and made clear the consequences for not complying—but presented abundant incentives for getting on board. The carrot-to-stick ratio felt 5-to-1, so this came across more like a movement within the company than bullying.

Flooding the zone need not be cruel or bombastic. Explain what change is happening, and why, and be a reliable source of information. Encourage influencers to tweak certain aspects of how the change shows up, but set boundaries and make bold decisions. Take the time to show gratitude. If anyone chooses to step off the train, treat them with dignity.

Advice to CEOs: You may want to flood the zone like some business and political leaders are doing today, but do so in the context of your values.

Final Thoughts

'A long-standing question is whether differences in management practices across firms can explain differences in productivity', wrote Nicholas Bloom.

The answer is yes.

In the first case study, we learned how a single field study can reveal important insights that directly impact business outcomes. We also saw that there was so much more to do and wished the company had a team to catch and continue the work. In the second case study, we learned about how the hiring managers and bar-raisers at Amazon helped build one of the most successful companies in history. We also saw opportunities for further improvement. In the third case study, we learned how shareholders and high performers at a bank were not getting maximum value for each dollar of bonus money distributed, and then how incentivizing managers to differentiate could lead to better outcomes.

In each of the studies, we saw how an evidenced-based mindset and willingness to experiment could lead to concrete business value. Behavioral science can be done in a centralized or decentralized manner. At its core, behavioral science can help leaders (1) generate new ideas to help a company outperform; (2) validate that those ideas can really work; (3) test those ideas for efficacy and unintended consequences; and (4) fine-tune them for maximum value.

THE AUTHOR

Minh Hua is a seasoned Chief People Officer, known for his practical application of behavioral economics to shape high-performing organizations. With over 25 years of leadership experience across technology and financial services, he has held key roles in private equity and at public companies such as Amazon Web Services (AWS), J.P. Morgan, GE, and Stanley Black & Decker (SBD). While at SBD, Minh worked with the Board of Directors on human capital priorities and founded the first People Analytics function for the company. At AWS, Minh was the go-to speaker for customers wanting to learn about the company's culture and people-practices. His ability to harness

⁵ The term "flooding the zone" was recently popularized by Steve Bannon, a political commentator. However, it's a change management tactic that can be found throughout history. The primary characteristic is making all the desired changes in rapid succession and starting with a larger change than needed.

behavioral insights from academia and commercial firms has made him a highly sought-after advisor and speaker.

REFERENCES

- Amazon. (2021). *Make great hiring decisions* [Unpublished internal training material]. Amazon.
- Bloom, N., Eifert, B., Mahajan, A., McKenzie, D., & Roberts, J. (2013). Does management matter? Evidence from India. *The Quarterly Journal of Economics*, 128(1), 1–51. <https://doi.org/10.1093/qje/qjs044>
- Bloom, N., & Van Reenen, J. (2007). Measuring and explaining management practices across firms and countries. *The Quarterly Journal of Economics*, 122(4), 1351–1408. <https://doi.org/10.1162/qjec.2007.122.4.1351>
- Cowgill, B. (2019). *Bias and productivity in humans and machines* (Upjohn Institute Working Paper 19–309). W.E. Upjohn Institute for Employment Research. <https://dx.doi.org/10.2139/ssrn.3433737>
- De Kock, F. S., Lievens, F., & Born, M. P., (2020). The profile of the “Good Judge” in HRM: A systematic review and agenda for future research. *Human Resource Management Review*, 30(2), 100667. <https://doi.org/10.1016/j.hrmr.2018.09.003>
- Gigerenzer, G., Reb, J., & Luan, S. (2022). Smart heuristics for individuals, teams, and organizations. *Annual Review of Organizational Psychology and Organizational Behavior*, 9(1), 171–198. <https://doi.org/10.1146/annurev-orgpsych-012420-090506>
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decisions under risk. *Econometrica*, 47(2), 263–292. <https://doi.org/10.2307/1914185>
- McDaniel, M. A., Whetzel, D. L., Schmidt, F. L., & Maurer, S. D. (1994). The validity of employment interviews: A comprehensive review and meta-analysis. *Journal of Applied Psychology*, 79(4), 599–616. <https://psycnet.apa.org/doi/10.1037/0021-9010.79.4.599>
- Oh, I.-S., Postlethwaite, B. E., & Schmidt, F. L. (2013). Rethinking the validity of interviews for employment decision making: Implications of recent developments in meta-analysis. In D. J. Svyantek & K. Mahoney (Eds.), *Received wisdom, kernels of truth, and boundary conditions in organizational studies* (pp. 297–329).
- Roth, P. L., Bobko, P., & McFarland, L. A. (2005). A meta-analysis of work sample test validity: Updating and integrating some classic literature. *Personnel Psychology*, 58(4), 1009–1037. <https://onlinelibrary.wiley.com/doi/10.1111/j.1744-6570.2005.00714.x>
- Schmidt, F. L., & Hunter, J. E. (1998). The validity and unity of selection methods in personnel psychology: Practical and theoretical implications of 85 years of research findings. *Psychological Bulletin*, 124(2), 262–274. <https://doi.org/10.1037/0033-2909.124.2.262>
- Syverson, C. (2004). Market structure and productivity: A concrete example. *Journal of Political Economy*, 112(6), 1181–1222. <https://doi.org/10.1086/424743>
- Toner-Rodgers, A. (2024). *Artificial intelligence, scientific discovery, and product innovation*. Massachusetts Institute of Technology. <https://doi.org/10.48550/arXiv.2412.17866>

Three thick, white, parallel diagonal lines slanting from the top-left to the bottom-right, positioned behind the word 'RESOURCES'.

RESOURCES

Postgraduate Programs



CITY
ST GEORGE'S
UNIVERSITY OF LONDON



Study MSc Behavioural Economics in central London.

Taught by experts in
the fields of Psychology,
Economics and Neuroscience.

School of Health
& Medical Sciences

citystgeorges.ac.uk

Course content

This MSc in Behavioural Economics will help you develop the analytical skills and the specialist knowledge required to understand human decision making.

The programme is taught by academic staff from the Departments of Psychology and Neuroscience, and Economics, with a strong emphasis on theoretical, methodological and practical aspects.

You will learn how people and institutions make decisions under conditions of constraint, such as time and knowledge, and under various influences such as social pressure. There are many practical implications of insights from behavioural economics and the field is widely accepted to deliver a powerful, cost-effective approach to improving human welfare.

Course structure

The MSc Behavioural Economics can be studied full-time or part-time.

Full-time students take four modules in each of the first two terms that can be tailored to your own interests and skills. Topics include research methods and statistics as well as modules focused on choice-behaviour and decision-making. This is followed by a written research dissertation, starting in the third term. Part-time students study all of their modules on the same weekday.

Benefits of studying at City St George's, University of London


- **Ideally situated in the heart of London;** providing valuable industry networking and employment opportunities.
- **Taught by experts in the fields of Psychology, Economics and Neuroscience;** find out more about our programme team below.
- **Networking opportunities with practitioners and experts;** our guest speakers seminar series hosts practitioners, academics and experts from across the field.
- **Emphasis on transparent and open science practices;** Students learn how to develop robust and replicable findings, working with real data.
- **Join a diverse student community;** our students come from all over the world, with a diverse range of educational and work experiences.

Career opportunities

Graduates are able to progress in careers such as:

- Economic consultants
- Public policy specialists
- Political campaign and public relations specialists
- Digital design consultants
- Behavioural science consultants
- Marketing specialists.

Some graduates also choose to pursue PhDs in economics or psychology.



Meet our programme team

Our team includes academic staff from our Department of Psychology and Neuroscience, and our Department of Economics. We have experts with the following expertise and experience:

- Game theory and experimental economics
- Computational modelling of cognition and decision-making
- Cognitive psychology and neuroscience
- Social psychology and individual differences
- Physics and philosophy

Our staff's research focuses on topics such as:

- Subjective well-being
- Perception of fairness
- Altruism and prosociality
- Decision-making under scarcity
- Intertemporal discounting
- Risk and uncertainty
- Learning and categorisation
- Female labour force participation and life course studies
- Developing and improving statistical models and data analysis methods for neuroimaging (EEG, fMRI) and behavioural data.
- Policy implications that decision making processes and biases may have
- Asset pricing and financial econometrics
- Labour economics, urban economics and local economic development
- Applied microeconomics and experimental economics.

Some of our programme team are also involved in consultancy work for big corporations such as the BBC.



Find out more
www.citystgeorges.ac.uk/behavioural-economics



Dr Claudia Civaï,
Programme Co-Director
and Senior Lecturer in
Psychology.

Claudia is an experimental psychologist and cognitive neuroscientist. She has particular expertise in social decision-making, especially the perception of unfairness and injustice, which she investigates by merging experimental game theory and cognitive neuroscience theoretical and methodological frameworks.

“The field of behavioural economics and, more broadly, behavioural science is constantly evolving; this MSc programme gives a solid grounding on the foundation of the discipline, as well as an up-to-date insight into the ample spectrum of applications thanks to the contribution of practitioners through our workshop and seminar series.”



Dr Sergiu Ungureanu,
Programme Co-Director
and Lecturer in
Behavioural Economics.

Sergiu is a theoretical and behavioural economist by training. He is interested in using both modelling and statistical analysis to look at interesting questions regarding decision-making behaviour. His research is in understanding the way in which behavioural effects (e.g. heuristics and biases) can be unified, or can be understood as results of deeper causes like bounded rationality.

“This programme is a great way to be introduced to the field of human decision-making. You will learn both the common theories we use to understand decisions, and the skills you need to critically analyse research, and also practical applications in the field – e.g. marketing approaches, choice architecture for consumers, etc.”



Department of
**Psychological and
Behavioural Science**

EXECUTIVE MSc
BEHAVIOURAL SCIENCE

UNCOVER THE SCIENCE BEHIND BEHAVIOUR

An increasing number of organisations now engage with the idea of applying behavioural insights to their organisational challenges.

The Executive MSc Behavioural Science, based in LSE's Department of Psychological and Behavioural Science, is taught by experts at the forefront of international research in behavioural science.

Our programme provides rigorous training for professionals who are seeking to expand their knowledge in this emerging and exciting field. Many of our alumni are now prominent behavioural science leaders and experts in a range of organisations around the world.



CONTACT US

pbs.emsc@lse.ac.uk

For more information, please visit lse.ac.uk/EMScBehaviouralScience



Department of
**Psychological and
Behavioural Science**



EXECUTIVE MSc BEHAVIOURAL SCIENCE

A UNIQUE AND DYNAMIC PROGRAMME FOR PROFESSIONALS

LSE's Executive MSc Behavioural Science is taught by specialists at the forefront of international research in behavioural science. Our programme provides the opportunity for full-time professionals working in any sector to obtain a graduate qualification in behavioural science, allowing you to pursue new and expanded opportunities within this emerging and exciting field.

The programme starts in September each year with teaching being delivered during three two-week intensive teaching blocks at the LSE campus in London. You are not required to be in attendance on campus outside of these weeks and can therefore continue to live and work outside of London and the UK. Between teaching sessions you work independently on various assignments. After the final teaching session you complete a dissertation on a topic of your choice with support from your supervisor.

The programme includes unique and innovative modules such as:

- Behavioural Science and Policy
- Behavioural Decision Science
- Research Methods for Behavioural Science
- Frontiers in Behavioural Science Methods
- Policy Appraisal and Ethics
- Behavioural Science in an Age of New Technology
- Corporate Behaviour and Decision Making
- Organisational Culture

OUR STUDENTS

Our students come from a wide range of academic and professional backgrounds from all over the world, but one thing binds them together: a passion for behavioural science and a desire to better understand how principles from behavioural science can be applied in their professional (and personal) lives.

CONTACT US

pbs.emsc@lse.ac.uk

For more information, please visit lse.ac.uk/EMScBehaviouralScience



Department of
**Psychological and
Behavioural Science**

WHAT OUR ALUMNI HAVE TO SAY ABOUT THE PROGRAMME



LSE's Executive MSc Behavioural Science is second to none in providing a complete insight into contemporary behavioural science debate and methodology, delivered by world-class experts. ”

Ana, 2021 graduate



The EMSc struck the perfect balance between teaching academic rigour and practical implementation, giving me solid foundations for a complete career change. ”

Nitish, 2022 graduate



The Executive MSc Behavioural Science has equipped me with tools to address some of the most pressing challenges with strong behavioural roots in the MENA region and the Global South. ”

Nabil, 2020 graduate



The network built during the EMSc is unmatched by any past professional or educational experience I've had, through faculty support, alumni connections, and lifelong professional and personal relationships. ”

Madeline, 2019 graduate

CONTACT US

pbs.emsc@lse.ac.uk

For more information, please visit lse.ac.uk/EMScBehaviouralScience

MA IN BEHAVIORAL ECONOMICS

DISCOVER THE SCIENCE BEHIND DECISION-MAKING

Join a world-class MA program at the intersection of Psychology and Economics

Why our program?

→ Interdisciplinary Curriculum

A unique combination of psychology, economics, policy, and statistical methods

→ Tailored Career Tracks

Choose between research (thesis) & applied (project) specializations

→ Global & Diverse Community

Study in Israel, one of the world's leading BE hubs, with peers from around the world



 "I learned from top researchers and passionate peers. This program shaped my UX career."

Yeva Lalayan

UXR, Infineon Technologies AG

Ready to shape the future?

www.runi.ac.il/en/schools/rris/graduate/behavioral-economics/

MA IN BEHAVIORAL ECONOMICS

**UNDERSTAND BEHAVIOR.
INFLUENCE DECISIONS.
LEAD CHANGE.**



“Studying at Reichman expanded my perspective and pushed me toward new horizons.”

Nick Mishkin

Customer Operations Professional



“The MABE program helped me turn curiosity into a career in behavioral economics.”

Lior Shem-Tov

Manager, IDeCision PhD Student,
Reichman University

The MA program in Behavioral Economics at Reichman University is designed for those passionate about understanding how people make choices, and how to help improve them. Our interdisciplinary program is a full collaboration between the Baruch Ivcher School of Psychology and the Tiomkin School of Economics, offering a rare blend of psychological insight and economic analysis. Our students gain strong theoretical depth and practical tools in a research-oriented, application-focused academic environment, led by world-renowned researchers and experienced practitioners.

Ready to shape the future?

www.runi.ac.il/en/schools/rris/graduate/behavioral-economics/

LIVE IN ISRAEL,
STUDY IN ENGLISH

REICHMAN UNIVERSITY FAST FACTS

Over **2,500** Raphael Recanati International School students from over **90** countries make up 1/3 of the Reichman University student body



Israel's first and only  **private university**

Over **1,200**  **worldwide companies** founded by Reichman University alumni

Reichman University:  **the most international university in Israel**

Over **48**  **research institutes** at Reichman University

Over **38,000**  **graduates worldwide**

The Israeli Council for Higher Education granted Reichman University permission to confer **doctoral degrees** 

GRADUATE PROGRAMS

- MA in Counter-Terrorism & Intelligence Studies
- MA in Diplomacy & Conflict Studies
- MA in Financial Economics
- MA in Human-Computer Interaction
- MSc in Machine Learning & Data Science
- One-Year MBA
- Global MBA AI & Big Data
- Global MBA Innovation & Entrepreneurship
- MBA in Healthcare Innovation
- MA in Behavioral Economics

For more information:

 +972 9 9602700  +972 54 5628121  rris.master@runi.ac.il



MSc in Behavioural Economics at University College Dublin

UCD School of Economics is Ireland's leading Economics department, and our MSc in Behavioural Economics is taught by a vibrant team of behavioural researchers and practitioners.

We provide in-depth training in the core concepts and theories of behavioural economics, with a focus on how to apply these tools in areas such as health, regulation, labour, environmental economics, and education. You can choose from a range of classes on empirical methodologies for conducting different types of controlled experiments (field, laboratory, and online).

You have the opportunity to design and conduct your own experiment, mastering the methods necessary for gathering causal evidence in a competitive world where distinguishing causation from correlation is crucial. You also have the opportunity to write a thesis or do a summer internship in an organisation of your choice.

Our MSc students are an integral part of the strong research culture in UCD and attend seminars by international researchers and workshops organised by the UCD Behavioural Science and Policy Group.



PROGRAM HIGHLIGHTS

- Rigorous training in core theories and methods for data gathering and analysis for fundamental research, as well as evidence-based policy.
- Strong policy focus and ample opportunity to interact with policymakers.
- Small class sizes and dedicated one-to-one supervision.
- Ireland's strategic location as a European hub for many large multinationals opens doors to global opportunities. Prominent employers include large tech, pharmaceutical and financial services companies.
- Career paths are diverse, including public policy, private sector and a PhD route into academia.
- Our graduates have found jobs in places such as: Behavioural Insights Team, Irish Government Economic and Evaluation Service, Grant Thornton, McKinsey, KPMG, Accenture, Novartis, Deloitte, Central Statistics Office, Central Bank of Ireland, European Central Bank, Economic and Social Research Institute, Competition and Consumer Protection Commission, AIB Bank, European Commission and other leading national and international organizations.

GRADUATE PROFILE: EMILY WHARTON-HOOD

It was an absolute honour and privilege to obtain my Masters in Behavioural Economics at UCD.

The MSc provides such a variety of modules while being taught by some of the best in the industry. Advancing my skills in Econometrics and Microeconomics was a highlight of the MSc, but also taking modules such as Energy Economics really broadened my perspective on where Economics is most impactful. Being able to run my own experiment allowed me to conduct my thesis in an area I was passionate about. Throughout the program, you become equipped with skills you would use in any industry you decide to move into. I am currently working as an Analyst at 2K Games Dublin!

I could not recommend this MSc program enough.



Contact us: economics@ucd.ie





MSc Economics: Behavioural Insights

Understand how people think. Shape how the world works.

The **MSc Economics: Behavioural Insights** at the University of Exeter combines cutting-edge economic theory with powerful behavioural science tools to help you understand, predict and influence decision-making in the real world.

Future-focused now.



University of Exeter
Business School

Contact us



Learn more and apply





Exeter: A hub for global behavioural economics research and collaboration

The University of Exeter's Department of Economics is considered home to one of the largest and most internationally recognised groups of behavioural and experimental economists in the world, with more than 25 faculty members who specialise in the field.

Our researchers have trained and worked at leading global institutions including Cambridge, Harvard, Melbourne, MIT, Monash, Oxford and Toronto, and bring a truly international

perspective to their work. Their expertise is recognised globally, working with the Behavioural Insights Team, the World Bank, government departments, and private organisations around the world.

We are also home to the FEELE Lab, a cutting-edge facility for experimental research, and host a regular research seminar series, featuring global thought leaders in behavioural economics.

“You’ll have support throughout your degree, working with an expert behavioural economist as your personal academic tutor.”

Dr Edwin Ip, Programme Director

A curriculum designed for impact

Gain deep theoretical knowledge and practical expertise through innovative modules, hands-on projects, and international opportunities:

- **Specialist modules:** Behavioural Finance, Experimental and Behavioural Economics, Optimisation Techniques and Study Design.
- **Capstone Applied Project:** Design your own real-world behavioural intervention with guidance from experts.
- **QTEM Programme:** Apply to join the QTEM network. Study abroad at a leading partner institution, and gain valuable experience through a data-focused international internship.
- **Optional modules:** Customise your degree to suit your career needs with a wide range of optional specialist modules covering areas such as environment and sustainability, international development, product innovation, and strategy.



University of Exeter
Business School

Contact us

Learn more and apply

Invest in your future

MSc Economics: Behavioural Insights equips you with the analytical thinking skills of an economist and the behavioural insight to understand and influence real-world decisions. You'll graduate with a strong set of in-demand skills, suitable for a range of careers in:

- Government and public policy
- Consulting and financial services
- NGOs and international organisations (e.g. UN, World Bank, OECD)
- Marketing, advertising, consumer analytics and product design
- Research and PhD study

With expert academic and careers support, you'll be prepared to lead in a world increasingly shaped by data and behaviour.



“Exeter truly offers more than just a degree; it’s a community that nurtures growth and fosters lifelong connections. The professors here are not just knowledgeable; they’re genuinely invested in our growth. They go the extra mile to ensure we grasp the concepts and are always there to lend a helping hand.”

Isha, University of Exeter Business School, Graduated 2024



University of Exeter
Business School

Contact us

Learn more and apply



University
of Glasgow

MSc Behavioural Science

Master Behaviour,
Shape Business,
Impact Society



This unique programme is a collaboration between the Adam Smith Business School and the School of Psychology & Neuroscience. It combines knowledge from Economics, Business, Finance, Psychology, and Neuroscience to develop future behavioural scientists. You will have access to cutting-edge methods, modern lab facilities, and strong industry partnerships.

Key Facts:

- **Duration:** 12 months full-time
- **Delivery:** In-person
- **Teaching starts:** September
- **Academic contact:** business-recruitment@glasgow.ac.uk



University
of Glasgow

Why choose MSc Behavioural Science at the Adam Smith Business School?

- **Multi-Disciplinary Approach:**

A unique decision-making programme combining insights from Economics, Management, Finance, Psychology, and Neuroscience.

- **Science and Business Tracks:**

A joint foundation for both tracks, ensuring consistency, with specialised skills based on your chosen track.

- **Theoretical & Practical Learning:**

Learn theory alongside quantitative, empirical, and market-driven research methods.

- **State-of-the-Art Lab Access:**

Use our new behavioural laboratory for real-world experiments, including dissertation research.

- **Industry Engagement:**

Take part in competitive, merit-based industry projects with real companies. Solve real business challenges, gain experience, and grow your professional network.

After graduation: What jobs can you get with this degree?

The programme will develop adaptable behavioural scientists with diverse skills, offering exceptional career opportunities across various sectors. You will be able to choose between career 2 tracks:


Science Track

Equips graduates with neuro-scientific skills and interdisciplinary expertise in Behavioural Science, ideal for pursuing a PhD or research roles in various sectors.

Business Track

Prepares graduates for careers in policy or business, focusing on decision science for practical applications in diverse industries and market systems.

Find out more and join us today:
glasgow.ac.uk/postgraduate/taught/behavioural-science



Penn
Master of Behavioral
and Decision Sciences



Learn the theory, apply the tools, and make a difference

Penn's Master of Behavioral and Decision Sciences (MBDS) program equips students with theoretical and practical tools to understand how individuals and groups make decisions, how to affect those decisions, and how social norms play a role in motivating and changing social behaviors. Led by world-renowned faculty, researchers, and practitioners, the MBDS program creates unique opportunities for students to engage with an exceptional advisory board, apply tools and knowledge in our annual Design Challenge, and pursue independent, cross-disciplinary research throughout Penn.

Meet our alumni:



"I loved learning tools and techniques to evaluate what people were saying and translate that into opportunities for the client. The [capstone] was a really exciting project. I never worked in the digital mental health space before. Having those opportunities to gain insight into different industries has helped me become a chameleon and learn to speak the languages of different clients."

Kathryn Ambroze, MBDS '22
Senior User Researcher, JPMorgan Chase & Co.



"When I was accepted into the MBDS, it was a meant-to-be moment. I felt like there was a link between me and the program. I could take the time to explore behavioral science in an academic setting. Some people love commercial spaces, some want to go into consulting, some people are really into research. I realized health and health outcomes are definitely what I'm interested in, personally and in my career."

Yuzhen (Valerie) Guo, MBDS '22
Behavioral Designer, Lirio, LLC



"One of the amazing things about Penn is that the faculty you work with are heavily involved in research—they're very much at the forefront of their field, so you can take part in a lot of research if you want to. Once I was done with the first project, there were other professors who needed help with different projects."

Max Spohn, MBDS '20
PhD candidate, Harvard Kennedy School of Government

Learn more about our engaged and well-connected alumni at

www.upenn.edu/mbds



Penn
Master of Behavioral
and Decision Sciences



Meet the Master of Behavioral and Decision Sciences program's founding director



Cristina Bicchieri

Founding Director, Master of Behavioral and Decision Sciences

S. J. Patterson Harvie Chair

Professor of Philosophy, Psychology and Legal Studies (Wharton)

*"Wherever there
is a human
group there are
social norms."*

-Cristina Bicchieri

Cristina Bicchieri is a world authority on social norms and has consulted with UNICEF, the World Bank, the Gates Foundation, the United Kingdom's Department for International Development, and many other organizations. She is the founder of the Master of Behavioral and Decision Sciences program and the Center for Social Norms and Behavioral Dynamics, a major research center at Penn that aims to support positive behaviors on a global scale. Cristina is the author of over 100 articles and seven books.

Unparalleled connections, exceptional opportunities

A defining feature of the University of Pennsylvania's Master of Behavioral and Decision Sciences program (MBDS) is its network of outstanding industry and research partners who help bring students exceptional practical experiences.

Meet the MBDS Advisory Board:

Charlotte Blank

*Director, Transformation & Analytics, Jaguar Land Rover
North America*

Claire Hobden

*Specialist on Vulnerable Workers, Domestic Work, International
Labour Organization*

Jeff Kreisler

*Head of Behavioral Science for JP Morgan Private Bank and
Founding Editor of PeopleScience.com*

Pavan Mamidi, PhD

*Director of the Centre for Social and Behaviour Change (CSBC),
Ashoka University*

Namika Sagara, PhD

*Co-Founder, Chief Behavioral Officer and Head of Consulting,
Syntoniq*

Neela Saldanha

*Executive Director, Yale Research Initiative on Innovation and Scale
(Y-RISE), Yale University*

Carlos Scartascini

*Head, Development Research Group, IDB Behavioral
Economics Group*

Greg Schwartz

Practice Lead – Healthcare Predictive Modeling, Deloitte Consulting

Piyush Tantia

Chief Innovation Officer, ideas42

Renos Vakis, PhD

Lead Economist, the Poverty and Equity Global Practice

Chiara Varazzani, PhD

*Lead Behavioral Scientist, Organisation for Economic Cooperation
and Development (OECD - OCDE)*

Scott Young

*Principal Advisor, Head of Private Sector, the Behavioural Insights
Team (BIT) North America*

Allison Zelkowitz

*Founder and Director, Center for Utilizing Behavioral Insights for
Children (CUBIC) at Save the Children International*

Learn more about our world-renowned faculty and researchers at:

www.upenn.edu/mbds



The MBDS Design Challenge

Every spring, the Master of Behavioral and Decision Sciences (MBDS) program organizes the Design Challenge, where our students partner with MBDS Industry Affiliates to apply cutting-edge knowledge from the fields of behavioral economics, decision sciences, network analysis, and public policy to solve real-life problems. We welcome world-leading clients in industries like health, wellness, sustainability, technology integration in marketing, and finance to collaborate with our students and provide guidance on solving the world's toughest challenges.

In the Design Challenge, MBDS students work to translate academic research, theoretical foundations, and applied frameworks into actionable insights toward a client-focused problem. At the end of the Challenge, students present their proposed solutions to the client's senior management and leadership.

The Design Challenge is an invaluable opportunity for our students to apply their MBDS education toward developing practical solutions while gaining real-world experience.



"Newristics has actually been a client for a couple of the Design Challenges. For many students who might be a little lighter on professional experience, the Design Challenge is a great way for them to talk about how they break down a challenge and use behavioral science to come up with a novel solution."

Michael Hayden II, MBDS '20

Consultant, MedTech & Applied Behavioral Insights

Learn more about how MBDS connects students and industry at:

www.upenn.edu/mbds



MASTER OF SCIENCE IN APPLIED ECONOMICS



UNIVERSITY OF SAN FRANCISCO

STUDY APPLIED ECONOMICS IN SAN FRANCISCO!

In the San Francisco Bay Area, new technologies are altering the way people live and work throughout the world. New marketplaces and platforms created here are fundamentally changing interactions between consumers, businesses, nonprofits and governments. As a next-generation economist with an M.S. in Applied Economics, you will gain the skills to take a leadership role in this dynamic environment.

PROGRAM HIGHLIGHTS:

- › Economics for the digitized age: Study market design, reputational systems, auction theory, pricing, behavioral economics and other concepts essential to understanding and shaping the new economy.
- › Practical programming and data skills: Learn to work with your data using major programming languages Python and R from your first semester and build these skills throughout the program.
- › Data analysis: Courses in experimental design, machine learning and econometrics take you beyond description and correlation to understand causal processes and economic mechanisms.
- › San Francisco Advantage: San Francisco's tech firms are leading the way in digitizing the global economy. Studying at USF positions you to take internships and jobs with global technology giants — or with startups on the way to becoming the next big thing. Graduates have achieved success in their careers, joining and leading teams at companies including Google, Intel, Apple, Amazon, Realtor.com, Morgan Stanley, Visa, Citi, Bloomberg, Disney, and more.
- › MSAE is a designated STEM program allowing eligible international students to work in the U.S. for three years after graduation under Optional Practical Training (OPT).

GRADUATES OF APPLIED ECONOMICS:

- › Analyze massive and often unstructured data sets.
- › Design research to draw causal inferences about consumer behavior and market structure.
- › Create new markets, platforms and reputation systems.
- › Optimize pricing, advertising, and investments.
- › Develop public policies that adapt to, and shape, the impact of new technologies.
- › Help non-profits and NGOs take advantage of opportunities that digitization creates.

CONTACT US



Take a snapshot of this information so that you can contact us at econ@usfca.edu or at 415.422.2711. Learn more at usfca.edu/economics



VISIT OUR
LINKEDIN PAGE!



MASTER OF SCIENCE IN DEVELOPMENT ECONOMICS



UNIVERSITY OF SAN FRANCISCO

DEVELOPMENT ECONOMICS PROGRAM IN SAN FRANCISCO

The rapid pace of globalization has increased the demand for professionals with training in international economics and economic development. Our one of a kind Master of Science in International and Development Economics (MS IDEC) provides students with the knowledge and skills to understand how market forces can be harnessed to empower developing countries to break from cycles of poverty.

PROGRAM HIGHLIGHTS:

- › Leading edge analytics: Study the latest programming languages, data acquisition and analytical techniques.
- › Learn the detailed econometric techniques practitioners use to analyze and enhance their programs.
- › Women and Development. Study the dramatic impact of programs that focus on women and their contributions to economic growth.
- › Climate change, health and environmental policy in developing countries: Learn about the impact of climate change in developing economies.
- › Study the methods of behavioral economics.
- › Overseas field study: Travel to a developing country to pursue your research interests with guidance and advice from award winning senior faculty.
- › Original research thesis and oral defense. You will have the opportunity present your research to faculty and students.
- › IDEC is a designated STEM program allowing eligible international students to work in the U.S. for three years after graduation under Optional Practical Training (OPT).

GRADUATES OF THE MS IDEC PROGRAM SERVE AS:

- › Leaders of intervention and research teams at the World Bank and regional development banks.
- › Senior managers of NGOs working in international development.
- › Data scientists at international social media companies.
- › Professors of Economics.

CONTACT US



Take a snapshot of this information so that you can contact us at idec@usfca.edu or at 415.422.2711. Learn more at usfca.edu/idec



VISIT OUR
LINKEDIN PAGE!

MSc in — **Behavioural and Economic Science**

Do you want to understand the choices people make, why they make them and what influences their behaviour?

The MSc in Behavioural Science and Economics at the University of Warwick combines multidisciplinary expertise from the departments of Psychology, Economics and Warwick Business School that is crucial to understand how influencing people's choices impacts a variety of sectors and industries.

The MSc suits those with a quantitative background (e.g. mathematics, sciences, economics, psychology).



Further details:

psychologyPG@warwick.ac.uk | +44 (0)24765 23096

www.warwick.ac.uk/bes

Why — study Behavioural and Economic Science

This innovative, interdisciplinary programme combines decision science and behavioural economics.

You will learn theory and real-world applications of behavioural economics and the cognitive science of judgement and decision making.

For those looking to pursue careers in business, public policy implementation or research, there are three core modules, a wide variety of optional modules to suit your interests and career goals and a research project.

- Modules span across the departments of Psychology, Economics and Warwick Business School, providing a thorough grounding of both the theory and real-world application of behavioural science.
- Modules on the design, conduction and analysis of behavioural experiments and the analysis of large-scale datasets.
- An empirical research project.

Our students have gone on to take positions at The Busara Center for Behavioral Economics, The UK Behavioural Insights Team, Google, Frontier Economics, Facebook, Ogilvy Change and more.



Further details:

psychologyPG@warwick.ac.uk | +44 (0)24765 23096

www.warwick.ac.uk/bes

Why — Warwick?

You will be taught by internationally recognised, world-leading researchers in the departments of Psychology, Economics and Warwick Business School.

We also have cutting-edge technology and laboratory facilities for conducting your behavioural research.

Warwick is consistently ranked highly, placing 5th in the UK for Economics (*The QS World University Rankings by Subject 2024*) and we are the 6th most targeted university by the UK's top 100 graduate employers (*The Graduate Market in 2024, High Fliers Research Ltd*). Behavioural Science was identified as an area of significant academic achievement in the Research Excellence Framework.

By studying at Warwick, you will be part of a global community of students from all over the world from diverse backgrounds. With students from South America, Asia, Europe, the USA and the Middle East, our supportive and inclusive community will enable you to get the most out of your studies.

**Find out more about
Postgraduate Study at Warwick**

www.warwick.ac.uk/postgraduate



Further details:

psychologyPG@warwick.ac.uk | +44 (0)24765 23096

www.warwick.ac.uk/bes

Other Resources

For the most up-to-date behavioral
science resources, please visit



BehavioralEconomics.com

Postgraduate Programs (complete listing)
LinkedIn Group
Job Board
Encyclopedia
Books
and more

APPENDIX

Author Profiles

Ernst Fehr (Introduction)



Ernst Fehr is Professor of Microeconomics and Experimental Economics at the University of Zurich and Director of the UBS Center for Economics in Society. His research examines the role of

social preferences and social norms for incentives, market behavior, and organizational outcomes. He has also contributed to the fields of neuroeconomics and behavioral finance. Ernst has published in leading

academic journals across economics and interdisciplinary science, and his work is widely cited. He is a member of the American Academy of Arts and Sciences, the John Kenneth Galbraith Fellow of the American Academy of Political and Social Sciences, and a Foreign Honorary Member of the American Economic Association. He also served as a president of the Economic Science Association and the European Economic Association. He co-founded FehrAdvice & Partners AG, a consultancy systematically applying behavioral economics to business and policy.

Suzanne B. Shu (Guest Editorial)



Suzanne B. Shu is the Dean of Faculty and Research for the Cornell University SC Johnson College of Business and the John S. Dyson Professor in Marketing at the Dyson School of Applied Economics and

Management. The types of decisions analyzed in her research include consumer self-control problems and consumption timing issues, with important

implications for both negative behaviors (such as procrastination) and positive behaviors (such as saving). Her work on financial decisions has focused specifically on decumulation during retirement (annuities, Social Security claiming) as well as on perceived fairness for financial products. She is also an NBER Research Associate and has served on the Academic Research Council for the Consumer Financial Protection Bureau and the Netspar Scientific Council.

Geoffrey Fisher (Guest Editorial)



Geoffrey Fisher is an Associate Professor of Marketing at Cornell University's SC Johnson College of Business. His research bridges marketing, behavioral economics, and

neuroeconomics, with a focus on how people make multi-attribute decisions. His current work integrates process-tracing methods into models of decision-making, revealing how cognitive and attentional dynamics shape consumer behavior.

David Just (Guest Editorial)



David Just is the Susan Eckert Lynch Professor in Science and Business in the Charles H. Dyson School of Applied Economics and Management at Cornell University and a fellow of the

Agricultural and Applied Economics Association. His research applies behavioral economics to policy issues in food and agriculture. Much of his recent research examines issues surrounding food insecurity and consumer behavior related to private and public food assistance programs.

Stephen Shu (Guest Editorial)



Stephen Shu is a Professor of Practice of Behavioral Economics at Cornell University, Dyson School of Applied Economics and Management. He also serves as the Academic Director for the Master of Professional Studies (MPS) and CEMS Programs at

Cornell. He has more than three decades of industry experience and is perhaps most recognized for helping to incubate some of the very first behavioral economics initiatives around the world in the commercial sector. His expertise is in the area of behavioral finance, and he has worked with numerous companies in retirement, wealth, investment, banking, advisory, and FinTech spaces.

Alain Samson (Editor)



Alain Samson is the editor of the Behavioral Economics Guide and founder of [BehavioralEconomics.com](https://behavioraleconomics.com). He has worked as a consultant, researcher and scientific advisor. His experience spans multiple sectors, including finance, consumer goods, media, higher education, energy and government. He was previously Chief Science Officer at Behave Technologies (formerly Syntoniq). Alain studied at UC Berkeley, the University of Michigan and the London School of Economics, where

he obtained a PhD in Social Psychology. His scholarly interests have been eclectic, including culture and cognition, social perception, consumer psychology and behavioral economics. He has published articles in scholarly journals in the fields of management, consumer behavior and economic psychology. He is the author of [Consumed](#), a *Psychology Today* online popular science column about behavioral science.

Alain is a Founding Member of the Global Association of Applied Behavioural Scientists (GAABS).

alain@behavioraleconomics.com

Contributing Organizations

Allianz

The Allianz Group is one of the world's leading insurers and asset managers with around 125 million private and corporate customers in nearly 70 countries. Allianz customers benefit from a broad range of personal and corporate insurance services, ranging from property, life and health insurance to assistance services to credit insurance and global business

insurance. In fiscal 2023, over 157,000 employees achieved a total business volume of 161.7 billion euros and an operating profit of 14.7 billion euros for the group. Allianz SE, the parent company, is headquartered in Munich, Germany.

www.allianz.com

Behavioral Science Group (UAE)

The UAE's Behavioral Science Group (BSG) is a specialized unit within the Office of Development Affairs. Our purpose is to support the UAE government in achieving its policy objectives using behavioral science as a fresh lens for new solutions.

To do this, the BSG combines behavioral science expertise with a deep understanding of our local policy objectives and context. By mixing international expertise with national talent, we provide government partners with a novel offer. This blend enables us to

design and test practical and innovative solutions to a range of local challenges.

Our mission is to draw on key success stories from the world to replicate and build on them within our nation. Our unit works closely with other government entities, helping them to embed behavioral science in their policy agendas, while building their capacity in the discipline.

www.bsg.ae

Dectech

Dectech strives to provide the most accurate and best value forecasts available on how people will behave in new situations. Founded in 2002, we've conducted more than 400 studies involving over three million participants. We hold that people make very different decisions depending on their context and struggle to self-report their beliefs and motives. So,

we developed Behaviourlab, a randomised controlled trial approach that immerses participants in a replica of the real-world decision environment. Over the years we've shown how Behaviourlab can provide higher accuracy forecasts and more actionable insights.

www.dectech.co.uk

Discovery Vitality

Vitality is a platform for behavior change, underpinning the insurance products of Discovery Limited and of leading insurers in 40 countries, impacting approximately 40 million lives. The Vitality model, established by Discovery Limited in South Africa, has been incentivizing behavior-change among its

clients for over 25 years. Vitality creates shared value by combining behavioral economics, clinical science, and financial incentives to encourage and reward members for taking steps to improve their health.

discovery.co.za/business/vitality

FehrAdvice & Partners

FehrAdvice & Partners, based in Zurich and Vienna, is a pioneering behavioral economics consultancy. The firm integrates behavioral science with real-world applications to design evidence-based interventions. Operating under its proprietary Behavioral Economics Approach (BEA™), FehrAdvice delivers actionable solutions that address cognitive biases, social norms, and human motivations. Their unique consulting process combines workshops, field experiments, and data-driven recommendations to ensure measurable

impact. Also, the process emphasizes collaboration, experimentation, and results, delivering prototypes tested in the field for maximum reliability. Known for their expertise in behavior change, FehrAdvice applies its insights across diverse domains such as sustainability, pricing, digital transformation, and culture, making them a trusted partner for organizations seeking innovative and effective strategies.

fehradvice.com

Lucerne University of Applied Sciences and Arts

The Lucerne University of Applied Sciences and Arts is the university of applied sciences of the six Central Swiss cantons and is the largest university level institution in Central Switzerland. More than 8,100 students are working towards a Bachelor's or a Master's degree, while an additional 12,000 participate in continuing and executive education programs.

The university's Institute of Tourism and Mobility ITM conducts research in the areas of tourism, mobility and sustainability. It provides degree programmes for those wishing to embark on a career in the tourism or mobility sector as well as those moving into the field from other backgrounds. Additionally, the ITM

offers consultancy services for practising tourism and mobility professionals.

In order to remain competitive, the tourism and mobility industry must always be customer-centric. This is why the ITM has pooled its expertise with a team of psychologists and behavioural economists to research the perception and decision-making behaviour of consumers in tourism and mobility and to recommend targeted measures based on this research.

www.hslu.ch/en/

Neovantas

Neovantas is a top international management consultancy focused on accelerating change through the power of behavioral data science. We focus on transforming behavior to assure business results in a sustainable way over time. Our consulting team is specialized by sector (retail banking, insurance, telecoms, airlines, utilities, etc.) and functions (advanced analytics, behavioral science, business

transformation and innovation). We build strong, lasting relationships with our clients through our commitment, integrity, professional excellence and creativity. Our international presence has been expanded with projects both in Europe and in Latin America.

www.neovantas.com

ThinkPlace Ethiopia

ThinkPlace is a globally recognized design and innovation consultancy dedicated to creating positive societal impact. Named one of Australia's 20 most innovative companies by the Australian Financial Review and a recipient of prestigious design awards, we specialize in tackling complex challenges in health, education, and sustainable development.

As a member of the UN Global Compact, our work aligns with the Sustainable Development Goals, using a blend of design-led innovation, behavioral science, and co-creation to drive systemic change. Our approach integrates behavioral science principles—such as social norm change, habit formation,

and gamification—to ensure sustainable adoption and long-term impact. By engaging stakeholders at every level, we co-create solutions that are inclusive, equitable, and practical.

With nine studios worldwide—including locations in Kenya, the U.S., Singapore, Australia, New Zealand, and the UK—ThinkPlace Ethiopia consists of ten Ethiopian designers, collaborating with a global network of 140+ experts. This cross-cultural exchange enriches our work, allowing us to develop innovative, contextually relevant solutions that transform lives.

www.thinkplace.com.au

Voya Financial

Voya Financial, Inc., is a leading health, wealth and investment company offering products, solutions and technologies that help individual, workplace and institutional clients become well planned, well invested and well protected. The Voya Behavioral Finance Institute for Innovation is focused on gaining deeper insights into the savings decisions of Americans to

help transform their retirement. The Institute's work is differentiated by the ability to leverage behavioral science, a spectrum of lab and field experiments, and the speed and scale of the digital world to improve financial wellness outcomes.

www.voya.com

