Open Minded

Searching for Truth about the Unconscious Mind

By: Ben R. Newell, David R. Shanks

Citation:

Open Minded: Searching for Truth about the Unconscious Mind
By: Ben R. Newell, David R. Shanks
DOI: 10.7551/mitpress/14922.001.0001
ISBN (electronic): 9780262375375
Publisher: The MIT Press
Published: 2023

The open access edition of this book was made possible by generous funding and support from MIT Press Direct to Open

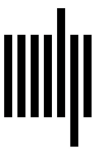

The MIT Press

11 The Mind Reclaimed?

On December 31, 2019, the Wuhan Municipal Health Commission informed the World Health Organization of some cases of "pneumonia of unknown etiology."[1] As is now common knowledge, these were the first recorded cases of COVID-19, a disease that has killed more than 6 million people and infected almost half a billion worldwide.[2] The impact of COVID-19 on society is difficult to overstate, not just for those of us who have been directly affected via illness and death of loved ones, but also via the changes to the way we work, live, socialize, and communicate. The pandemic has brought into sharp focus the urgent need to understand and influence human behavior. In many ways, it presents an opportunity for psychological and behavioral science to shine—to show their worth in helping us adapt to the "new normal" of living with COVID-19. In this final chapter, we use the pandemic to illustrate the importance of having well-developed and falsifiable theories of behavior if we are to use them to guide us, and we argue that research on unconscious thinking has largely failed to provide such theories. We also highlight that in the weak theories that have been put forward, unconscious mental processes seem to operate like dark matter in the universe—the residue that we infer "must" be there, but for which we have little direct evidence and even less theoretical understanding.

Promoting Social Distancing

A very early piece of health advice from governments around the world was to socially distance. In 2019, few of us would have known what that meant, but now the idea that we should maintain space between ourselves and others in order to stop the spread of the virus has become well established. If we want to understand why people may or may not comply with

the request to socially distance then we need to have a theory of the psychological drivers of social distancing behavior. Absent a solid theory, our attempts to influence and maintain behavior changes will be futile.

In chapter 1 we discussed the theory of planned behavior (TPB); it turns out that along with its many successes across a broad spectrum of behaviors, it can be applied to understanding social-distancing compliance.[3] The TPB proposes that the extent to which people will socially distance depends on their *intention*, which in turn comes from three sources: their *attitudes* toward social distancing, their *subjective norms* regarding social distancing, and their *perceived behavioral control* over social distancing. Figure 11.1 shows how these aspects are related to one another. The important aspect of the theory, for our argument, is that behavior is determined by conscious knowledge, beliefs, and attitudes. There is no box or arrow in figure 11.1 for unconscious mental processes.

How does TPB fare when used to predict social distancing? To answer this question, in April 2020, when much of the world was under stay-at-home orders or lockdowns, Laura Gibson from the University of Colorado and colleagues asked a group of US adults about their attitudes, subjective norms, and perceived control toward social distancing. Attitudes were assessed via questionnaire items such as whether they found social distancing "healthy" or "unhealthy," subjective norms inquired about whether they thought most of their friends were engaging in social distancing, and perceived control measured confidence in being able to socially distance. For each item, a numerical rating was given, for instance from 1 (strongly disagree) to 7 (strongly agree). Participants then answered questions on their intentions to social distance over the coming weeks and months, and their actual distancing behavior over the preceding two weeks. In July 2020, when many lockdowns had been lifted, the same participants were approached and asked the same set of questions. This longitudinal design allows us to see how people's intentions, perceived control, and behavior are related across time.

A clear pattern emerged: the baseline measures, taken in the first survey in April, showed that attitudes, subjective norms, and perceived behavioral control predicted intentions to socially distance. Moreover, these intentions also predicted the extent of social distancing at follow-up in July. Finally, the level of perceived control measured in April was related to behavior in July, suggesting that control can bypass intentions to influence behavior. In

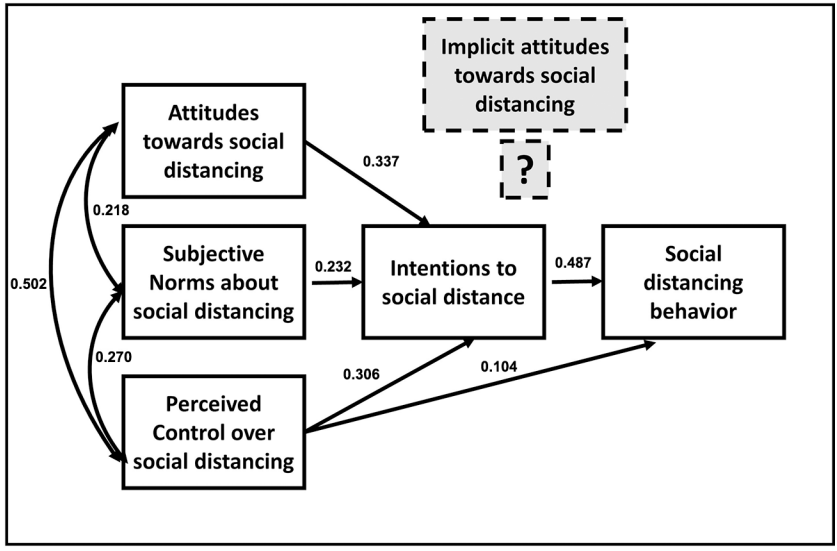

Figure 11.1
Theory of planned behavior applied to social distancing in the COVID-19 pandemic. Hundreds of studies have established that intentions to behave are driven by attitudes toward that behavior, subjective norms about those behaviors (what others do and think), and perceived control over engaging in the behavior. In turn, these intentions explain reported behaviors. The numbers indicate the strength of each link, where 1.0 is a perfect correlation and 0 is no relationship. The gray box marked "implicit attitudes" and a question mark illustrate uncertainty about how an implicit measure would be added to the theory. The data are from Laurel P. Gibson, Renee E. Magnan, Emily B. Kramer, and Angela D. Bryan, "Theory of Planned Behavior Analysis of Social Distancing during the Covid-19 Pandemic: Focusing on the Intention–Behavior Gap," *Annals of Behavioral Medicine* 55, no. 8 (2021): 805–12. https://doi.org/10.1093/abm/kaab041.

other words, the conscious, reportable beliefs and attitudes elicited by the questionnaires provide a good account of reported behavior. The numbers on the links in figure 11.1 indicate the strength of each link.

The appeal of such a clear set of results is that they provide direct guidance to policymakers wishing to develop health communications or interventions. Public health campaigns aimed at increasing social-distancing compliance should target attitudes, subjective norms, or perceptions of control (or some combination of these). Exactly how such targeting should be done is, of course, a question for another research program. One obvious

extension to the Gibson study would be to use experimental designs (or randomized controlled trials) in which different groups of participants are presented with messages that are aimed to influence the different factors thought to underlie behavior. For example, one group might receive messages that emphasize subjective norms ("80 percent of people in your local area are socially distancing"), while another group see messages relating to perceived control ("going shopping early or late in the day helps avoid crowds"). Such designs would take us beyond the correlational findings and allow for causal inferences about the strength of different interventions.

One might ask whether the addition of unconscious factors into the TPB framework would improve its ability to predict and explain behavior. For example, one could imagine adding a box for "implicit attitudes" to figure 11.1 that is connected to intentions. The challenge then would be to demonstrate that the inclusion of such factors produced a better fit to the observed data or a better explanation of behavior than a model that omitted them. In essence, we'd be asking what the value-added, in terms of understanding behavior, is of assuming that some determinants of social distancing are unconscious. For example, maybe people's explicit attitudes to social distancing are less predictive than their implicit ones. People might say that social distancing is not unpleasant when asked directly, but if we could tap into their implicit attitudes, we might find a very different opinion. To do this, we need to measure those implicit attitudes, and that is not as straightforward as it might seem.

The Validation Crisis

The Implicit Association Test (IAT) discussed at length in chapter 5 provides a way to measure implicit attitudes. It is useful to return to this discussion because it provides an example of how one might test the idea that implicit attitudes are an additional predictor of behavior beyond explicit ones. Recall that the IAT attempts to measure automatic evaluations by combining two types of responses concurrently (see figure 5.1). For the purposes of illustration, let's imagine we want to measure unconscious ageism. In the experiment, you would see a sequence of randomly interspersed words and faces and have to press the left key for negative words and older faces and the right key for positive words and younger faces. If you have an unconscious bias against the elderly, it will be relatively easy to respond rapidly

in these circumstances as the separate decisions are compatible in the sense that the left key is used for both negative words and the relatively disliked (older) faces and the right key for both positive words and the relatively liked (younger) faces. Your button presses should be faster on average under these circumstances than in another part of the experiment in which you have to press the left key for negative words and younger faces and the right key for positive words and older faces, where the separate decisions are incompatible. The difference between speed of responding in the compatible and incompatible stages is the IAT's estimate of your unconscious ageism.

Just as in the COVID-19 social-distancing example, we can build a theory in which unconscious ageism as measured by the IAT is assumed to cause some important aspect of decision making such as favoring younger compared to older job applicants. If we have IAT scores and job applicant ratings for a large sample of individuals, we can ask whether these measures are correlated. We can make our model more complex by adding a second potential cause of discriminatory behavior, namely, conscious or explicit ageist attitudes, which we could measure using a standard questionnaire (for example, the Expectations Regarding Aging scale).[4] Now we can ask whether unconscious, implicit attitudes (measured by the IAT) or conscious, explicit attitudes (measured by the questionnaire) correlate better with our measure of behavior.

This sounds relatively straightforward, but in reality, a great deal of further work is needed to validate our implicit attitude measure. Establishing the validity of a construct and the ability to measure it accurately is not trivial. Indeed, even though the notion of construct validity has a venerable history in psychology, it is often given insufficient attention in the development of our theories. This reticence to engage properly with the definition, refinement, and measurement of constructs has contributed substantially to the parlous state that the behavioral sciences now find themselves in.

To illustrate how concepts develop, consider an example from outside psychology. The concept "electron" was introduced to physics in the 1890s. Initially the term meant an elementary unit of charge. But over time, with experimentation, theoretical advances, and step changes in our understanding brought about by quantum theory, the meaning has changed radically. An electron now refers to an elementary particle that is a fundamental constituent of matter, with a negative charge of 1.602×10^{-19} coulombs, a mass of 9.108×10^{-31} kilograms, a spin of ½, and so on. Achieving such a level of

precision in the definition of psychological concepts and linking them in precise ways to other theoretically related concepts might seem unattainable, but that is no reason to abandon our attempts to do so.[5]

One of the reasons construct validation is so challenging in psychology is that just as with electrons, we cannot directly observe mental processes. As we discussed in chapter 7, much of psychology and cognitive science is dedicated to inferring the impact of these latent mental states on behavior. We cannot observe perceived control, but we can draw inferences about the level of control someone has over a particular behavior. This can be done either by asking them directly or by measuring the impact on behavior of a manipulation that we think will affect their control. This is the approach that has been taken in hundreds of studies testing TPB and has enabled us to build confidence in the model depicted in figure 11.1. But these attempts at measurement depend fundamentally on our assumptions that our tools measure what they are intended to measure. This is the essence of establishing construct validity.

Returning to our ageism IAT example, let's suppose that our findings suggest that implicit but not explicit attitudes are strongly associated with behavior. Before concluding that our behavior toward other people is strongly influenced by unconscious ageist attitudes, we need reassurance that we measured conscious attitudes in a sound manner. Perhaps our questionnaire items are simply not very appropriate and fail to sensitively discriminate between people with and without ageist attitudes. Or perhaps our questionnaire is insensitive because the respondents were wary about revealing their true views. There are any number of ways in which a test may be a poor assessment of the construct it is designed to measure. Note that validity is not quite the same as exhaustiveness. In chapter 2 (see table 2.1) we discussed some of the criteria that adequate tests of awareness need to meet before we can deem them sufficiently exhaustive and sensitive to be usable; for instance, awareness tests should be administered at or very close to the time at which behavior is being measured. Validity includes exhaustiveness but encompasses other attributes too, such as convergence with other tests designed to measure the same construct (awareness in this case).

Against the standard criteria for validation, it is debatable whether the IAT provides a good measure of unconscious attitudes.[6] Primarily this is because its incremental predictive validity—the extent to which it predicts

behavior over and above consciously reportable attitudes—is tiny.[7] But this is not the key point. At least with respect to the IAT, considerable efforts have been made to validate it as a measure of unconscious attitudes, and its properties have been fairly thoroughly assessed. When we turn to other domains in which unconscious processes have been investigated, we find that the picture regarding the IAT is very much the exception. In research on subliminal perception, priming, decision making, and many other areas, virtually no efforts have been made to establish that the constructs being measured have any validity. There is, in short, very little reason to believe that something like money priming (that we reviewed extensively in chapters 9 and 10) can be measured validly and distinguished from conscious attitudes. Despite hundreds of publications on money priming, the relevant research needed to establish its soundness as a measurable and distinct psychological construct has simply not been undertaken.

The consequences of inadequate construct validation are hard to exaggerate. Different measurements, superficially probing the same construct, can yield completely contrasting results. In a major investigation led by Justin Landy, teams of investigators were asked to devise their own methods for answering several different hypotheses, one of which related to unconscious thoughts:[8]

Hypothesis: People explicitly self-report an awareness of harboring negative automatic associations with members of negatively stereotyped social groups.

This hypothesis conjectures that people tend to have some insight into their own biases and prejudices against certain social groups such as the elderly or racial minorities. On the face of it, the hypothesis looks as if it should be fairly easy to test. A sample of randomly chosen individuals could be given a question such as

Q1: Although I don't necessarily agree with them, I sometimes have prejudiced feelings (like gut reactions or spontaneous thoughts) that I don't feel I can prevent,

and asked to what extent they agree or disagree using a scale from 1 ("strongly disagree") to 7 ("strongly agree"), and with 4 ("neither agree nor disagree") being the crucial midpoint. We could then ask whether the average rating on this scale is greater than 4, indicating that indeed most people explicitly

self-report an awareness of harboring negative automatic associations with members of negatively stereotyped social groups, or instead is close to or even below 4, in which case the hypothesis is falsified.

Suppose we obtain the former outcome (which was indeed the case): Does that mean the hypothesis is correct? Bear in mind that although our question (Q1) seems intuitively to probe beliefs about the construct we're interested in, we could equally measure those same beliefs in a virtually infinite number of other ways. Without undertaking a validation exercise, we have no way of knowing whether these different ways would all yield similar results and we don't actually know that Q1 probes the target construct at all. If we find an average rating of, say, 5 in response to our question, what does that tell us?

The results of Landy's project were stark: different teams, constructing different questions to test the same hypothesis, reached completely divergent conclusions. For example, one team employed this alternative question wording:

Q2: Regardless of my explicit (i.e., conscious) beliefs about social equality, I believe I possess automatic (i.e., unconscious) negative associations towards members of stigmatized social groups.

With this wording, the majority of respondents did *not* report harboring negative automatic associations. Hence two questions, both of which seem superficially to get at the same psychological state, yield quite contrasting results. Indeed across the many teams that participated in Landy's project, this was the typical pattern found for all of the research hypotheses investigated. Each team constructed (on the face of it) completely reasonable questions to test the same hypothesis, but the conclusions reached by the various teams did not converge. The message is plain: it is not sufficient to appeal to plausibility when constructing test items; they need to be subjected to a thorough validation exercise. If we want to determine whether Q1 or Q2 or indeed any other question provides a valid means of measuring beliefs about negative automatic associations, we need to establish convergent validity (the question should yield scores that are similar to other questions designed to measure the same beliefs), discriminant validity (the question should yield scores that are unrelated to the scores for questions designed to measure different beliefs), and predictive validity (it should predict some meaningful aspect of behavior that we expect to be related to

beliefs about negative automatic associations, such as willingness to attend an unconscious bias training course). The question should also be reliable, in the statistical sense of yielding roughly similar scores when any given individual answers the question on different occasions.

This neglect of validation is widespread across psychological science to such an extent that it has been described as a "validation crisis."[9] Instead of carefully measuring and characterizing latent psychological states, beliefs, or dispositions, researchers focus unduly on very simplistic questions such as, "Does an increase in A cause an increase in B?," answering such questions solely by reference to statistical significance and the $p < .05$ threshold.

Theory Building and Severe Tests

There is an important distinction to be made here though. Some phenomena (for example some forms of priming, facial feedback effects) have proven to be very hard to replicate, leading the research community to be skeptical about their existence. This contrasts with phenomena that are highly robust, but for which explanations for their properties are still hampered by poor construct validation. Take the example of the bat-and-ball problem presented at the start of chapter 7 (if a bat and a ball cost $1.10 in total and the bat costs $1 more than the ball, how much does the ball cost?). It is undisputed that the majority of people get this problem wrong when they first see it, and this behavioral regularity requires an explanation. The standard flavor of this explanation follows the dual-system framework depicted in figure 7.1. People fail to engage their deliberative system 2 processes, and so the automatic system 1 provides a plausible but incorrect answer. Performance on the bat-and-ball problem along with two similar questions is then purported to measure the degree to which people can inhibit system 1 when they should be listening to system 2, a construct described as *cognitive reflection*.

But is cognitive reflection an ability or a disposition to think in particular ways? And if it is an ability, how does it relate to other general cognitive abilities such as intelligence or working memory? Some progress has been made toward answering these questions with the general consensus that the Cognitive Reflection Test (CRT) measures something more than simple mathematical or reasoning ability.[10] The test also appears to have incremental predictive validity over similar measures in predicting other types of behavior such as people's willingness to gamble. But this is not the main

problem. Even if one accepts that cognitive reflection is an independent psychological construct that is measured validly by the CRT, the broader question of how it relates to all of the other constructs and capabilities listed in figure 7.1 remains unanswered.

The dual-system framework outlined in figure 7.1 is a far cry from the detailed *nomological net*—a term used to describe how the different constructs in a theory are related to one another—shown in the TPB COVID-19 example of figure 11.1. The construction of these nets is the hard but fundamentally important part of theory development. Without clear statements of how and why our different, validated, psychological constructs interrelate, we remain in a quagmire of loose relationships with untestable and unfalsifiable theories. As we noted in chapter 7, many of the dual-system models include collections of different constructs—capacity and automaticity—that are assigned to different systems with little attempt to establish a coherent understanding of how these constructs link to one another or indeed whether the measurements of them are valid. To reiterate our conclusion from chapter 7, the result of this incoherence is that the attempts to develop dual-system theories become akin to the epicycles devised by Greek astronomers to rescue the theory that the Earth revolves around the sun; their contact with hard facts becomes more and more distant while precise predictions that could be subjected to experimental tests are scarce.

At the start of this book, we suggested that the science of psychology needs to be rebuilt from the ground up on firmer foundations. Those firm foundations are our theories of behavior. So why is theory building so difficult? To give a sense of the challenges, imagine that we propose a theory that could explain some of the more surprising findings that we've reviewed in earlier chapters. What kind of theory might explain why we'd walk more slowly down a corridor after solving anagrams related to the concept of old age, or rate cartoons as funnier when we hold a pen between our teeth, or rate someone as warmer on a personality scale after holding a hot cup of coffee rather than a cold can of Coke? We now know that many of these findings do not replicate, and thus the need to find explanations becomes moot, but for a moment, let's consider a world in which these or similar findings were real. How might we explain them?

We might propose a general theory called "embodiment priming theory."[11] The core assumptions of this theory are as follows. First, abstract concepts like warmness toward an individual are grounded in bodily states,

sensations, or movements, like the physical sensation of warmth. Second, inducing that bodily state or sensation, for instance, by giving someone a warm cup, activates or primes that concept. And finally, the concept influences downstream behavior related to that concept, such as rating the warmth of a stranger. The tenets of the theory are similar to the ideas we explored in chapter 3's discussions of the ripples of activation. How would we test embodiment priming theory? The first challenge is defining the scope of the theory or the space of potential hypotheses that we might want to test. There are many abstract concepts that can be grounded in innumerable bodily states and sensations that can be induced experimentally in a host of different ways. There is also a variety of downstream behaviors that could be influenced by the primed concept. Hypotheses are limited only by the ingenuity of the experimenter. For example, we might test the hypothesis that a person in the UK who stands on their head for three minutes will feel more connected to people in Australia by priming the concept of being Down Under.

The problem here is that the combinatorial explosion that results from pairing concepts, groundings, their manipulations and measurements presents an asymmetry in tests of the theory. Let's imagine we ran the Down Under experiment and found no effect. Does that falsify the theory? Not necessarily; we could argue that we chose the wrong amount of time for standing on your head—perhaps 1 minute is better because after 3 minutes, people become woozy and don't feel connected to anyone. Or perhaps it only works for Brits who have relatives in Australia. There are many ways to argue ourselves out of a failure to find an effect but to remain convinced that the embodiment priming theory is a space of hypotheses worth exploring. Of course, if we are lucky enough to find the effect (perhaps with the aid of *p*-hacking), then our confidence in the general applicability of the theory increases. But should it?

The answer lies in how tightly our hypotheses are linked to our theories. If a theory strongly implies a given hypothesis, then attempts to test those hypotheses are useful in the sense that they are diagnostic about the boundaries of the theory. This brings us back to the crucial importance of construct validity. If the elements in our theory are loosely defined or poorly validated, then any tests of the theory will be hamstrung from the start. If we do not have a well-validated concept of Down-Underness, then finding an effect of standing on our heads will be meaningless. A useful

distinction here is between discovery-oriented and theory-testing research (we met similar ideas before in chapter 9 when discussing exploratory and confirmatory research). Discovery-oriented research is what we are doing when we are searching, perhaps stumbling, through the space of hypotheses that our theory implies. It describes situations where we focus on testing a single hypothesis with a single manipulation and care less about how that collection of effects (and their absences) integrates into a broader understanding. Theory-testing research, in contrast, is where we focus on strengthening the inferential links between well-established empirical findings and formalized theories that can explain them.

As we have argued throughout the preceding chapters, psychology has been stuck for a long time in cycles of discovery-oriented research, and much of this has been fueled by increasingly implausible claims about the power of unconscious mental processes. Most frustrating, the unconscious is treated in many of these explanations like dark matter in the universe—the residue that we infer must be there, but for which we have no theoretical understanding. To make real progress, we must abandon these black-box explanations and construct genuinely testable explanations of human behavior.

A strong theory-led approach, combined with the improvements in research practices that we've discussed in this book, is the way forward to a true science of human behavior. A piecemeal approach in which we address one aspect but not the other will be insufficient. For example, a pure focus on improving methods will not be enough to strengthen theories. The recent enthusiasm for preregistration provides a good illustration. The act of preregistering the hypothesis that standing on your head in the UK makes you feel more connected to people in Australia does not in and of itself make it a stronger or better hypothesis. Moreover, if a researcher happened to find an effect that supported the hypothesis, it should not necessarily increase its credibility. What makes a hypothesis credible is how closely it is connected to or constrained by the theory that motivates it. Without a detailed nomological net outlining how our (validated) construct of Down-Underness relates to the physiological state induced by being inverted, which in turn influences our feelings about distant others, a statistically significant effect is no more informative than a lucky guess. And preregistering that lucky guess won't help.

This is not to say that preregistration is without merit. As discussed in chapters 9 and 10, it can help to negate the impact of HARKing (hypothesizing

after the results are known) and forces researchers to be explicit about how, why, and how strongly their predictions are justified by their theory before seeing any data. But absent the strong theory to begin with, preregistration appears to be best thought of as a cure for the symptoms of our current malaise than a solution for the core problems.[12]

One way that our predictions could be made more precise and subject to severe and decisive tests is to use the tools of computational modeling. In addition to wanting to understand the magnitude of relationships between well-measured constructs, we also often want to know what exactly the mental machinery is and what the algorithms are that implement these relationships. Computational modeling, in the form of computer programs designed to capture the essence of these algorithms, plays an important role in addressing such questions.[13] We might speculate, for instance, that the process of choosing between two options, A and B, involves an evidence-accumulation process in which, moment by moment, internal counters accumulate reasons for choosing each option, and a decision is made as soon as one of these counters reaches a prespecified threshold. Such a mechanism could be turned into a computer model and used to generate testable predictions about decision making. Computational models have made an enormous contribution to modern psychological science, yielding deep mechanistic explanations for decision making, perception, memory, and other cognitive processes. Usually the aim is to compare two or more models to a rich set of data to see which fits the data more closely rather than to focus on a single question about whether an effect is or is not statistically significant. Arguably, the best way to test verbal theories is in quantitative terms—that is, through models.

As with construct validation, however, the computational approach has been underemployed in the study of unconscious mental processes. The general approach would be to construct a computational model of whatever aspect of behavior is under investigation (response times or choices, for example) and then compare a model that includes an unconscious aspect with one that does not to gauge whether the former appreciably improves explanatory power. In the rare cases where this approach has been taken, little evidence for such an outcome has been found, but again that is not the main point.[14] Rather, the key issue is the disappointingly limited use of the approach. When a well-validated tool is available to help us answer crucial questions about human behavior, it is depressing that the tool is not used.

This does not mean that computational approaches are a panacea. A construct does not necessarily become more valid simply because it is instantiated in a formal mathematical model. Neither is a theory that lacks formalization necessarily bad. A case in point is Darwin's theory of evolution. The theory has very clearly defined concepts or constructs, but the relationships between them remain qualitative rather than quantitative.[15] The central point remains, though, that for our understanding of human behavior to improve, we must not shy away from the difficult task of theory building.

Opening Our Science and Minds

This book is about the need to be open about our science, but more fundamentally, it is about opening our minds to what on the face of it seems like an uncontroversial idea: we are the authors of our own actions, we typically have good intuitions and understanding about the reasons for why we behave as we do, and thus we also have the ability to change our behavior. The fact that this idea brooks any controversy is testament to the reach of the powerful unconscious mind meme. Realizing that the notion of unconscious thinking is built on very shaky foundations has significant implications for how we lead our lives. In one sense, it should be deeply empowering: we are not constantly being buffeted by forces beyond our control, nor do we have deeply rooted unconscious biases and prejudices lurking at the bottom of the iceberg. But in another sense, this realization is unsettling: we now have to take responsibility for our failures rather than blaming them on our "brain" acting unconsciously. We also have to acknowledge that our prejudices are often conscious—in the sense of being accessible to our awareness—even if we seek to deny their existence or lay the blame elsewhere.

Just over one hundred years ago, the world was in the grip of another pandemic, the influenza pandemic that killed at least 50 million people and infected almost a third of the world's population.[16] George Soper, a major in the Sanitary Corps of the US Army, wrote an article at the time for the journal *Science*, "The Lessons of the Pandemic."[17] In a telling line, he noted that "the measures which were introduced for the control of the pandemic were based upon the slenderest of theories."[18] How far have we come in the intervening century? How much fatter have our theories become?

Soper argued that successful prevention required overcoming the public's indifference or underappreciation of the risks of spreading the disease. He

suggested that much of the transmission occurred "unconsciously, invisibly, unsuspectingly." Soper's recommendations for the most essential behaviors to stop the spread are eerily reminiscent of the much-repeated slogans of the COVID-19 era: avoid needless crowding, smother coughs and sneezes, wash hands before eating, and open windows when practicable. Has a century's worth of psychological research brought us any closer to understanding how to achieve those changes in behavior? At the onset of the current pandemic, many researchers were swift to point out a raft of relevant findings from psychology that could help to mitigate the impact of COVID-19.[19] These ranged from research on basic science communication to work on leadership, stress, threat perception, and social context. There is no doubt that many of these studies could be applied to understanding and changing behavior. But taken as a whole, how much do they contribute to genuinely deep theory development? Do they constitute real progress toward achieving a true science of the mind, or are they an ad hoc collection of effects found in underpowered studies, using inappropriate methods and questionable statistical approaches? The answer must lie between these extremes. It would be churlish to argue that we've made no progress (take TPB for example) but naive to claim that our theories have allowed us to advance much beyond George Soper's recommendations. If we don't want to find ourselves in the same position in another one hundred years, then we must address the challenges outlined in this book. Abandoning the myth of the smart unconscious is a good place to start.

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Newell, Benjamin R., 1972– author. | Shanks, David R.
Title: Open minded : searching for truth about the unconscious mind / Ben R. Newell and David R. Shanks.
Description: Cambridge, Massachusetts : The MIT Press, [2023] | Includes bibliographical references and index.
Identifiers: LCCN 2022038725 (print) | LCCN 2022038726 (ebook) | ISBN 9780262546195 (paperback) | ISBN 9780262375368 (epub) | ISBN 9780262375375 (pdf)
Subjects: LCSH: Subconsciousness. | Thought and thinking. | Self-consciousness (Awareness)
Classification: LCC BF315 .N479 2023 (print) | LCC BF315 (ebook) | DDC 154.2—dc23/eng/20230316
LC record available at https://lccn.loc.gov/2022038725
LC ebook record available at https://lccn.loc.gov/2022038726