CHAPTER 18
IS RITUAL BEHAVIOR A RESPONSE TO ANXIETY?

Martin Lang, Jan Krátký, John Shaver, Danijela Jerotijević, and Dimitris Xygalatas

Introduction

Imagine the following scenario: You are a student sitting in front of a committee during finals, waiting for the professors to come up with a question that will take the next hour or so to answer. There are several factors playing in your favor: you know the professors, their interests, and their favorite questions; the list of usual exam topics was circulated well in advance; and you have prepared for weeks. Everything should go smoothly. You hope. Now imagine a slightly different scenario. At the Department for the Scientific Study of Religions in Brno, Czech Republic (where our group works), the faculty decided to introduce a new, supposedly unbiased mechanism governing the topic-selection process, namely a thirty-two-sided die. The rationale for the use of the die is simple: because topics are chosen at random, no one can complain about fickle and eccentric professors giving tricky questions. But the die substantially changes the game for the students. They cannot reasonably anticipate the question based on professors' interests or previous exams; the only thing they know for sure is that they will face an uncontrollable, random process that will select the questions for them.

So, although you prepared really well for the exam, you couldn't read all the materials. And surely there is a question about Heideggerian metaphysics waiting for you. No one in their right mind knows about such things! If only there was a way to influence the die and make it choose anything but Heidegger. Note that we have nothing against Heideggerian philosophy; it just seems complicated to students at times. Perhaps doing something will help. Anything! You can't just do nothing! And so, before rolling the die, you reach into your pocket to touch your lucky rabbit's foot, or perhaps you quickly recite a short prayer—"Please, God, no Heidegger." Perhaps this may help a bit, you think. It's certainly better than doing nothing, right? Indeed, many of us would behave in a similar way. Few of us would be so blindly optimistic as to just take the die and roll it without doing *something*.

Stressful situations are common in our lives: public performances at schools or in sports are frequent, serious family situations befall us all too commonly, and natural disasters, unfortunately, occur. People cannot have full control over every possible situation. Think about the last time you experienced stress. Your heart probably raced and your body tensed. This physiological state is a useful mechanism, since it motivates

people to react immediately and, in most everyday situations, to adopt rational and instrumental actions to avert undesirable but likely outcomes (e.g., avoiding entering dark forests in favor of well-lit streets). However, as a response to events that cannot be prevented, stress and anxiety can be detrimental for our coping abilities and, in the long term, for our health. Attempting to influence the odds by doing any kind of action can prevent us from feeling utterly powerless and help us manage our anxiety, despite the fact that such actions usually have unknown or only dubious causal connections to the desired outcome (such as rubbing a rabbit's foot prior to an exam). Often recognized as ritual behaviors due to their opaque causal relationship with desired goals, such actions are of interest to anthropologists, psychologists, and cognitive science of religion.

Perhaps the most widely recognized theory of ritual and anxiety comes from Bronislaw Malinowski (1948/1992), one of the first anthropologists conducting fieldwork. Malinowski spent two years living in the Trobriand Islands with the locals, studying their language, their customs, and their culture in general. While observing Trobrianders' rituals, Malinowski uncovered a startling pattern: the rituals usually occurred before events that were beyond an individual's control. For example, Trobrianders would perform their magic before going to fish on the open sea, which was risky, uncontrollable, and uncertain. In contrast, they would not perform rituals before fishing in the local lagoon, which was easy to navigate and guaranteed a safe catch—so much so that even kids engaged in lagoon fishing. On the basis of these and other similar observations, Malinowski concluded,

> [A Trobriander] knows that a plant cannot grow by magic alone, or a canoe sail or float without being properly constructed and managed, or a fight be won without skill and daring. He never relies on magic alone, while, on the contrary, he sometimes dispenses with it completely, as in fire-making and in a number of crafts and pursuits. But he clings to it, whenever he has to recognize the impotence of his knowledge and of his rational technique. (1948/1992: 32)

Malinowski is here describing a situation similar to our test-taking: everyone knows that having a rabbit's foot, while completely ignoring exam materials, will not yield an A+. But when you've done as much as you could and there is still some lingering uncertainty surrounding the exam, you might as well try to do something extra (*like a ritual*) to bring good fortune. From these and similar examples, we can derive a general prediction: when people find themselves in uncertain situations that do not allow direct control of future events through clearly linked actions, they will engage in ritualized behavior in an effort to influence the outcomes.

We have just made a step from ethnographic observations[1] to a general scientific prediction. But how can we go about testing the prediction? One way would be to carry out more systematic fieldwork, such as in the study conducted by Giora Keinan (1994) in Israel during the Gulf War. Keinan surveyed 174 Israeli citizens living in two cities: Tel Aviv and Jerusalem. An important difference between these two cities is their proximity to Iraqi borders: while Jerusalem was safe from Iraqi missiles, participants in Tel Aviv

lived in constant fear of unpredictable and uncontrollable missile attacks. As a result of these threats, Keinan predicted that people from Tel Aviv would perform more magical rituals. Through a series of questionnaires on magical thinking and behavior (e.g., "It is best to step into the sealed room right foot first"), Keinan showed that the uncertain and uncontrollable conditions of living in Tel Aviv were associated with a greater frequency of ritual behavior. Other researchers have investigated the relationship between anxiety and ritualized behavior in various contexts, such as among athletes, gamblers, or students during test-taking, and they all arrived at the same conclusion: anxiety-inducing situations tend to be associated with an increased frequency of ritual behavior (Felson & Gmelch 1979; Schippers & Van Lange 2006; Sosis & Handwerker 2011).

While these studies make very valuable contributions, they are limited in their explanatory power because they investigate correlations (statistical association) among variables but do not help to understand causal patterns. In other words, the previous studies (e.g., Keinan's) cannot tell us whether it is really the anxiety-inducing situation in Tel Aviv that leads people to perform more rituals, or whether people performing rituals tend to live in Tel Aviv for some other reasons. In order to disentangle the causal relationships, it is necessary to introduce an experimental manipulation, which will allow us to examine the phenomenon in question in the assumed logical sequence. Furthermore, while introducing an experimental manipulation, we need to make sure that no other variable is influencing (or confounding) the predicted relationship; that is, researchers need to control for other external variables that could be driving the effect. Laboratory experiments address these problems through a strenuous control of the lab environment, well-designed procedures, and by manipulating one thing at a time. In the next section, we turn to experimental methods and explain how we used a laboratory experiment to investigate the relationship between ritual behavior and anxiety (Lang et al. 2015).

Methodology

Laboratory experiments are an extremely useful tool for scientists, yet they also present a significant challenge: in order to control for the numerous variables that naturally occur in concert with the variables under investigation, researchers must reduce the phenomenon of interest into a simplified model of the real world and perform their tests on this model. However, this step comes at a steep cost—our models are mere simulations (approximations) of reality and not reality itself. Thus, researchers have to make sure that the way they operationalize their questions (i.e., how they create their models of reality) sufficiently captures all of the important aspects of the phenomenon of interest. For example, if we operationalized anxiety as the amount of sweat participants produce (recall sweaty hands when stressed), we would need to make sure that this measure sufficiently captures the levels of anxiety that people experience in uncertain situations (aside from controlling for more basic confounds like varying temperature in the lab).

This problem grows in proportion to the complexity of the studied phenomenon, especially when such a phenomenon is deeply embedded in its cultural context. While increased sweating is a reliable indicator of physiological response to anxiety (due to the inner workings of the sympathetic nervous system), operationalizing and measuring ritual behavior is much more challenging. First of all, what is ritual? How can we distill ritual's essential aspects from its cultural context without missing its crucial components? And if ritual is culturally specific, what if some participants do not belong to that specific culture/religion? Will the desired effect occur? These are difficult questions with no clear answers.

In our research, we turned to a methodological paradigm from clinical psychology, where researchers identify the main aspects of a particular behavior by examining its exaggerated version as it manifests in some psychiatric disorder. Specifically, we turned to obsessive-compulsive disorder (OCD), in which patients exhibit pathological ritualization. For example, OCD patients repeatedly check whether they have locked the doors or turned off a stove, or they may perform elaborate routines before starting to prepare a meal. People suffering from OCD need to follow such routines down to the last painstaking detail in order to assuage anxiety. Interestingly, the behaviors that are distinctive in OCD resemble some of the rituals described by anthropologists, exposing the important aspects of the ritual form, which is characterized by an invariable sequence of actions (Dulaney & Fiske 1994). Synthesizing the literature on OCD with anthropological insights, Pascal Boyer and Pierre Liénard identified core aspects that are shared by both cultural and OCD rituals, namely, their rigidity, repetitiveness, and redundancy (Boyer & Liénard 2006; Liénard & Boyer 2006). For example, ritual participants put emphasis on the correct, prescribed way of performing ritual actions (the ritual script cannot be changed, much like in OCD food preparation); the sub-actions of rituals are often repeated several times (as in OCD door lock–checking); and, finally, it is not clear how the particular sub-actions involved in rituals are connected to the desired outcome (why is locking the door ten times more effective than just once?). Consider, for instance, the Muslim prayer (*salat*). Ritual participants follow an unchangeable and purportedly ancient form of praying: first they join hands while standing, then they bow, kneel, and touch the floor with their foreheads, and finally kneel with their backs upright. This sequence is repeated many times during each prayer session, yet it is not clear why this specific form (and not any other) leads to the successful worship of Allah.

Apart from human psychopathology, we also sought insights from animal models of anxiety. By exposing various animals, such as mice or voles, to an image of a predator, ethologists (scholars studying animal behavior) have elicited behaviors that include a high degree of ritualization (Eilam, Izhar, & Mort 2011; Eilam et al. 2006). Assuming that an analogy in responses to anxiety between humans and nonhuman animals is appropriate, these ethological studies motivated us to focus on observable behavior rather than self-reports, which can be subject to all sorts of biases (cf. Lang et al. 2017; Xygalatas & Lang 2017). In other words, we were interested in the behavioral responses to anxiety that may be shared across various cultures and therefore likely have deep evolutionary roots. Of course, such an approach is not appropriate when cultural variation is the subject of a

study! One can be interested in why two different religious traditions have various ritual forms and the historical pathways leading to those forms. In our research, we do not claim that ritual content is unimportant; quite the contrary, it can sometimes entirely change the nature of the ritual (e.g., by introducing anxiety itself!). However, we decided to put cultural variation aside for the moment and to focus on the behavioral forms that we predicted to be the main reaction to anxiety. Thus, our basic prediction was that when people experience anxiety, they will display behaviors characterized by rigidity, repetitiveness, and redundancy.

However, this prediction still leaves many questions open: How can we manipulate anxiety, and how can we make sure that this manipulation works, that is, that participants are really anxious? And how can we measure ritual behavior? As you can imagine, these questions could be answered in various ways. This stage of reducing real-world phenomena to laboratory models is often the most difficult step in research. Researchers need to make sure that their materials will simulate the real world as accurately as possible, while still controlling for any variation that is not of interest.

At this point, a survey of previously used methodologies (often in different disciplines) can be most helpful. Experimenters can get an idea of what worked before and how it could be usefully adapted in their own research. We chose to employ the public speaking paradigm, a common psychological method used to induce anxiety in research participants. We divided our subjects into two groups and asked one group to prepare a five-minute speech about an object of art; they would then have to present the speech in front of an expert committee of art critics. (Recall our example of exams and how stressful they can be.) In contrast, the other group's task was just to think about the same object for three minutes.[2]

It is also important to make sure that one's experimental manipulation worked as expected. As mentioned above, the sympathetic nervous system carries anxiety states, and one of the manifestations is increased heart rate, which we typically experience when we feel stressed. Therefore, to assess whether our manipulation was effective, each of our participants wore a heart-rate monitor during the experiment. This allowed us to see whether there was an increase in their heart rates during our anxiety manipulation (preparing the speech) compared to other times. See Figure 18.1(a) for changes in heart rates during the experiment.

Measuring ritualization was more challenging than measuring anxiety. How can we quantify such things as rigidity, repetitiveness, and redundancy and compare ritualization between our two groups? We could have let people wait before giving the speech, observe their behavior, and simply describe with words what they were doing (e.g., "walked back and forth"). However, such an approach would be very imprecise and would introduce a lot of between-participant variation that would be hard to account for in a statistical analysis. Instead, we decided to standardize the performed action by giving participants a specific task: cleaning the artistic object they were supposed to talk about (see Figure 18.1b). But there was one caveat—the object was already clean. Thus, cleaning could serve as a placeholder or platform for expressing participants' ritualization needs. We could have chosen other tasks or objects, but we opted for

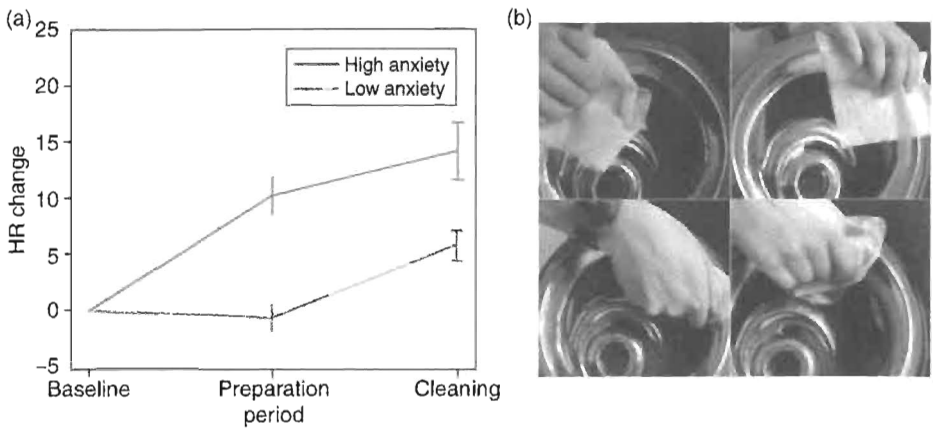

Figure 18.1 **(a)** Mean change in participants' heart rates during various stages of the experiment. Baseline heart rate (preexperiment) was subtracted from other periods to reflect change in anxiety levels. The mean heart-rate change (with standard errors) shows that our manipulation elicited more anxiety in the high-anxiety condition during the preparation and cleaning periods. **(b)** An illustration of cleaning patterns. A participant wears ActiGraph accelerometers on his wrists, which capture the hand-movement acceleration that was used to quantify ritualization during cleaning.

cleaning because it is often part of both cultural and OCD rituals (Dulaney & Fiske 1994; Zor et al. 2009). The specific object that we used was circular with small creases, which allowed participants to engage in both variable and ritualized movements while cleaning. (Had we used a square object, all movements would be restricted by its sharp angles and would artificially appear as ritualized.) See Figure 18.1B for an illustration of cleaning patterns on the work of art.

To address the problem of imprecision in observation methodology, we employed accelerometers positioned on participants' wrists to measure various hand-movement characteristics. The isolation and measurement of elementary behavioral expressions was motivated by action-parsing theory, which attempts to describe the ways humans understand actions (Zacks & Swallow 2007; Zacks, Tversky, & Iyer 2001). Human minds divide ongoing actions into smaller units, and this division facilitates an understanding of actions and allows for better prediction of what will happen next. That is, parsing provides constant feedback on whether the smaller action units follow an expected trajectory toward the final goal. For example, *complex scripts* like dressing up are parsed into *behavioral episodes* like putting on pants, which can be further parsed to basic *gestural expressions* such as reaching for the pants with the right hand, lifting the left leg to push it through the pants, and so on. If one of the gestural expressions failed (e.g., pants were not grabbed), the whole complex script would be violated and in need of updating (grab again, more precisely). Such fine-grained parsing is necessary when evaluating the efficacy of actions and their trajectory to the overall goal. Interestingly, when humans lack an understanding of the mechanistic workings behind desired

outcomes (as in ritual behavior), they focus on every detail of gestural expression (e.g., a specific finger configuration when throwing rice into the fire). Thus, in the current study, we measured behavioral ritualization at the lowest level of gestural expressions, where the ritual form manifests itself most apparently. Focusing on the characteristics of fine movements allowed us to quantify hand-movement acceleration and examine whether the acceleration patterns exhibited the ritual form.

We used recurrence quantification analysis (RQA), a statistical technique for analyzing the behavior of nonlinear dynamical systems (Marwan et al. 2007; Webber & Zbilut 2005), to assess ritualization in the cleaning task. Put simply, we did not look only at mean acceleration (as in linear statistical approaches), but rather we examined how the acceleration signal evolved over time and measured the signal's rates of recurrence and predictability. RQA is a useful tool for measuring any dynamics, for example, the height of the waves at a beach. Imagine placing a pole in the water at a specific distance from the beach and recording the highest point each wave reached on the pole. After one minute, you could see that wave height is dynamic (it changes over time) but also that wave height recurs in predictable patterns. Now imagine measuring the amplitude of the whole wave (instead of only the peaks) and then comparing it to the amplitudes of all of the other waves during the same minute. The result is a matrix of amplitude points, some of which will be identical; that is, some of them will repeat over time in predictable patterns (see Figure 18.2(a)). This is exactly what RQA measures: how much a signal is recurrent (repetitive) and how much it is deterministic/predictable (rigid in our terms). Figure 18.2(b) illustrates the usage of recurrence plots in the current study (see also Lang et al. (2016) for an application of RQA in a different context). Finally, to quantify redundancy, we measured how much time and how many movements people used while cleaning the art object. After participants finished cleaning, we informed them that they were not selected to give the speech and asked them to fill out a final questionnaire about individual religiosity, age, sex, gender, and perceived anxiety during the study. In summary, we moved from ethnographic observations to a specific experimental hypothesis: we expected that participants in the high-anxiety condition (public speaking) would display more rigid, repetitive, and redundant movements (ritualized behavior) while cleaning an object, compared to participants in the low-anxiety condition.

Nevertheless, what looks like a clear and logical sequence of decisions was actually a product of many discussions and a lot of trial-and-error testing. Although we determined the main contours of our manipulation, manipulation check, and the types of measuring devices at the very beginning of the study, the way we calibrated these elements changed substantially from the first pilot testing to the final data-collection procedure. For example, we first designed our manipulation as a fake TV studio, pretending that we were shooting a video of the participants, which in the high-anxiety condition would later be evaluated by experts, but would only be for our own use in the control condition. This manipulation employed five researchers in total: a cameraman, an assistant, and three evaluators in the other room, connected via a wireless link. The manipulation worked exceedingly well—participants' heart rates went through the roof! However, the entire setup and procedure stressed participants in both the experimental and control

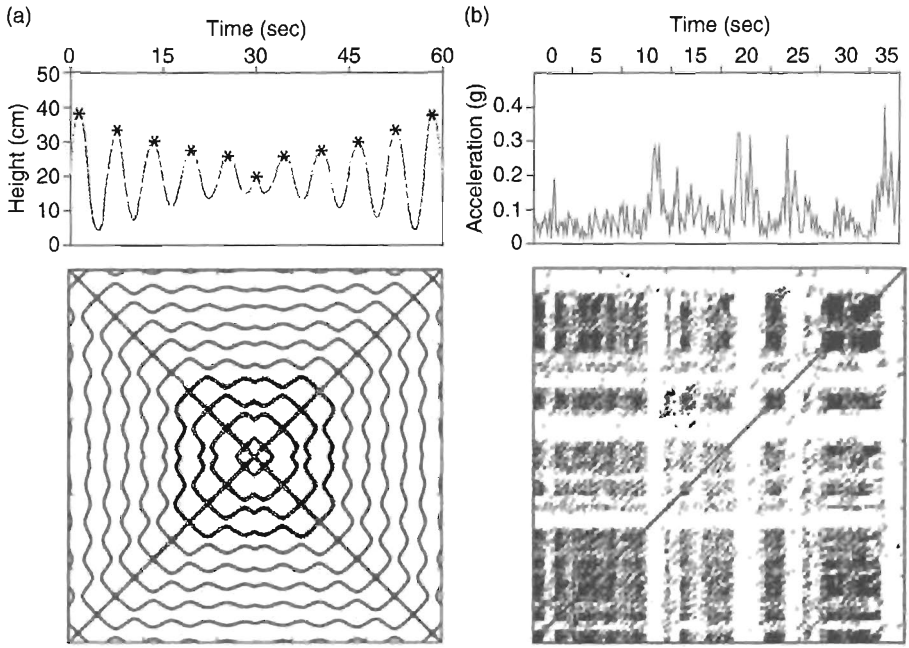

Figure 18.2 **(a)** A simulated wave height across a one-minute measurement. Asterisks mark the peaks of these imaginary waves, illustrating that wave height is dynamic yet recurs in predictable patterns. Plotting the recurrence of all the points comprising waves (and not only peaks) produces harmonic patterns in the recurrence plot (waves are repeating in a sine-like fashion). **(b)** Hand-movement acceleration during cleaning. While much more chaotic than the wave example, there are noticeable recurrent patterns in the signal as illustrated by the recurrence plot below the signal. The black dots represent points where the signal repeats itself, and the chunks of black dots represent the predictability of repetitive patterns (rigidity in repetition).

conditions, thus leaving us without the desired contrast between the two groups. From the first pilot test to the final experiment, the procedure changed substantially: we changed the setting, slightly altered the cover story, and excluded the actual speaking. These changes also saved personnel—the final design only required one research assistant, which is important because it minimized variation potentially arising from a complicated procedure and/or having too many cooks in the kitchen. In other words, the more complex an experimental situation, the more fragile this situation will be. In the end, we opted for a simpler design.

Results and Analysis

As expected, we observed higher heart rates among participants in the high-anxiety condition. These participants also reported that they were more anxious than participants in the low-anxiety condition. These results suggest that our manipulation

was successful, although there was, of course, individual variation within each condition: some participants in the high-anxiety condition were not petrified by public speaking, while some participants in the low-anxiety condition were quite anxious just from the experimental situation. But overall, the experimental manipulation seems to have worked very well.

In terms of our main variable of interest—ritualized behavior—we observed significant differences between the conditions. Participants in the high-anxiety condition had more recurrent points in their acceleration patterns (they engaged in more repetitiveness) and also displayed more predictable patterns (an index of rigidity). We did not observe any statistical differences in redundancy (time spent cleaning and number of movements used during cleaning). These results largely supported our hypothesis—participants in the high-anxiety condition displayed more rigidity and repetitiveness in their behavior, that is, they differed in two out of three ritual characteristics.

When we examined our results in more detail, two concurrent processes seemed to generate our data: one conscious and one subconscious. Participants who reported being more anxious during the speech preparation/thinking phase (independent of condition) spent more time and engaged in more movements when cleaning the object, but this self-reported anxiety did not predict rigidity and repetitiveness. Those characteristics were, however, predicted by an increase in heart rate during the speech preparation/thinking phase (independent of condition). Thus, ritualization provoked by anxiety seems to be driven by two processes: one at a perceivable level such as the action length, which is more conscious, and one driven by physiological processes at a subconscious level of fine-grained motoric action.

Discussion

Our experimental approach helped us tackle the direction of causality between ritual and anxiety: we found that anxious people performed more ritual-like behavior. This conclusion supports ethnographic observations suggesting that rigidity, repetitiveness, and redundancy may be appealing to people in anxiety-inducing situations. Why might this be so? Ritual actions, be they individual or social, rely upon general psychological proclivities that evaluate the efficacy of possible actions, especially in situations where it is not clear what the person should do in order to meet the goal. Due to its repetitiveness and redundancy, ritualization appears to be helpful under uncertain situations, tapping into intuitive evaluations of efficacy (Legare & Souza 2012). Further formalization of these actions via social convention (i.e., cultural rituals) can amplify the perceived efficacy. For example, anthropologist Roy Rappaport (1999) emphasized the invariable, formal aspects of rituals as one of their most important elements, suggesting that the rigidity of this behavior allows certain types of actions to be stabilized and perpetuated, giving them an aura of eternal efficacy ("it has always been done this way"). Together, the combination of repetitive, redundant, and rigid action may seem the most appealing behavior in situations where people lack control.

This conclusion is compatible with theories of the inner workings of the human brain under anxiogenic situations (Hirsh, Mar, & Peterson 2012). In uncertain and uncontrollable environments, the hierarchical cortical structure tries to minimize errors in predicting future affordances. In other words, the human brain works like a predictive machine, trying to guess what will happen next on several perceptual levels in order to prepare appropriate responses (Clark 2013). When such prediction possibilities are limited (as in uncertain situations), anxiety motivates people to take precautionary actions and decrease the possible prediction errors (e.g., choose the well-lit street over a dark alley). However, this is not always possible, especially in situations where people lack control (recall those missile attacks), and anxiety can become overwhelming; hence, a predictable action that is repeated over and over can help decrease anxiety, despite seeming nonfunctional from a pragmatic perspective. Repeating rigid actions, in other words, can help minimize the prediction errors that arise in the hierarchy of human cortical structures. (For further discussion of this mechanism, see Krátký et al. 2016.)

The ultimate goal of this line of research is to investigate whether ritual behavior actually decreases anxiety. In the study under current discussion, we focused on the first logical part of the prediction: Do people perform more rituals when they are anxious? In order to answer the question of whether or not rituals help decrease anxiety, we will need to design follow-up studies where, for instance, both groups experience an anxiety treatment, such as public speaking, and then the members of one group will perform a ritual while the members of the other will not. We then predict that those in the ritual group will exhibit lower anxiety levels (after the ritual treatment) compared to those in the control group. As for the ritual activity, we need a task that involves predictable and repetitive movements (or verbalizations such as those uttered during prayer), which we can compare with more variable movements. This design could help us answer the *why* question about rituals and anxiety.

Note that in the current study, we measured very simple behavioral gestures, behaviors that are substantially detached from naturally occurring ritual behavior. Our simulation of ritual behavior, and our modelling approach, disentangled the complexities of the real world and allowed us to investigate rituals in a controlled manner. While we believe that we successfully simulated a real-world phenomenon in the lab, the ultimate test would be to perform a similar experiment in naturalistic settings. For example, our group works in Mauritius, where we observe ritual behavior in local Catholic and Hindu populations. We could recruit participants from one of those populations and subject them to the public speaking tasks, and we could measure their heart rate, track their hand movements, and record their speech. After the anxiety treatment, participants could then be asked to perform their rituals or prayers. In comparison with a control group not subjected to public speaking, we would expect to see more ritualization in the anxious group. On the one hand, this approach has the advantage of letting people perform their own habituated rituals, which may manifest the link between rituals and anxiety more effectively. On the other hand, such an approach will also bring a lot of variability, which is minimized under controlled laboratory settings. For example, is there a difference in the usual length of prayer between different traditions? Does one of those traditions

comprise more rituals than the other? How frequently, on average, are those rituals performed? Are they performed individually or collectively? These and other questions will be crucial when considering the generalizability of findings from such a study. Moreover, most laboratory studies (including the current one) have been conducted on university students in Western countries (so-called WEIRD populations; see Henrich, Heine, & Norenzayan 2010; Sears, 1986) and thus present a very homogenous sample. In the real world, we will also have to deal with other potential confounding variables such as age, income, socioeconomic status, and so forth. These can be measured and controlled for in statistical models, but they also add more complexity, making it more difficult to interpret the results.

To conclude, there are costs and benefits to both laboratory and field experiments. Ideally, we should strive for the combination of both approaches, moving back and forth between the field and the lab. Only under such methodological collaboration can we be sure that our laboratory simulations represent the real world and our real-world studies are not confounded by unobserved factors. As this chapter hopes to illustrate, we live in exciting times when new methodologies and technologies can help us answer old questions and test classical theories in our disciplines, such as the one by Bronislaw Malinowski (1948/1992).