



Making Behavioral Ethics Research More Useful for Ethics Management Practice: Embracing Complexity Using a Design Science Approach

Jeroen Maesschalck¹

Received: 11 September 2020 / Accepted: 17 July 2021 / Published online: 26 July 2021
© The Author(s) 2021

Abstract

Research on behavioral ethics is thriving and intends to offer advice that can be used by practitioners to improve the practice of ethics management. However, three barriers prevent this research from generating genuinely useful advice. It does not sufficiently focus on interventions that can be directly designed by management. The typical research designs used in behavioral ethics research require such a reduction of complexity that the resulting findings are not very useful for practitioners. Worse still, attempts to make behavioral ethics research more useful by formulating simple recommendations are potentially very damaging. In response to these limitations, this article proposes to complement the current behavioral ethics research agenda that takes an ‘explanatory science’ approach with a research agenda that uses a ‘design science’ approach. Proposed by Joan van Aken and building on earlier work by Herbert Simon, this approach aims to develop field-tested ‘design propositions’ that present often complex but useful recommendations for practitioners. Using a ‘CIMO-logic’, these propositions specify how an ‘intervention’ can generate very different ‘outcomes’ through various ‘mechanisms’, depending on the ‘context’. An illustration and a discussion of the contours of this new research agenda for ethics management demonstrate its advantages as well as its feasibility. The article concludes with a reflection on the feasibility of embracing complexity without drowning in a sea of complicated contingencies and without being paralyzed by the awareness that all interventions can have both desirable and undesirable effects.

Keywords Behavioral ethics · Design science · Ethics training

Introduction

Partly in response to calls from society following high-profile scandals, research on ethics in organizations has flourished in the last two decades. Within that body of literature, behavioral ethics research has become a very prominent strand (Bazerman & Gino, 2012; Trevino et al., 2006). Drawing from findings in other disciplines, including cognitive and moral psychology, it has produced many valuable insights. Particularly interesting are observations about the counterproductive effects of well-intended ethics management interventions. For example, behavioral ethics research showed how mandatory disclosure of conflicts of interests as

well as strict sanctioning systems can inadvertently generate unethical behavior (Bazerman & Tenbrunsel, 2011). With insights like these, behavioral ethics research aims to offer advice that can be used by practitioners to improve ethics management (Haidt & Trevino, 2017). However, as this article will argue, behavioral ethics research does not deliver on this promise. It does not sufficiently focus on interventions that can be directly designed by management. The typical research designs used in behavioral ethics research require such a reduction of complexity that the resulting findings are not very useful for practitioners. Worse still, attempts to make behavioral ethics research more useable by simplifying recommendations are potentially very damaging.

The article’s central proposition is that behavioral ethics research will only become genuinely useful for practice if its current research agenda that takes an ‘explanatory science’ approach is complemented with an agenda that takes a ‘design science’ approach. Proposed by van Aken (2004), who built on the work of Simon (1969), the design science

✉ Jeroen Maesschalck
Jeroen.maesschalck@kuleuven.be

¹ Leuven Institute of Criminology, KU Leuven, Hooverplein 10, Box 3418, 3000 Leuven, Belgium

approach is increasingly seen as a promising way to make management research more relevant (Hodgkinson & Starkey, 2011). By focusing research on finding solutions for field problems, instead of developing causal models to explain phenomena, this approach offers more direct guidance for practitioners. Thus if the aim is to use research to improve the practice of ethics management, then the current explanatory science approach to behavioral ethics is useful, but will not be enough. Huff et al., (2006, p. 414) explain this using an analogy. Practice-oriented management research, they state, relates to behavioral and social sciences as medicine does to life sciences. While fundamental biological research about cells might generate knowledge that is essential to develop new medical treatments (Huff et al., 2006, p. 416), we would not expect those treatments to come from that research directly. For their development, we would expect a complementary research program aimed at the design and testing of those treatments. Arguing that the current explanatory approach to behavioral ethics research somewhat resembles fundamental research in biology, this article will propose an agenda for a complementary research program that aims to develop useful ethics management interventions using insights from behavioral ethics research. Specifically, this research agenda aims to produce recommendations in the form of field-tested ‘design propositions’ that contain generic solutions (van Aken, 2004; Van Aken & Romme, 2012) to actual problems in the field of ethics management. The article will propose to formulate those design propositions in the language of the ‘CIMO-logic’. Inspired by Pawson and Tilley (1997)’s classic on realist evaluation, Denyer et al. (2008) state that design propositions should specify in which *Context* an *Intervention* generates, through a *Mechanism* (or series of mechanisms), particular *Outcomes*. Hence, this approach moves beyond the usual focus on the intervention and its outcome, and takes a particular interest in the underlying mechanism(s) and the context. Rather than simply understanding ‘whether’ an intervention works, the approach aims to understand ‘why’ it works and under which circumstances (Pawson & Tilley, 1997). Thus, in contrast to the experimental research designs that are typical for behavioral ethics research, this approach chooses to embrace complexity instead of reducing it.

The article starts out with a brief discussion of the contributions of behavioral ethics research to the practice of ethics management. It then shows how its usability is hampered by three barriers and argues that these can only be partially addressed by improvements within the current explanatory science approach to behavioral ethics research. This is followed by an explanation of the design science approach and then an illustration that demonstrates how the approach helps to generate design propositions that are of real use to practitioners. The subsequent discussion of the contours of a design science research agenda for ethics management

shows the advantages as well as the feasibility of this project. The article concludes with a discussion of the risks of embracing complexity in a design science approach to ethics management. It argues that it is possible to do so without drowning in a sea of complicated contingencies and without being paralyzed by the awareness that all interventions can have both desirable and undesirable effects.

Behavioral Ethics Research’s Contribution to the Practice of Ethics Management

Behavioral ethics research studies “systemic and predictable ways in which individuals make ethical decisions and judge the ethical decisions of others, ways that are at odds with intuition and the benefits of the broader society” (Bazerman & Gino, 2012, p. 90). It does not focus on unethical behavior by actors who consciously and deliberately choose to do so. Instead, it focuses on “ordinary unethical behavior” (Gino, 2015) committed by actors who usually value morality. Typically, a distinction is made between two types of such behavior (Bazerman & Gino, 2012). First, in a surprisingly large number of situations people seem to behave in ways they know to be unethical, but without having originally planned to do so and without being aware of the forces that guided them to that behavior. Second, there also seems to be a considerable number of situations in which people behave unethically without even being aware that they are doing so. Both cases concern a gap between intended and actual behavior. Inspired by Simon (1957)’s concept of ‘bounded rationality’, Chugh et al. (2005) describe this phenomenon as “bounded ethicality”. It is hard for people to be fully ethical all the time. Their ethicality is ‘bounded’, particularly by their tendency to protect their self-view as ethical (Chugh et al., 2005, p. 80). Most people want to see themselves and to be seen by others as ethical. However, they find that this need conflicts with their self-interested motivation to gain, for example from cheating, if such an opportunity arises (Mazar et al., 2008). Tensions like these generate “ethical blind spots” (Bazerman & Tenbrunsel, 2011; Sezer et al., 2015). For example, people’s unconscious attitudes can lead them to act against their own ethical values, e.g. by implicit discrimination (Bertrand et al., 2005). Temporal distance can also be a source of blind spots (Sezer et al., 2015, p. 78). For example, after having behaved unethically, people tend to rationalize their behavior or to ‘forget’ the moral rules they broke (Shu et al., 2011). Behavioral ethics research not only studies these internal psychological mechanisms, but also addresses situational as well as social factors that lead to unethical behavior (Bazerman & Gino, 2012). Situational factors like physical characteristics of the environment (e.g. the amount of litter) or an abundance of wealth (Gino & Pierce, 2009) have indeed been found to impact

the incidence of unethical behavior. Even more interesting for organizational research are social factors. People's actions are influenced by the actions of those around them. For example, when exposed to unethical behavior by an in-group member, they will tend to align with that behavior (Gino et al., 2009).

Such findings are useful for practitioners working in the field of ethics management (e.g. ethics officers, auditors, investigators, compliance officers, etc.), henceforth 'ethics practitioners'. Particularly interesting are observations about the counterproductive effects of well-intended ethics management interventions. A well-known example is the research showing that the mandatory disclosure of conflicts of interest by advisors can be counterproductive (Bazerman & Tenbrunsel, 2011, pp. 115–117). Having properly warned their clients of their conflict of interest, advisors seem to feel licensed to be less objective in their advice (Cain et al., 2011; Loewenstein et al., 2011). This is an example of the more general mechanism of moral compensation or moral licensing: having behaved in an ethical way in the past can act as a 'license' for unethical behavior in the future, as if "being good frees us to be bad" (Merritt et al., 2010, p. 344). This moral licensing mechanism also has other counterintuitive implications. It helps to explain the paradox that people with a strong moral self-concept ('organizational saints' (Ashforth & Lange, 2016)) are sometimes more likely to commit unethical behavior than others. This offers an important qualification to the self-evident advice that the more serious a job candidate is about ethics, the more ethical he or she will eventually behave on the job. The moral licensing mechanism also helps to explain possible counterproductive effects of ethics training: having done the training, trainees might feel licensed to be less concerned about ethics afterwards in the way that Bohnet (2016, p. 53) describes for diversity training programs. These are just a few examples of insights from behavioral ethics research with important implications for practice. However, practitioners who really want to use this research as a basis for designing ethics management systems will either find it very difficult to deduce practical advice from the existing studies or, when they do find advice, they might be misled by its deceptive simplicity. The next paragraph addresses these barriers to usability in behavioral ethics research.

Why the Explanatory Science Approach to Behavioral Ethics Research will Always be Weak on Usability

Ethics practitioners who, inspired by the fascinating and sometimes counterintuitive findings of behavioral ethics research, want to use this research as a basis for designing

ethics management interventions will face at least three barriers.

First, one very obvious problem is that most behavioral ethics research focuses on explaining (un)ethical behavior, not on explaining how interventions can impact that behavior. Many of the studied explanatory variables are not ethics management interventions that can be designed by management. A case in point is survey research about the impact of ethical climate (Victor & Cullen, 1988) and ethical culture (Treviño et al., 1998). While both do indeed help to explain (un)ethical behavior (Treviño et al., 2006), their usability for practice is limited. As for ethical climate, Kuenzi et al. (2020) point out that Victor and Cullen (1988)'s Ethical Climate Questionnaire only measures perceptions of ethical decision making, not procedures and instruments of ethics management. Hence, it will be difficult to translate the findings from ethical climate research into actual ethics management interventions. As for Treviño et al. (1998)'s 'ethical culture' concept, this amalgamates factors that can be directly designed by management (e.g. ethics codes, reward systems and training programs) with factors that are less susceptible to direct intervention by management such as peer behavior and informal ethical norms. In an explanatory logic, such broad concepts can be useful, but in a research agenda that aims to support the design of interventions, they are not very helpful. Because they do not distinguish between intervention and context, they do not really help to understand which interventions would work in particular contexts.

Second, the usability of the findings is also limited by methodological choices and particularly the preference in behavioral ethics research for (quasi-)experimental research designs. Typically set in the laboratory or in online settings, such designs allow researchers to control for confounding variables (Mitchell et al., 2020, p. 6) and thus to reduce complexity. Yet these designs are often weak on ecological validity, with limited relevance for management interventions in the real world (Houdek, 2019). For example, the consequences of cheating in an experiment (e.g. losing a game) are of a very different nature than those of cheating in real life (e.g. being fired). In real life, even relatively subtle differences in the context can explain why the same intervention is sometimes effective, sometimes without any effect and sometimes even counterproductive, as complexity theory (Byrne & Callaghan, 2014) would predict. Such subtleties are lost in rigid experimental research and replication. A case in point is Verschuere et al. (2018)'s replication of Mazar et al. (2008)'s finding about the effect of moral reminders. Mazar et al. (2008) had found that students were less likely to cheat on a task after they had recalled a moral reminder (the Ten Commandments). This study was the basis for recommendations of Ayal et al., (2015, p. 739) and others to use moral reminders as a means to reduce unethical

behavior. Verschuere et al. (2018) aimed to replicate Mazar et al. (2008)'s finding by means of a collaborative multisite replication study in 25 different labs across the world, using a detailed protocol to ensure similar conditions in all 25 labs. They could not replicate Mazar et al. (2008)'s finding: the moral reminder did not reduce cheating. While this replication study might be important from an explanatory science perspective, its use for practitioners is very limited. It showed practitioners that this particular intervention in those particular standardized laboratory settings had no effect. Much more useful for practitioners would have been a multisite study that would have looked at various different versions of the intervention and that would have taken an interest in the role of various conditions, instead of standardizing them away with a rigid protocol.

The third problem does not have to do with the research itself, but with the recommendations provided on the basis of it. Because of the two previous obstacles, recommendations based on behavioral ethics research can only be formulated very cautiously. Researchers should emphasize that (quasi-)experimental research can only generate decontextualized recommendations that are merely relevant for the controlled circumstances of the experiment. The effect that was identified in the experiment might be absent or even reversed in the real world because of differences in context. While scientifically correct, such cautious recommendations will not be very helpful for practitioners. With both the researchers and their practitioner-readers unsatisfied, researchers are sometimes tempted to formulate recommendations with more confidence than their study actually allows for. The editorial policy of many academic journals to stimulate or even require an 'implications for practice' section at the end of the article further adds to this temptation. As a consequence, researchers might translate their decontextualized findings into general recommendations, ignoring the possible impact of contextual factors and also ignoring contradictory recommendations formulated in other research. These unambiguous quick fixes then turn into 'proverbs' (Simon, 1946) to be followed by practitioners. A case in point is the well-known study of Shu et al. (2012) on the impact of signing a self-report form at the beginning instead of at the end. On the basis of both laboratory and field experiments, they had found that if people are asked to sign before, rather than after an opportunity to cheat, they are less likely to actually cheat. Shu et al. (2012) hypothesized that this is because the act of signing makes ethics salient. This led to many recommendations to managers such as Treviño et al., (2014, p. 639)'s advice "to ask employees to sign a form before important events, such as the beginning of the annual compliance process, stating that employees will read the code and will answer compliance questions truthfully". However, in a series of studies, Kristal et al. (2020) failed to replicate these findings: signing at the beginning does

not seem to be more effective in reducing dishonesty than signing at the end. Kristal et al., (2020, p. 7106) concluded with the recommendation "that practitioners take this finding out of their intervention "tool-kit" as it is unlikely to increase honesty". While this is an uplifting example of the self-reflective capacity of the explanatory approach to science, it is also an example of the risks of attempting to generate 'quick fixes' from experimental research. In fact, Kristal et al. (2020) run the risk of repeating that error with their advice, following the failed replication, to remove the recommendation from the toolkit. While the place of signing might not have an impact in those particular controlled experimental conditions, their research does not say anything about the possible synergetic effect of combining signing upfront with other measures to increase salience of ethics. Practitioners or practice-oriented researchers simply adopting Kristal et al. (2020)'s advice to remove this from their intervention toolkit, would unnecessarily exclude such potential synergies. A second illustration of the dangers of quick recommendations is offered at the conclusion of Belle and Cantarelli (2017)'s review of (quasi-)experimental research on the causes of unethical behavior. Having selected 19 experimental studies for a sophisticated meta-analysis (measuring Hedges's g effect sizes and checking for publication bias), they conclude that monitoring individuals' behavior reduces their unethical behavior. This leads them to advice for public administration practice: "meta-analytic evidence that monitoring reduces unethical behavior provides support for government initiatives around the globe aimed at increasing monitoring of public service providers through transparency and openness" (Belle & Cantarelli, 2017, p. 335). However sound that conclusion might be supported in the meta-analysis, it does overlook the many risks and limitations of monitoring that are identified elsewhere in the behavioral ethics literature (e.g. Moore & Gino, 2013, p. 17) and beyond. Monitoring might generate disengagement or even reactance to the moral message (Zhang et al., 2014, p. 69) and mandatory transparency might generate unethical behavior (Loewenstein et al., 2011). As with the previous example, this illustration shows that the problem here is not so much the usability of behavioral ethics research itself, but the way in which it is translated into recommendations. Combining overconfidence in (quasi-)experimental research designs or the technique of meta-analysis with a questionable way of moving from findings to recommendations, this style of reporting can be seriously damaging. Such overconfident recommendations can lead well-meaning practitioners on the course to ineffective or even counterproductive results. Ironically, practitioners' confidence in the 'scientific' nature of the advice, might lead them to suppress the warning signs that come from their own intuitions and experience. Practitioners are used to complexity in their daily practice. Most of them are well-aware that well-intended interventions can

also have counterproductive effects. It would be painful if they would ignore that awareness. That would be research-informed management at its worst: replacing experience-based professional knowledge by an unfounded trust in simplified ‘research-based’ recommendations. When failure then eventually comes, this might undermine practitioners’ trust in the very possibility of using research to improve ethics management.

To some extent, these barriers to usability can be addressed within the explanatory science approach to behavioral ethics. The first barrier can be addressed by stimulating researchers to focus on explanatory variables that can be directly manipulated by management. While helpful, this will still generate generic recommendations about standardized interventions, not the creative context-sensitive solutions that can be produced by design science’s reflective cycle (see below). As for the second barrier, some methodological improvements within the confines of explanatory science can increase usability. Behavioral ethics researchers themselves offer various ways to do so: field experiments in real-life organizations (Brief & Smith-Crowe, 2016), “mixed context” experiments in both the lab and the field (Zhang et al., 2014, p. 74), long-term field experiments that can help to assess habituation to interventions (Houdek, 2019, p. 51), and conceptual instead of literal replication studies (Amir et al., 2018). Because such research designs improve ecological validity and allow for more complexity, they are indeed likely to generate insights that are more relevant for practice. Yet as long as they remain within the logic of explanatory research, their usability will remain weak because they still reduce context to a set of confounding variables that should be controlled for and because they do not focus on developing and field-testing solutions. To address the third barrier, behavioral ethics researchers could be more cautious in their recommendations for practice. While such modest recommendations would replace dangerously overconfident proverbs, they would also be much less relevant for practice. This might leave practitioners disillusioned about the possibilities to research-based ethics management. As the next section will argue, the design science approach instead allows researchers to make genuinely useful recommendations without having to produce unfounded and potentially damaging proverbs.

A Design Science Approach to Ethics Management

For the development of a design science approach to ethics management, this article draws from the work of van Aken and colleagues (e.g. Denyer et al., 2008; van Aken, 2004, 2005; van Aken & Romme, 2009, 2012). In a seminal article, van Aken (2004), drawing from Herbert Simon Simon

(1969)’s classic *‘the sciences of the artificial’*, contrasts the design science approach with that of explanatory science. The latter aims to explain the world as it is, values knowledge as an end in itself and typically aims to produce causal models (van Aken, 2004; Van Aken & Romme, 2012). In contrast, design science research invites researchers to think as engineers, architects or medical doctors. It is driven by field problems and uses knowledge instrumentally to help solving these field problems. Hence, its focus is more on finding solutions (or rather more generic ‘solution concepts’) than on describing and explaining problems (Romme, 2003). It aims to produce field-tested and theoretically grounded ‘design propositions’. A design proposition (or technical rule) is a “chunk of general knowledge, linking an intervention or artefact with a desired outcome or performance in a certain field of application” (van Aken, 2004, p. 228). In the context of management, design propositions usually do not take the form of an ‘algorithmic’ prescription: “if you want to achieve Y in situation Z then perform action X” (van Aken, 2004, p. 227). They are typically heuristic instead of algorithmic, taking the form “if you want to achieve Y in situation Z, then something like action X will help” (van Aken, 2004, p. 227). By merely claiming that ‘something will help’, rather than that something should be done, this formulation is more cautious and thus less susceptible to the misplaced confidence that can be found in the practical recommendations of some behavioral ethics research (see above). In addition, by referring to ‘something like action X’ instead of simply ‘action X’, this format implies that the prescription does not offer a standard recipe that should simply be applied in all cases. It is not to be used as an instruction but as a “design exemplar” (van Aken, 2005, p. 27) that still has to be translated to specific circumstances. It is “a proposal to professionals to use this generic solution concept to design a variant of it to address the specific field problems of this type” (Van Aken et al., 2012, p. 64).

The development of a design proposition would typically start from a general statement. For example, it might claim that, in a particular context, a particular type of ethics training might help to reduce particular types of unethical behavior. This general statement would then be complemented with more specific information, drawn from empirical research, e.g. about underlying mechanisms or relevant contextual factors. Hence, rather than simple and straightforward recipes, design propositions tend to become elaborate, possibly in the format of an article, a manual or even a whole book (van Aken, 2005, p. 23). Although design propositions are meant to be an important source for practitioners, the expectation is not that they will be the only source. Even when they provide much detail, they will never look like an instruction manual that has to be followed to the letter. However detailed they are, design propositions are still general statements that have to be translated by practitioners to

a specific situation. For this translation, practitioners will combine the information from design propositions with information from other sources such as their “experience, creativity and deep understanding of [their] local setting” (van Aken, 2004, p. 238).

How does this all translate into a design science research agenda that can really support ethics management? More specifically, how can design propositions in the field of ethics management be developed and improved by research? The current article follows Denyer et al. (2008)’s recommendation to formulate design propositions in the language of the ‘CIMO-logic’. Inspired by Pawson and Tilley (1997)’s classic on realist evaluation, Denyer et al. (2008) state that design propositions should specify in which *Context* an *Intervention* generates, through a *Mechanism* (or series of mechanisms), particular *Outcomes*. Hence, this approach moves beyond the usual focus on the intervention and its outcome, and takes a particular interest in the underlying mechanism(s) and the context. Rather than simply understanding ‘whether’ an intervention works, the approach aims to understand ‘why’ it works and under which circumstances (Pawson & Tilley, 1997). As for the latter, this approach does not consider the context of the intervention as something to be controlled for, but as something of which the impact should be included in the analysis. Any claim that an intervention leads to a particular outcome only becomes useful for practice when it is complemented with a specification of the context and its possible role in this process. The other contribution of the CIMO-logic is its focus on mechanisms that link the intervention with the outcome. It opens the black box and aims to explain why, i.e. through which process, the intervention generates a particular outcome. Behavioral ethics research of course provides a rich source of such mechanisms as well as underlying theories that help to explain why a particular intervention generated a particular outcome.

To realize this research agenda, lab or field experiments will not be sufficient. It requires observations of real interventions in real-life contexts: “holistic field testing” as van Aken (2005, p. 29) would call it. The most obvious, but certainly not the only possible (van Aken & Romme, 2009, pp. 9–10), research strategy to acquire this knowledge would be the multiple case study (van Aken, 2005, p. 24). This would generate detailed case descriptions of real-life interventions, their context, the mechanisms they trigger and the outcomes they generate. Such case studies would often combine various methods. They would include qualitative research such as observation or critical incident interviews (Butterfield et al., 2005), but might also imply the collection of quantitative data, for example to quantify patterns in the outcome over time. Through such a multiple case study, knowledge would gradually evolve and design propositions would be refined: “step by step one learns how to produce the desired

outcomes in various contexts” (van Aken, 2005, p. 25). During this cycle, researchers might limit their role to observing and analyzing, but they might also be involved in the actual design of the intervention. As such, they contribute to what van Aken describes as the reflective cycle: designing and implementing an intervention in a case, reflecting on the results and then developing design propositions to be tested and further refined in subsequent cases (van Aken, 2004, p. 229). Importantly, while such multiple case studies might include detailed descriptions and holistic field testing, the aim is not simply to accumulate descriptions of particular interventions in particular cases. Instead, the aim is to develop generic design propositions, identifying patterns that return in various cases (van Aken & Romme, 2009, p. 8).

An Illustration: What the Design Science Perspective can Add to a Typical Behavioral Ethics Study

To illustrate how a design science perspective can complement the explanatory science approach of current behavioral ethics research, this paragraph uses Zhang et al. (2018)’s study on the impact of creativity on moral insight as an example. In a series of experiments, Zhang et al. (2018) found that people who face an ethical dilemma exhibit more propensity to moral insight when they are asked to answer the question ‘what could I do?’ than when they are asked ‘what should I do?’. They argue that the ‘could mindset’ that results from being asked ‘what could I do?’ generates divergent thinking (and hence creativity), which in turn increases the propensity to reach moral insight. They also offer an explanation for the apparent contradiction between this finding and earlier research that found creativity to actually increase the propensity to act unethically (Gino & Arieli, 2012). Specifically, they argue that the impact of creativity depends on the type of dilemma. In situations where people face ‘right vs. right’ dilemmas, creativity will indeed increase moral insight, but in situations where people have to choose between ethical as well as self-interested options (‘right vs. wrong’ dilemmas), creativity might have no effect or even an opposite effect. In the latter situation, creativity can help people to justify self-interested options and to come up with ways to get away with unethical behavior (Zhang et al., 2018, p. 880). To some extent, Zhang et al. (2018)’s study does make efforts to address the three barriers to usability identified above. It proposes an intervention that can be directly applied by managers and trainers, it addresses contradictory findings by identifying an important contextual factor and its recommendations are formulated cautiously. Yet in spite of all this, the practical relevance of Zhang et al. (2018)’s study remains limited. It is useful

now to illustrate how a design science approach can build on the work of Zhang et al. (2018) to develop genuinely useful advice, based on a rich understanding of potential interventions, of varying contexts and of potentially contradictory mechanisms. These are now each discussed in turn.

First, while Zhang et al., (2018, p. 879) offer some brief reflections on how their insights might be translated in specific interventions such as training and counseling, a design science perspective would place such translation at the core of its research agenda. It would, for example, specify in much more detail how a training aimed at a 'could mindset' could actually look like. For this, it might turn to other sources such as Mumford et al. (2008). For example, the training could contain group discussions of cases (Mumford et al., 2008) in which trainees learn to use forecasting as a sensemaking strategy (Thiel et al., 2012, p. 57). It would also explicitly address many other aspects of the training: how the training would be announced, what the profile of the trainer would be, how training groups would be composed, etc. The impact of these aspects of the training would be formulated as design propositions using the CIMO logic introduced above. These would be further developments of the hypothesis Zhang et al. (2018) confirmed in the lab: an ethics training that focuses on 'what could I do?' (*intervention*) will, for 'right vs. right dilemmas' (*context*), generate the ability to consider multiple solutions or 'divergent thinking' (*mechanism*), which in turn will increase the propensity to reach moral insight (*outcome*) (Zhang et al., 2018, p. 863).

Second, a design science perspective would strongly emphasize the importance of context and particularly the awareness that the same intervention can have opposite effects depending on variations in that context. Zhang et al. (2018) did this to some extent by hypothesizing that the creativity-inducing training they propose might generate unethical behavior for 'right vs. wrong' dilemmas. Yet they did not really consider the implications of this understanding for the design of an ethics training. Given that in real life, trainees will be confronted with all kinds of dilemmas, how can we avoid that they use the creative thinking skills they learned in the training to justify unethical behavior in 'right vs. wrong' dilemmas? Simply explaining this to the trainees is unlikely to be sufficient. People are not always immediately aware which type of dilemma they face. Moreover, even if trainees are aware that they face a 'right vs. wrong' dilemma, they might still deliberately choose to use their acquired creative thinking skills to justify unethical behavior in their own interest. A design science perspective would address these concerns by looking for solutions. One option could be to complement the creativity-inducing group discussions with a more one-directional message from the trainer that emphasizes sanctions that would follow if people behave unethically. Previous research has indeed shown that the

communication of the threat of sanctions can have a deterrent effect on unethical behavior (Nagin, 2013), also in the context of organizations (Rorie & West, 2020). In the CIMO-logic, this part of the training (*intervention*) is expected to generate deterrence (*mechanism*) for 'right vs. wrong' dilemmas (*context*) and thus reduce the risk of unethical behavior (*outcome*). Ideally, this expansion of the training would ensure that the acquired divergent thinking skills are only used for 'right vs. right' dilemmas.

Third, the latter illustration also hints at another important advantage of the design science approach: its openness for the complexity of ethics management in real life and particularly for the possibility of contradictory mechanisms. While one might hope that communication about sanctions indeed deters unethical behavior, it might also have other effects. The communication of the threat of sanctions (*intervention*) can also crowd-out people's intrinsic motivation to behave ethically (*mechanism*) (Bazerman & Tenbrunsel, 2011, pp. 109–113; Tenbrunsel & Messick, 1999), which might reduce the trainees' propensity to apply moral insight (*outcome*) in right vs. right dilemmas (*context*) and thus undermine the original intention of the training. Hence, the challenge now is to design a training that can generate both the divergent thinking mechanism in order to increase moral insight and the deterrence mechanism to prevent unethical behavior, without those mechanisms undermining each other. Of course, this argument for combining complementary but contradictory approaches to ethics management such as compliance and integrity (Paine, 1994; Weaver & Treviño, 1999) or values-oriented and structure-oriented (Zhang et al., 2014) is not new to the behavioral ethics literature. Yet in spite of the many claims in that literature that such strategies should be combined in practice, there is actually very little specific theorizing, let alone empirical research, on the specifics of this challenging balancing act. This is not surprising given the explanatory science perspective of the behavioral ethics tradition and its preference for the experimental method. From such a perspective, a research agenda that focuses on how contradictory messages generate contradictory mechanisms does not seem very attractive. In contrast, from a design science perspective this would be an obvious research avenue, as it concerns the contradictions faced by practitioners in the real world. Design science research would not be content with the mere identification of basic mechanisms such as divergent thinking or deterrence in controlled conditions. Instead, it would focus on how those mechanisms interact with each other in the complexity of a real-life environment. It would develop design propositions which would specify how particular combinations of training elements can, in particular contexts, generate both the divergent thinking and deterrence mechanisms in such a way that they do not undermine and instead perhaps even reinforce each other. Those initial CIMO configuration

would then be field-tested and further developed through the iterative, reflective cycle described above.

A Design Science Research Agenda on Ethics Management

The previous paragraph offered merely one illustration of how a design science approach could generate useful recommendations for ethics practitioners. The broader research agenda could be structured around types of intervention such as ethics training or whistle-blowing arrangements. Those research findings could then be presented somewhat like a drugs catalogue in medicine. It would present the different forms and shapes of the intervention type (e.g. ethics training) and then present propositions about the various ways in which it might impact the outcome through various mechanisms and given a particular context. As in medicine, the development of this catalogue would be a permanent process, continuously adding more complexity and depth.

Like medical drug catalogues, the ethics management intervention catalogue would pay ample attention to unintended consequences. In medicine these are usually described as side-effects. Yet as Reynolds and Bae (2019, p. 248) rightly note, this a bit of a misnomer as it misleadingly suggests effects that are weaker. In actual fact, side-effects might be more powerful than the expected effect. While the urge to offer simple fixes might drive researchers to ignore those side-effects in their recommendations to practice, the discussion above illustrated that this is not helpful for ethics practitioners. In medicine, even when there is only a small chance that a particular side-effect of a particular drug treatment might occur, we still hope that our physician's drug catalogue addresses this. A research agenda that intends to be useful for ethics management practitioners should do the same. Specifically, for each intervention it should generate at least three types of information about potential side-effects. First, as in medicine, a practice-oriented research agenda would have to specify the conditions under which particular side-effects are more likely. An example of this, provided in the illustration in the previous paragraph, is the expectation that creativity-inducing training might inadvertently lead to unethical behavior in 'right vs. wrong' dilemmas (Zhang et al., 2018). This is similar to a drug catalogue's prediction that a particular drug treatment will have adverse effects for patients with a particular medical condition. Second, research should identify other interventions that can prevent or reduce undesirable side-effects of ethics management interventions, in the same way as medical drug catalogues propose medication to alleviate the side-effects of other medication. The suggestion above to complement creativity-inducing group discussions with an emphasis on the threat of sanctions for unethical behavior is an example of such an

attempt to alleviate side effects. Third, the catalogue should also contain information about the right 'dosage' of the intervention. It will be important to administer enough of the intervention, without administering too much of it (Pierce & Aguinis, 2013) so as to avoid side-effects. Referring back to the illustration: what would, given a particular context, be the right dosage of the creativity-inducing group discussions on the one hand and of the communication of the threat of sanctions on the other? In practice, this will require a skilled ethics trainer who finds a context-specific balance between both approaches. The research agenda proposed here provides a path to translate this professional trainer expertise in design propositions that can then be continuously further field-tested, specified and nuanced.

Conclusion

With "bounded ethicality" as one of its core concepts, behavioral ethics draws from the seminal work of Herbert Simon on the bounded rationality of decision making (Simon, 1957). As argued in this article and elsewhere (Barzelay, 2019; Moynihan, 2018), it is useful to combine that understanding of the importance of psychological insight with another important insight from Simon's work: the idea of management as a design science, generating knowledge directly useful for practice (Simon, 1969). Specifically, this article proposed to complement the current behavioral ethics research agenda with a design science research agenda aimed at generating design propositions that are useful for ethics practitioners. It argued that, as long as current behavioral ethics research remains within the explanatory science perspective, its usefulness will remain limited. Focusing on variables that can be changed by management, improving ecological validity of the experimental designs and being more modest in recommendations: it will simply not be sufficient. Rather than being more modest, researchers should be more ambitious in the sense that they should aim for advice that is also useful for the "swamp of practice" (van Aken, 2004, p. 225). Such advice, this article argued, can be formulated in the form of design propositions. Those propositions will typically not take the form of algorithm-like instructions that should be diligently applied by practitioners. In fact, the article argued for abandoning the hope that research would ever be able to provide ethics practitioners with easy fixes that always work. Instead, researchers can use the CIMO framework to formulate design propositions that identify the mechanisms and thus outcomes generated by an intervention under study in various contexts. Moreover, because the design science approach does not think of interventions as standardized off-the shelf packages, this agenda also allows for much flexibility and hence complexity in the design of interventions. The emphasis is on replicating

mechanisms, not on replicating interventions (Barzelay, 2019, pp. 106–107). Hence, instead of diligently imitating an intervention that has worked elsewhere, practitioners should reflect on how those mechanisms might play out in their own context. Their conclusion might be that, in their context, the intervention should be adapted in order to generate the same mechanisms and hence desirable outcomes. In sum, a design science research agenda will provide practitioners with nuanced design propositions providing useful knowledge that they can translate to their own situation. Somewhat paradoxically, these nuanced, complex propositions will be more useful than the simple quick fixes currently offered on the basis of decontextualized research.

The proposed research agenda's emphasis on complexity will not only increase the relevance of research for practice. It can also open up important avenues for theory building that are underexplored by current behavioral ethics research. A case in point is the recommendation, formulated in the illustration discussed above, to design and test an ethics training that combines creativity-inducing discussions with a communication that emphasizes the threat of sanctions. Really understanding the complex dynamics that will follow from such a paradoxical intervention will require sophisticated theorizing that goes beyond current understanding in behavioral ethics and might draw, for example, from theorizing on organizational paradox (Smith & Lewis, 2011; Smith et al., 2017). This example shows that an emphasis on practical relevance does not have to exclude theory building, but can in fact compel researchers to expand their theoretical scope. Another example of a preference of current behavioral ethics research that would be challenged by the aim to be practically relevant, is its focus on 'ordinary unethical behavior' committed by actors who usually value morality. While such a focus on one type of behavior might be a smart choice within the logic of explanatory science, it is not tenable for research that aims to be relevant for real-life organizations in which a broad range of types of behavior can occur. Thus, design propositions about the impact of a particular intervention should not only describe its impact on ordinary unethical behavior but also its impact on calculative employee behavior that is deliberately unethical.

Having argued that research can offer more useful advice for practitioners if it allows for more complexity, it is also important now to address two important risks of embracing complexity. The first risk is that of paralysis. Systematically reflecting on all possible undesirable side-effects of ethics management interventions might lead to fear of generating those side-effects and thus inertia. Referring back to the illustration above, warned that both a creativity-inducing training and a sanctions-focused training might have counterproductive effects, practitioners might conclude it to be safest not to organize any training at all. While not implausible, this risk of inertia remains limited because of the design

science approach's pragmatic emphasis on solving field problems. A design science research agenda would focus on thinking about training programs that smartly combine both types of training so that they reinforce each other or at least not cancel each other out. Thus, the energy would be aimed at creatively designing and field-testing solutions in an iterative cycle, not at agonizing over their consequences.

The second risk of embracing complexity is that it might generate a sea of details and complicated contingencies in which practitioners would drown, leaving them disillusioned with the possibility of research-based ethics management. There are at least two ways in which this risk can be reduced. First, it will be important to remain realistic in ambition (Pawson, 2013, p. 85). The research agenda proposed in this article will never generate detailed, unequivocal instructions prescribing exactly what practitioners should do in a particular situation. It will simply never be possible to theorize all possible contingencies that might be relevant. Instead, the proposed research program will produce design propositions that identify the most important mechanisms and outcomes that might be generated by an intervention under various conditions. Those design propositions will never work as cooking book recipes that are simply to be followed. They will always have to be combined with local information and experience. While they will never be the only source of information for practitioners, the aim should be for them to become at least an important source of information (van Aken, 2004, p. 239). With that notion of modesty in mind, it is easier to see how complexity can remain manageable. While the empirical research in this agenda will provide detailed context-specific information about the impact of a particular intervention in a particular context, the ultimate goal of the research is to generate design propositions that offer a chunk of *general* knowledge (van Aken, 2004, p. 228). Those propositions will not include all idiosyncrasies found in case studies, but present the recurring, observable patterns (van Aken, 2005, pp. 31–32). Iterative field testing of interventions in an increasingly growing number of conditions (Pawson, 2013, p. 83), will help to identify those interventions and mechanisms that have the most causal power. The result is knowledge that is complex, but not idiosyncratic. It does leave a considerable degree of doubt and uncertainty (Pawson, 2013, p. 85), and practitioners will still need local knowledge to know what the intervention will do in their specific context. Second, complexity is also constrained because the proposed research agenda is relatively limited in focus. As befits a design science approach, it focuses on interventions rather than on the many other factors that can explain ethical and unethical behavior in organizations. Those other factors only come into play if they are needed to understand why a particular intervention generates a particular outcome. More generally, the conceptual framework offered by the CIMO-logic also helps to

constrain complexity. It forces researchers to translate all observations in the language of those four generic categories. If, over time, researchers would also manage to increase agreement about the concepts used in this CIMO-logic, this would further help to constrain complexity. For example, it would help if researchers could agree on a limited number of outcomes intended by ethics management interventions (Smith-Crowe & Zhang, 2016). An obvious candidate would be something like '(un)ethical decision making and (un)ethical behavior of individuals within organizations'. Researchers would then have to agree who can decide what is considered ethical and what is considered unethical. Will it be the researcher (perhaps using standards from moral philosophy) or institutions (e.g. treaties, legislation, professional codes); or should we avoid taking a normative position by simply focusing on reasons why people behave in ways inconsistent with their own values (Bazerman & Tenbrunsel, 2011, p. 22)? Similar questions can of course be asked when theorizing interventions, mechanisms and context. Even if such debates do not lead to shared definitions, they might at least generate some agreement about the language in which to specify the different concepts. All this would help constraining complexity and developing a genuinely cumulative research agenda. It would thus contribute to the ambition of turning business ethics into the kind of 'cumulative science' Haidt and Trevino (2017) hope it to become.

Assuming that this article's appeal for a design science approach to study ethics management is valid, it is useful to conclude with a reflection on ways in which the publication of such research can be stimulated. It would certainly help if editors of academic publications can be convinced that the proposed approach to research not only generates useful advice, but that it can do so in a methodologically rigorous way (Van Aken et al., 2012, 2016). Indeed, many indicators of methodological quality in explanatory science research are also relevant for design science research: transparency about methodological choices, awareness of the role and possible biases of the researcher, active investigation of rival explanations, etc. (Van Aken et al., 2016, p. 6). Likewise, the actual techniques of data-collection and data-analysis are the same as in explanatory science research and their quality can therefore be assessed in the same way. There are of course novel elements to the proposed approach that might look at odds with conventions for traditional academic publications such as its focus on the design of interventions and on iterative field-testing. Yet those aspects can also be executed in a rigorous way and, as shown above, can provide avenues for theory development that will also interest an academic audience. One way in which academic journals could open up to the proposed approach, is by introducing the format 'design science research add-on'. This is put forward by Van Aken et al. (2016) as an alternative to the 'implications for practice' section that often features at the end of explanatory

science articles. According to Van Aken et al. (2016), such an add-on should improve on those traditional recommendations in three important ways: it should be sufficiently specific about the proposed intervention and the actors that are expected to implement it; its recommendations should rely on a minimum of field-testing of the proposed intervention (e.g. a small pilot implementation); and it should be sufficiently specific about all four components present in the CIMO logic, including the underlying mechanisms. Although they are a useful way of launching this research agenda in traditional academic publication outlets, such add-ons still remain rather modest. A methodologically rigorous study that includes a cycle of field-tests to develop design propositions deserves publication as a full-length academic journal article. Editors of periodicals like the Journal of Business Ethics can play a major role here by actively inviting such design science research, for example in a dedicated section or in special issues. When a series of such studies eventually delivers something that looks like the ethics management intervention catalogue proposed above, then this deserves publication as an academic book. Importantly, such a book would not only be at its place in the academic library, it would also have a good chance to be found on the management shelves in the airport bookshop.

Acknowledgements The author wishes to thank two anonymous reviewers for very helpful comments and suggestions.

Declarations

Conflict of interest There are no conflicts of interest to be reported for this publication.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Amir, O., Mazar, N., & Ariely, D. (2018). Replicating the effect of the accessibility of moral standards on dishonesty: Authors' response to the replication attempt. *Advances in Methods and Practices in Psychological Science*, 1(3), 318–320. <https://doi.org/10.1177/2515245918769062>
- Ashforth, B. E., & Lange, D. (2016). Beware of organizational saints: How a moral self-concept may foster immoral behavior. In D.

- Palmer, K. Smith-Crowe, & R. Greenwood (Eds.), *Organizational wrongdoing: Key perspectives and new directions* (pp. 305–336). University Press Cambridge.
- Ayal, S., Gino, F., Barkan, R., & Ariely, D. (2015). Three principles to REVISE people's unethical behavior. *Perspectives on Psychological Science*, 10(6), 738–741. <https://doi.org/10.1177/1745691615598512>
- Barzelay, M. (2019). *Public management as a design-oriented professional discipline*. Edward Elgar Publishing.
- Bazerman, M. H., & Gino, F. (2012). Behavioral ethics: Toward a deeper understanding of moral judgment and dishonesty. *Annual Review of Law and Social Science*, 8(1), 85–104. <https://doi.org/10.1146/annurev-lawsocsci-102811-173815>
- Bazerman, M. H., & Tenbrunsel, A. E. (2011). *Blind spots: Why we fail to do what's right and what to do about it*. Princeton University Press.
- Belle, N., & Cantarelli, P. (2017). What causes unethical behavior? A meta-analysis to set an agenda for public administration research. *Public Administration Review*, 77(3), 327–339. <https://doi.org/10.1111/puar.12714>
- Bertrand, M., Chugh, D., & Mullainathan, S. (2005). Implicit discrimination. *The American Economic Review*, 95(2), 94–98.
- Bohnet, I. (2016). *What works. Gender equality by design*. Harvard University Press.
- Brief, A. P., & Smith-Crowe, K. (2016). Organizations matter. In A. G. Miller (Ed.), *The social psychology of good and evil* (2nd ed., pp. 390–414). The Guilford Press.
- Butterfield, L. D., Borgen, W. A., Amundson, N. E., & Maglio, A.-S.T. (2005). Fifty years of the critical incident technique: 1954–2004 and beyond. *Qualitative Research*, 5(4), 475–497. <https://doi.org/10.1177/1468794105056924>
- Byrne, D., & Callaghan, G. (2014). *Complexity theory and the social sciences: The state of the art*. Routledge.
- Cain, D. M., Loewenstein, G., & Moore, D. A. (2011). When sunlight fails to disinfect: Understanding the perverse effects of disclosing conflicts of interest. *Journal of Consumer Research*, 37(5), 836–857. <https://doi.org/10.1086/656252>
- Chugh, D., Bazerman, M. H., & Banaji, M. R. (2005). Bounded ethicality as a psychological barrier to recognizing conflicts of interest. In D. A. Moore, D. M. Cain, G. Loewenstein, & M. H. Bazerman (Eds.), *Conflicts of interest: Challenges and solutions in business, law, medicine, and public policy* (pp. 74–95). Cambridge University Press.
- Denyer, D., Tranfield, D., & van Aken, J. E. (2008). Developing design propositions through research synthesis. *Organization Studies*, 29(3), 393–413. <https://doi.org/10.1177/0170840607088020>
- Gino, F. (2015). Understanding ordinary unethical behavior: Why people who value morality act immorally. *Current Opinion in Behavioral Sciences*, 3, 107–111. <https://doi.org/10.1016/j.cobeha.2015.03.001>
- Gino, F., & Ariely, D. (2012). The dark side of creativity: Original thinkers can be more dishonest. *Journal of Personality and Social Psychology*, 102(3), 445–459. <https://doi.org/10.1037/a0026406>
- Gino, F., Ayala, S., & Ariely, D. (2009). Contagion and differentiation in unethical behavior: The effect of one bad apple on the barrel. *Psychological Science*, 20(3), 393–398. <https://doi.org/10.1111/j.1467-9280.2009.02306.x>
- Gino, F., & Pierce, L. (2009). The abundance effect: Unethical behavior in the presence of wealth. *Organizational Behavior and Human Decision Processes*, 109(2), 142–155. <https://doi.org/10.1016/j.obhdp.2009.03.003>
- Haidt, J., & Trevino, L. (2017). Make business ethics a cumulative science [Comment]. *Nature Human Behaviour*, 1, 1–2. <https://doi.org/10.1038/s41562-016-0027>
- Hodgkinson, G. P., & Starkey, K. (2011). Not simply returning to the same answer over and over again: Reframing relevance. *British Journal of Management*, 22(3), 355–369. <https://doi.org/10.1111/j.1467-8551.2011.00757.x>
- Houdek, P. (2019). Is behavioral ethics ready for giving business and policy advice? *Journal of Management Inquiry*, 28(1), 48–56. <https://doi.org/10.1177/1056492617712894>
- Huff, A., Tranfield, D., & van Aken, J. E. (2006). Management as a design science mindful of art and surprise. A conversation between Anne Huff, David Tranfield, and Joan Ernst van Aken. *Journal of Management Inquiry*, 15(4), 413–424.
- Kristal, A. S., Whillans, A. V., Bazerman, M. H., Gino, F., Shu, L. L., Mazar, N., & Ariely, D. (2020). Signing at the beginning versus at the end does not decrease dishonesty. *Proceedings of the National Academy of Sciences*, 117(13), 7103–7107. <https://doi.org/10.1073/pnas.1911695117>
- Kuenzi, M., Mayer, D. M., & Greenbaum, R. L. (2020). Creating an ethical organizational environment: The relationship between ethical leadership, ethical organizational climate, and unethical behavior. *Personnel Psychology*, 73(1), 43–71. <https://doi.org/10.1111/peps.12356>
- Loewenstein, G., Cain, D. M., & Sah, S. (2011). The limits of transparency: Pitfalls and potential of disclosing conflicts of interest. *American Economic Review*, 101(3), 423–428. <https://doi.org/10.1257/aer.101.3.423>
- 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>
- Merritt, A. C., Effron, D. A., & Monin, B. (2010). Moral self-licensing: When being good frees us to be bad. *Social and Personality Psychology Compass*, 4(5), 344–357. <https://doi.org/10.1111/j.1751-9004.2010.00263.x>
- Mitchell, M. S., Reynolds, S. J., & Treviño, L. K. (2020). The study of behavioral ethics within organizations: A special issue introduction. *Personnel Psychology*, 73(1), 5–17. <https://doi.org/10.1111/peps.12381>
- Moore, C., & Gino, F. (2013). Ethically adrift: How others pull our moral compass from true North, and how we can fix it. *Research in Organizational Behavior*, 33, 53–77. <https://doi.org/10.1016/j.riob.2013.08.001>
- Moynihan, D. (2018). A great schism approaching? Towards a micro and macro public administration. *Journal of Behavioral Public Administration*, 1, 1–8.
- Mumford, M. D., Connelly, S., Brown, R. P., Murphy, S. T., Hill, J. H., Antes, A. L., Waples, E. P., & Devenport, L. D. (2008). A sensemaking approach to ethics training for scientists: Preliminary evidence of training effectiveness. *Ethics & Behavior*, 18(4), 315–339. <https://doi.org/10.1080/10508420802487815>
- Nagin, D. S. (2013). Deterrence in the twenty-first century. *Crime and Justice*, 42(1), 199–263. <https://doi.org/10.1086/670398>
- Paine, L. S. (1994). Managing for organizational integrity. *Harvard Business Review*, 72(2), 106–117.
- Pawson, R. (2013). *The science of evaluation: A realist manifesto*. Sage.
- Pawson, R., & Tilley, N. (1997). *Realistic evaluation*. Sage.
- Pierce, J. R., & Aguinis, H. (2013). The too-much-of-a-good-thing effect in management. *Journal of Management*, 39(2), 313–338. <https://doi.org/10.1177/0149206311410060>
- Reynolds, S. J., & Bae, E. (2019). The dark side: Giving context and meaning to a growing genre of ethics-related research. In D. M. Wasieleski & J. Weber (Eds.), *Business ethics* (pp. 239–258). Emerald Publishing Limited.
- Romme, A. G. L. (2003). Making a difference: Organization as design. *Organization Science*, 14(5), 558–573. <https://doi.org/10.1287/orsc.14.5.558.16769>

- Rorie, M., & West, M. (2020). Can “focused deterrence” produce more effective ethics codes? An experimental study. *Journal of White Collar and Corporate Crime*. <https://doi.org/10.1177/2631309X20940664>
- Sezer, O., Gino, F., & Bazerman, M. H. (2015). Ethical blind spots: Explaining unintentional unethical behavior. *Current Opinion in Psychology*, 6, 77–81.
- Shu, L. L., Gino, F., & Bazerman, M. H. (2011). Dishonest deed, clear conscience: When cheating leads to moral disengagement and motivated forgetting. *Personality and Social Psychology Bulletin*, 37(3), 330–349.
- Shu, L. L., Mazar, N., Gino, F., Ariely, D., & Bazerman, M. H. (2012). Signing at the beginning makes ethics salient and decreases dishonest self-reports in comparison to signing at the end. *Proceedings of the National Academy of Sciences*, 109(38), 15197–15200.
- Simon, H. A. (1946). The proverbs of administration. *Public Administration Review*, 6(1), 53–67. <https://doi.org/10.2307/973030>
- Simon, H. A. (1957). *Models of man: Social and rational: Mathematical essays on rational human behavior in a social setting*. Wiley.
- Simon, H. A. (1969). *The sciences of the artificial* (3rd ed.). MIT Press.
- Smith, W. K., & Lewis, M. W. (2011). Toward a theory of paradox: A dynamic equilibrium model of organizing. *The Academy of Management Review*, 36(2), 381–403. <https://doi.org/10.5465/AMR.2011.59330958>
- Smith, W. K., Lewis, M. W., Jarzabkowski, P., & Langley, A. (2017). *The Oxford handbook of organizational paradox*. Oxford University Press.
- Smith-Crowe, K., & Zhang, T. (2016). On taking the theoretical substance of outcomes seriously: A meta-conversation. In D. Palmer, K. Smith-Crowe, & R. Greenwood (Eds.), *Organizational wrongdoing. Key perspectives and new directions* (pp. 17–46). Cambridge: Cambridge University Press.
- Tenbrunsel, A. E., & Messick, D. M. (1999). Sanctioning systems, decision frames, and cooperation. *Administrative Science Quarterly*, 44(4), 684–707. <https://doi.org/10.2307/2667052>
- Thiel, C. E., Bagdasarov, Z., Harkrider, L., Johnson, J. F., & Mumford, M. D. (2012). Leader ethical decision-making in organizations: Strategies for sensemaking. *Journal of Business Ethics*, 107(1), 49–64. <https://doi.org/10.1007/s10551-012-1299-1>
- Treviño, L. K., Butterfield, K. D., & McCabe, D. L. (1998). The ethical context in organizations: Influences on employee attitudes and behaviors. *Business Ethics Quarterly*, 8(3), 447–476. <https://doi.org/10.2307/3857431>
- Treviño, L. K., Den Nieuwenboer, N. A., & Kish-Gephart, J. J. (2014). (Un)ethical behavior in organizations. *Annual Review of Psychology*, 65, 635–660. <https://doi.org/10.1146/annurev-psych-113011-143745>
- Trevino, L. K., Weaver, G. R., & Reynolds, S. J. (2006). Behavioral ethics in organizations: A review. *Journal of Management*, 32(6), 951–990. <https://doi.org/10.1177/0149206306294258>
- van Aken, J. E. (2004). Management research based on the paradigm of the design sciences: The quest for field-tested and grounded technological rules. *Journal of Management Studies*, 41(2), 219–246. <https://doi.org/10.1111/j.1467-6486.2004.00430.x>
- van Aken, J. E. (2005). Management research as a design science: Articulating the research products of mode 2 knowledge production in management. *British Journal of Management*, 16(1), 19–36. <https://doi.org/10.1111/j.1467-8551.2005.00437.x>
- van Aken, J. E., & Romme, G. (2009). Reinventing the future: Adding design science to the repertoire of organization and management studies. *Organization Management Journal*, 6(1), 5–12. <https://doi.org/10.1057/omj.2009.1>
- Van Aken, J. E., & Romme, A. G. L. (2012). A design science approach to evidence-based management. In D. M. Rousseau (Ed.), *The Oxford handbook of evidence-based management* (pp. 43–57). Oxford University Press.
- Van Aken, J., Berends, H., & van der Bij, H. (2012). *Problem solving in organizations. A methodological handbook for business and management students*. Cambridge University Press.
- Van Aken, J., Chandrasekaran, A., & Halman, J. (2016). Conducting and publishing design science research: Inaugural essay of the design science department of the Journal of Operations Management. *Journal of Operations Management*, 47, 1–8. <https://doi.org/10.1016/j.jom.2016.06.004>
- Verschuere, B., Meijer, E. H., Jim, A., Hoogesteyn, K., Orthey, R., McCarthy, R. J., Skowronski, J. J., Acar, O. A., Aczel, B., & Bakos, B. E. (2018). Registered replication report on Mazar, Amir, and Ariely (2008). *Advances in Methods and Practices in Psychological Science*, 1(3), 299–317. <https://doi.org/10.1177/2515245918781032>
- Victor, B., & Cullen, J. B. (1988). The organizational bases of ethical work climates. *Administrative Science Quarterly*, 33(1), 101–125.
- Weaver, G. R., & Treviño, L. K. (1999). Compliance and values oriented ethics programs: Influences on employees’ attitudes and behavior. *Business Ethics Quarterly*, 9(2), 315–335. <https://doi.org/10.2307/3857477>
- Zhang, T., Gino, F., & Bazerman, M. H. (2014). Morality rebooted: Exploring simple fixes to our moral bugs. *Research in Organizational Behavior*, 34, 63–79. <https://doi.org/10.1016/j.riob.2014.10.002>
- Zhang, T., Gino, F., & Margolis, J. D. (2018). Does “could” lead to good? On the road to moral insight. *Academy of Management Journal*, 61(3), 857–895. <https://doi.org/10.5465/amj.2014.0839>

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.