

Marketers Project Their Personal Preferences onto Consumers Overcoming the Threat of Egocentric Decision Making

Herzog, Walter ; Hattula, Johannes; Dahl, Darren W.

Document Version

Final published version

Published in:

Journal of Marketing Research

DOI:

[10.1177/0022243721998378](https://doi.org/10.1177/0022243721998378)

Publication date:

2021

License

CC BY-NC

Citation for published version (APA):

Herzog, W., Hattula, J., & Dahl, D. W. (2021). Marketers Project Their Personal Preferences onto Consumers: Overcoming the Threat of Egocentric Decision Making. *Journal of Marketing Research*, 58(3), 456-475. <https://doi.org/10.1177/0022243721998378>

[Link to publication in CBS Research Portal](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us (research.lib@cbs.dk) providing details, and we will remove access to the work immediately and investigate your claim.

Download date: 16. Apr. 2024



Marketers Project Their Personal Preferences onto Consumers: Overcoming the Threat of Egocentric Decision Making

Journal of Marketing Research
2021, Vol. 58(3) 456-475
© The Author(s) 2021



Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/0022243721998378
journals.sagepub.com/home/mrj



Walter Herzog, Johannes D. Hattula, and Darren W. Dahl

Abstract

This research explores how marketers can avoid the so-called “false consensus effect”—the egocentric tendency to project personal preferences onto consumers. Two pilot studies show that most marketers have a surprisingly strong lay intuition about the existence of this inference bias, admit that they are frequently affected by it, and try to avoid it when predicting consumer preferences. Moreover, the pilot studies indicate that most marketers use a very natural and straightforward approach to avoid the false consensus effect in practice, that is, they simply try to “suppress” (i.e., ignore) their personal preferences when predicting consumer preferences. Ironically, four subsequent studies show that this frequently used tactic can backfire and increase marketers’ susceptibility to the false consensus effect. Specifically, the results suggest that these backfire effects are most likely to occur for marketers with a low level of preference certainty. In contrast, the results imply that preference suppression does not backfire but instead decreases the false consensus effect for marketers with a high level of preference certainty. Finally, the studies explore the mechanism behind these results and show how marketers can ultimately avoid the false consensus effect—regardless of their level of preference certainty and without risking backfire effects.

Keywords

attitude certainty, customer centricity, customer insights, debiasing, false consensus effect, managerial decision making, new product development, thought suppression

Online supplement: <https://doi.org/10.1177/0022243721998378>

A series of influential studies in marketing have established that marketing executives are affected by a multitude of undesirable behavioral tendencies in their decision making and have investigated strategies to avoid these tendencies (e.g., Blattberg and Hoch 1990; Mittal, Ross, and Tsiros 2002; Perkins and Rao 1990; Puto 1987; Qualls and Puto 1989; Ross 1991). However, none of the existing studies have examined how marketing managers can avoid the so-called “false consensus effect” (FCE)—that is, the tendency to project personal preferences onto target consumers (Hattula et al. 2015; Krueger and Clement 1994). Yet, this research question is highly relevant from a managerial perspective. Indeed, we contribute to the literature by showing that most marketing executives (1) have a surprisingly strong lay intuition about the existence of the FCE, (2) admit that they are frequently affected by the FCE, and (3) try to avoid the FCE when predicting the preferences of target consumers. Moreover, our research shows that marketing managers consider the ability to separate personal preferences from the preferences of target consumers (i.e., the ability to avoid the FCE) to be one of the most important skills in the marketing

profession. These findings suggest that the extant academic literature in marketing may have overlooked the managerial relevance of the FCE and, in particular, the managerial relevance of avoiding the FCE.

Furthermore, we find that most marketers use a very natural and straightforward approach to avoid the FCE when predicting consumer preferences; specifically, they simply try to ignore or “suppress” their personal preferences during prediction processes.¹ This finding is in line with the general

¹ Individuals frequently attempt to avoid biases by suppressing bias-inducing factors (e.g., Petty et al. 2007; Wilson and Brekke 1994). According to Wegner (2009), the term “suppression” specifically refers to the attempt to ignore the bias-inducing factor. The main idea behind suppression-based debiasing is

Walter Herzog is Professor of Marketing, WHU – Otto Beisheim School of Management, Germany (email: walter.herzog@whu.edu). Johannes D. Hattula is Assistant Professor of Marketing, Imperial College Business School, Imperial College London, UK (email: j.hattula@imperial.ac.uk). Darren W. Dahl is BC Innovate Professor, Sauder School of Business, University of British Columbia, Canada (email: darren.dahl@sauder.ubc.ca).

literature on debiasing techniques, which suggests that the suppression of biasing factors is one of the most intuitive and therefore common tactics to avoid biases (Petty et al. 2007; Strack and Mussweiler 2001; Wegner 2009). However, existing research on the FCE has not yet explored whether this frequently used approach is an effective way to avoid this inference bias.

Our contribution centers on addressing this gap in the existing literature by exploring whether preference suppression effectively reduces the FCE and, if so, under which conditions. We find that the effectiveness of preference suppression is inextricably linked to managers' level of preference certainty, that is, the extent to which their personal preferences are clear and held with conviction (Rucker et al. 2014). We show that efforts toward preference suppression reduce managers' susceptibility to the FCE for high levels of preference certainty; however, we also demonstrate that for lower levels of preference certainty, preference suppression attempts backfire and actually *increase* managers' susceptibility to the FCE. The latter finding is particularly important from a managerial point of view, because marketers frequently predict consumer reactions to novel stimuli (e.g., new products or new technologies; Bolton 2003)—a situation known to result in lower levels of preference certainty (Rucker et al. 2014). Thus, our findings suggest that the way marketing managers attempt to avoid the FCE in practice may frequently backfire and increase the FCE when predicting consumer preferences.

Importantly, our research also suggests that marketing managers' general susceptibility to the FCE can be greatly reduced if they are advised to suppress their personal preferences for high, but not for low, levels of preference certainty. This finding is important in light of the considerable robustness of the FCE demonstrated in previous studies. For example, Krueger and Clement (1994) concluded that the FCE is an extremely robust and potentially even “ineradicable . . . bias in social perception” (p. 596). However, our results suggest that this conclusion may have been premature, as the FCE can indeed be reduced via preference suppression when considering preference certainty as a moderating factor.

We explain the effects documented in this research by building on and extending Wegner's (1994, 2009) theory on the effects of mental suppression. In particular, we extend Wegner's work by introducing preference certainty as a moderator of the effects of preference suppression. In the following sections, we first present the results of two pilot studies to motivate our research. Drawing on the results of the pilot studies, we develop our theorizing and then test our predictions in a series of studies with experienced marketing executives.

Pilot Studies

Pilot Study 1

Pilot Study 1 was designed to explore three questions: (1) Do marketing managers have an intuitive knowledge of the tendency to project personal preferences onto consumers (i.e., the FCE)? (2) If so, do they try to avoid this tendency? (3) If they try to avoid this tendency, how do they do this? Each of these questions was explored in a separate part of the study (see Web Appendix A for materials and measures). We recruited 100 marketing executives ($M_{\text{age}} = 43.0$ years; 41.0% female) from the survey pools of two European business schools to take part in the study.

Part 1: Do marketing managers have an intuitive knowledge of the FCE? In the first part of the study, we assessed whether participants have an intuitive knowledge of the FCE and three other behavioral tendencies (confirmation bias, overconfidence, and unrealistic optimism; see Web Appendix A). We included three additional behavioral tendencies to conceal our primary interest in the FCE, thereby reducing potential demand effects.

Following Liu, McFerran, and Haws (2020), we measured participants' intuitive knowledge of the FCE by asking whether marketing managers' personal preferences have a positive, a negative, or no impact on their predictions of consumer preferences (the order of the response options was randomized). Overall, 82.0% of the participants (correctly) assumed that managers' personal preferences have a positive impact on their predictions of consumer preferences, 9.0% assumed a negative impact, and 9.0% assumed no impact. Moreover, participants' intuitions were better calibrated for the FCE than for the other three behavioral tendencies (correct intuitions for confirmation bias: 34.0%; overconfidence: 61.0%, unrealistic optimism: 53.0%). Overall, these results suggest that most marketers are intuitively aware of the FCE.

Part 2: Do marketing managers try to avoid the FCE? Participants with an intuitive knowledge of the FCE ($N = 82$) were asked to complete the second part of the study, while the remaining participants were directed to the end of the survey. In the second part of the study, participants were informed that we had randomly drawn one of the intuitions they provided in the first part of the study and would ask follow-up questions regarding this specific intuition. In fact, the relevant intuition was the FCE in all cases. We opted for this procedure to conceal our primary interest in the FCE, thereby reducing potential demand effects.²

Participants were first asked follow-up questions related to the prevalence of the FCE. Specifically, 92.7% of the participants witnessed that other managers in their company had been recently affected by the FCE, and 86.6% admitted that they

succinctly expressed by Petty et al. (2007, p. 273): “On some level, if it were possible, keeping biased thoughts from coming to mind at all would be the ideal way to avoid bias.” In our research context, marketers attempt to avoid the FCE by suppressing their personal preferences.

² Furthermore, this procedure enabled us to reduce the length of the management survey. Pretests indicated that the survey would have been too long if we had asked follow-up questions for all intuitions from the first part of the study.

themselves had been recently affected by the FCE. These results suggest that the FCE is a prevalent behavioral tendency among marketing managers. Moreover, additional follow-up questions revealed that most marketing managers consider the FCE a “bias” as defined by Wilson and Brekke (1994)—that is, an “unwanted” tendency from the perspective of the decision maker. Specifically, 79.3% of the participants indicated that marketing managers should avoid the FCE, and 79.3% stated that they attempt to avoid the FCE when predicting consumer preferences.

Part 3: How do marketing managers try to avoid the FCE? Participants who indicated that they attempt to avoid the FCE ($N = 65$) were asked to complete the third part of the study, while the remaining participants were directed to the end of the survey. In the third part of the study, participants were asked how they attempt to avoid the FCE. Specifically, extant research on debiasing tactics suggests that managers might be able to avoid inference biases such as the FCE by (1) attempting to suppress the bias-generating factor (here, the managers’ personal preferences; Petty et al. 2007), (2) relying on market research (Blattberg and Hoch 1990), (3) consulting other managers (i.e., group decision making; Larrick 2004), (4) considering the opposite (i.e., considering that consumers might have preferences opposite to their own; DeRosia and Elder 2019), (5) attempting to adjust in the direction opposite to the bias (Petty et al. 2007), and/or (6) avoiding “primacy” of the biasing factor (here, considering personal preferences only after other information has been explored; DeRosia and Elder 2019). Participants were provided with a list of these debiasing methods (in randomized order) and asked to indicate which methods they tend to use (multiple answers were possible). Moreover, participants were encouraged to enter any additional methods not found on the list.

The results reveal that most participants (75.4%) use a very natural and straightforward approach to avoid the FCE—that is, they simply try to ignore or “suppress” their personal preferences when making consumer predictions. This finding is consistent with the existing literature on debiasing tactics, which implies that decision makers frequently attempt to reduce the influence of a biasing factor (in our case, their personal preferences) by mentally suppressing this factor (e.g., Petty et al. 2007; Strack and Mussweiler 2001; Wegner 2009).

Furthermore, our findings suggest that some marketing managers also rely on “indirect” debiasing instruments,³ such as consulting other managers (23.1% of the participants) or using market research data (18.5%). We assume that these methods are less popular than preference suppression because they are costlier in terms of time and/or money. Moreover,

some marketing managers also try to avoid the FCE with strategies such as considering personal preferences only after other information has been explored (27.7%), considering the opposite (15.4%), and/or attempting to adjust in the direction opposite to the bias (12.3%).⁴

Finally, participants indicated that they consider it important to avoid the FCE ($M = 4.2$; 1 = “not important,” and 5 = “extremely important”) and expressed an interest in learning more about how to avoid the FCE. Specifically, 82.8% of the participants stated that they would like to read a management summary of an academic article on how to avoid the FCE, and 75.4% would even be willing to participate in an online workshop to learn more about this topic.

Overall, Pilot Study 1 suggests that many marketers (1) have a strong lay intuition about the existence of the FCE, (2) admit that they are affected by the FCE, and (3) consider it important to avoid the FCE by using a diverse portfolio of methods (with preference suppression being the most prevalent among these methods).

Pilot Study 2

We conducted a supplementary pilot study to provide further evidence that marketers consider it important to avoid the FCE (for study details and main results, see Table 1). For this study, we recruited 64 marketing managers ($M_{\text{age}} = 39.4$ years; 56.3% female) from the survey pools of two European business schools and asked them to complete Moorman’s (2020) scale on the nine most important abilities of marketing managers. Notably, we added one additional ability to this scale: the ability to separate personal preferences from the preferences of target consumers (i.e., the ability to avoid the FCE). Our findings revealed that the ability to separate personal from consumer preferences was rated as more important than the nine original abilities, and it therefore tops the list of the most important skills of marketing managers.

A natural question arising from the two pilot studies is whether marketing managers are indeed able to separate personal from consumer preferences and, thus, avoid the FCE. We explore this question in the remainder of the article, with a particular focus on the most prevalent debiasing tactic used by marketing managers (i.e., preference suppression).

Theoretical Background

Drawing on the results of the pilot studies, this research examines whether marketing managers can avoid the FCE by suppressing their personal preferences. In this section, we first show that prior research has not yet addressed this question. We then develop our hypotheses.

³ According to Kahn, Luce, and Nowlis (2006), “direct” debiasing refers to directly altering decision makers’ cognitive processes (preference suppression is thus a direct debiasing strategy). In contrast, “indirect” debiasing refers to altering decision makers’ environment (thus, market research data and group decision making are indirect debiasing measures).

⁴ Three participants provided additional debiasing methods not included in the list. Directly talking to target consumers was mentioned twice, and getting information on consumers’ preferences from various sources was mentioned once.

Table 1. Perceived Importance of Marketing Managers' Skills (Pilot Study 2).

Skills of Marketing Managers	Importance Rating ^a Mean (SD)	Top-Two-Box Score ^b
Ability to differentiate between their personal, subjective preferences and the preferences of target customers	1.88 (1.59)	81.25%
Ability to pivot as new priorities emerge	2.23 (1.63)	70.31%
Creativity and innovation skills	2.30 (1.74)	70.31%
Natural leadership abilities	2.45 (1.46)	56.25%
Emotional intelligence	2.48 (1.70)	57.81%
Navigating ambiguity	3.03 (1.57)	42.19%
MarTech platform experience	3.27 (1.87)	39.06%
Financial acumen	3.30 (1.57)	32.81%
Curiosity	3.33 (1.95)	42.19%
Data science background	4.59 (2.00)	9.38%

^a1 = "skill is very important," and 9 = "skill is not important."

^bPercentage of marketing managers who assigned an importance score of 1 or 2. Notes: N = 64 marketing managers participated in the study. Participants were asked to rate the importance of ten skills of marketing managers. The order of the skills was randomized. We adopted all (i.e., nine) items/skills from Moorman's (2020) scale of key abilities of marketing managers. Furthermore, we added the following item/skill to Moorman's list to measure the perceived importance of avoiding the FCE: "Ability to differentiate between their personal, subjective preferences and the preferences of target customers."

Related Literature

Literature on biases of marketing managers. An impressive amount of empirical work in marketing has established that marketing managers are affected by a multitude of biases in their decision making. For example, a series of influential studies in this area have demonstrated that marketing managers are susceptible to framing effects (i.e., positively or negatively framing an otherwise identical business situation has strong effects on managers' subsequent decisions; Mittal, Ross, and Tsiros 2002; Puto 1987; Qualls and Puto 1989; Ross 1991). Moreover, additional studies have shown that marketers are susceptible to the confirmation bias and tend to favor information that is in line with their preexisting assumptions (Bolton 2003; Lee, Acito, and Day 1987), show signs of unrealistic optimism when evaluating new product ideas (DeRosia and Elder 2019), display an overly strong commitment to decisions that turn out to be suboptimal (Biyalogorsky, Boulding, and Staelin 2006), and tend to weight market research information inappropriately (Blattberg and Hoch 1990; Perkins and Rao 1990; Roggeveen and Johar 2004).

Moreover, Hattula et al. (2015) have shown that marketers are also susceptible to the FCE and that the size of the FCE depends on marketers' level of empathy with consumers. Specifically, their studies indicate that instructing (vs. not instructing) marketing managers to empathize with consumers increases their susceptibility to the FCE. Furthermore, Study 4 in this article shows that instructing managers to be empathic *and* to suppress their personal preferences does not increase their susceptibility to the FCE. Unfortunately, this study did not include an experimental condition with preference suppression instructions only. Thus, it is not possible to isolate the specific effect of preference suppression on the FCE based on this study.

Overall, the existing literature in this area has not investigated whether marketers can avoid the FCE by suppressing their personal preferences. Similarly, we were not able to identify studies that have explored the effectiveness of other remedies for the FCE. In other words, this line of research has generally neglected the important question of how marketing managers can avoid the FCE (for an overview of key studies in this area, see Web Appendix B). In a next step, we extend our literature review to the general psychological literature on the FCE.

General literature on the FCE. In more than four decades of empirical research, a large number of studies have provided evidence for the FCE (for an overview of important studies, see Web Appendix C). Reviews of this literature frequently emphasize three aspects of the FCE that relate to the results of our pilot studies and our research question (e.g., Ostrom et al. 2017). First, the FCE is one of the most prevalent biases identified in social psychology (Ross and Ward 1996). People tend to project all kinds of opinions onto others, such as their personal opinions of products (David 2018), advertisements (Ross and Ward 1996), or music (Gilovich 1990). Moreover, intra-individual analyses indicate that most people are affected by the FCE when engaged in social inference processes (Krueger and Zeiger 1993). These findings are consistent with the results of our pilot studies, which suggest that most marketers are affected by the FCE.

Second, the FCE is driven by factors unrelated to the actual similarity between decision makers' preferences and the preferences of the target group, such as the momentary activation of decision makers' preferences (Dunning and Hayes 1996), decision makers' desire for social support of their preferences (Ostrom et al. 2017), or decision makers' perceived influence over the target group (Overbeck and Droutman 2013). Thus, the FCE is frequently considered a biasing factor that can increase prediction errors (e.g., Flynn and Wiltermuth 2010). These findings are in line with the results of our pilot studies, which indicate that most marketers consider the FCE an undesirable tendency that should be avoided.

Third, reviews of this research area frequently emphasize that the FCE is a fairly robust inference bias (e.g., Ostrom et al. 2017) and refer to a classic study by Krueger and Clement (1994) in this context. Krueger and Clement's findings suggest

that the FCE remains unchanged when decision makers are (1) educated about the bias, (2) provided with feedback on the accuracy of their predictions, or (3) provided with statistical information on the target group.

However, it is important to note that the studies presented by Krueger and Clement (1994) do not imply that preference suppression is an ineffective remedy for the FCE. First, our Pilot Study 1 shows that creating awareness of the FCE does not necessarily trigger preference suppression. Specifically, the study shows that (1) not all marketing managers who are aware of the FCE attempt to avoid it, and (2) not all marketing managers who attempt to avoid it engage in preference suppression. Second, providing accuracy feedback is known to increase participants' focus on both diagnostic *and* nondiagnostic sources of information (such as their personal preferences in social inference processes; Kunda 1990). Thus, providing accuracy feedback in the present research context may in fact reduce the likelihood of preference suppression. Third, Pilot Study 1 suggests that preference suppression attempts in practice occur independent of the use of market research on the target group; specifically, these two approaches for avoiding the FCE are uncorrelated ($r = -.004$, $t = -.03$, $p > .97$).

Overall, Krueger and Clement's (1994) investigation implies that the FCE is a particularly robust inference bias—but it did not explore whether preference suppression is an effective or ineffective remedy for the FCE. We were not able to identify other, more direct studies on the effectiveness of preference suppression in the general literature on the FCE.

Summary. Existing research in marketing and psychology has not yet investigated the effect of preference suppression on the FCE. Our research addresses this important gap in the literature (see our pilot studies). To do so, we draw on and extend the most established theory on the effectiveness of suppression-based debiasing: Wegner's (1994, 2009) ironic process theory.

Theoretical Framework

Overview. According to Wegner (1994, 2009), suppression refers to the attempt to avoid a specific bias by ignoring or "suppressing" the factor that induces the bias (see also Petty et al. 2007; Wilson and Brekke 1994). Consistent with this general definition, we conceptualize preference suppression as decision makers' attempt to avoid the FCE by ignoring their personal preferences. However, note that Wegner's model has not yet been applied to the analysis of the effects of preference suppression on the FCE. We therefore extend his model to the context of preference suppression.

Importantly, as discussed in detail in the next section, our adapted version of Wegner's model implies that the effects of preference suppression on the FCE are inextricably linked to managers' level of preference certainty—that is, the extent to which their personal preferences are clear and held with conviction (Rucker et al. 2014). Specifically, our model suggests that preference suppression attempts reduce managers' susceptibility to the FCE for high levels of preference certainty;

however, it also suggests that for low levels of preference certainty, preference suppression attempts backfire and *increase* managers' susceptibility to the FCE.

Our analysis of the effects of preference suppression for different levels of preference certainty is important for at least three reasons. First, due to the opposing effects of preference suppression for high (vs. low) preference certainty, an analysis omitting preference certainty as a moderator might (mistakenly) lead to the conclusion that preference suppression has no effect on the FCE. Second, marketing managers frequently predict consumer reactions to new stimuli that are not yet available in the market (e.g., product innovations; Bolton 2003)—a situation known to result in lower levels of preference certainty (Rucker et al. 2014). Thus, our model indicates that in practice, preference suppression may frequently backfire and increase the FCE when predicting consumer preferences. Third, preferences that are held with low certainty levels are often labeled "weak preferences" because they are easily affected by situational factors and can change quickly (Rucker et al. 2014). Therefore, it seems particularly unwise to project these types of preferences onto consumers. However, our model implies that managers' well-intentioned attempts to inhibit such preferences will backfire and inflate their impact on consumer predictions.

Preference suppression and the role of preference certainty. Wegner's (1994, 2009) model suggests that preference suppression attempts initiate two concurrent cognitive processes that work in concert to avoid the FCE. First, a so-called operating process is invoked whereby managers attempt to infer the target consumer's preferences (e.g., reactions to a new product feature) without referring to their personal preferences. At the same time, a monitoring process is started that tests whether the inferences during the operating process are indeed not self-referential—that is, it essentially compares managers' personal preferences to their predictions of consumer preferences. If self-reference is detected, the operating process is adjusted accordingly, thereby reducing decision makers' susceptibility to the FCE (Petty et al. 2007; Wilson and Brekke 1994).

A large number of empirical studies indicate that "under normal circumstances" (Wilson and Brekke 1994, p. 128), operating and monitoring processes interact efficiently, thereby leading to a successful suppression of bias-inducing factors (Abramowitz, Tolin, and Street 2001; Wegner 2009). However, Wegner's model also implies that suppression attempts will fail if decision makers are not able to properly execute the aforementioned monitoring process—that is, if their ability to compare personal and predicted preferences is limited. Such a situation is most likely to occur if the "search target" of the monitoring process—in our case, a manager's personal preference—is vague and not well-defined (Sternberg 1966; Wegner 1994; Wilson and Brekke 1994). In the context of preference suppression, decision makers' ability to successfully execute the monitoring process might therefore depend on their level of preference certainty (Rucker et al. 2014).

Specifically, a meaningful comparison between personal and predicted preferences requires that managers can precisely articulate their personal preferences regarding a target stimulus (e.g., a new product feature). If this condition is not met, managers' consumer predictions may—or may not—coincide with their personal preferences and it is, strictly speaking, impossible to arrive at a definitive conclusion regarding this comparison process. We therefore propose that managers are able to execute Wegner-type monitoring processes for high, but not for low, levels of preference certainty. Thus, preference suppression will be an effective remedy for the FCE for high, but not for low, levels of preference certainty.

Ironic consequences for low preference certainty. Further, Wegner's (1994, 2009) theory suggests that when a suppression attempt "goes wrong, it goes very wrong" (Wilson and Brekke 1994, p. 128), meaning that preference suppression may be not only ineffective but *counterproductive* for low levels of preference certainty. Thus, managers' well-intentioned efforts to suppress their personal preferences may backfire and *increase* the FCE in this case.

Specifically, our preceding discussion suggests that for low levels of preference certainty, monitoring processes pose a question that decision makers cannot answer (i.e., "Do my personal preferences affect my preference predictions?"). However, from the perspective of decision makers, this relationship may not be readily apparent, and they may therefore continue their futile attempts to infer potential effects of their personal preferences. This assumption is consistent with the existing literature on self-regulation, which shows that decision makers tend to invest a considerable amount of effort into tasks that have no objective solution (so-called "unsolvable tasks") before noticing the futility of their efforts (e.g., Baumeister et al. 1998). Similarly, extant research on suppression-based debiasing indicates that decision makers are surprisingly relentless in their efforts to detect biases, even though it can be difficult (if not impossible) to master this task (Duguid and Thomas-Hunt 2015; Wilson and Brekke 1994). We therefore expect that for a low level of preference certainty, preference suppression will trigger an effortful, but eventually futile, attempt to infer whether personal preferences affect consumer predictions.

According to Wegner (1994, 2009), such an attempt (i.e., an effortful but futile attempt to infer whether personal preferences affect consumer predictions) is likely to produce "ironic" (i.e., counterintentional) effects, and we thus refer to it as "ironic monitoring process." Specifically, decision makers' well-intentioned but futile effort to infer potential effects of their personal preferences automatically generates a strong focus of attention on their personal preferences (which contradicts decision makers' primary goal to ignore their personal preferences).⁵ In other words, the ironic monitoring process

is considered equivalent to a "prime" of decision makers' personal preferences in the existing literature on suppression attempts (Higgins 1996; Macrae et al. 1994; Wegner 2009). Moreover, it is well-documented that priming personal preferences inflates the effects of personal preferences on preference predictions (i.e., the FCE increases; Hattula et al. 2015; Higgins 1996; Krueger and Clement 1994). We therefore expect that for low levels of preference certainty, preference suppression attempts will trigger a Wegner-type ironic monitoring process, which in turn increases the FCE.

Hypotheses. Wegner's (1994, 2009) ironic process theory suggests that preference suppression is likely to reduce the FCE as long as it does not trigger an ironic monitoring process. If it does, preference suppression will not only be ineffective but also backfire and increase the FCE. We extend Wegner's model by assuming that the former situation is likely to occur for high levels of preference certainty, whereas the latter situation is expected to occur for lower levels of preference certainty. Thus, we anticipate the following relationships:

H₁: When preference certainty is high, marketing managers' attempts to suppress their personal preferences decreases the FCE.

H₂: When preference certainty is low, marketing managers' attempts to suppress their personal preferences increases the FCE.

H₃: The effect of preference suppression for low preference certainty (H₂) occurs because (a) preference suppression triggers an ironic monitoring process (i.e., an effortful but futile attempt to infer whether personal preferences affect consumer predictions), and (b) the ironic monitoring process increases the FCE. This mediation does not occur for high preference certainty, because in this case preference suppression does not trigger an ironic monitoring process.

Furthermore, H₃ implies a rather straightforward way to avoid the backfire effect for low preference certainty (H₂). Specifically, it suggests that this effect is driven by a Wegner-type ironic monitoring process. Thus, encouraging decision makers to abandon their suppression efforts at the first sign of ironic monitoring (i.e., when they find it difficult to tell whether they are affected by the FCE) should eliminate the effect postulated in H₂.

H₄: The backfire effect postulated in H₂ can be avoided by encouraging marketing managers to abandon their suppression efforts in cases where they find it difficult to tell whether they are affected by the FCE.

⁵ Wegner (1994) sometimes uses the terms "monitoring process" and "ironic monitoring process" interchangeably even though not all monitoring processes

produce ironic (i.e., counterintentional) results (e.g., Abramowitz, Tolin, and Street 2001; Wegner 2009). Strictly speaking, the attribute "ironic" only applies to monitoring processes that increase the decision maker's focus on the to-be-suppressed concept (in our case, the decision maker's personal preference) rather than assisting the decision maker to suppress it.

Overview of Studies

In the following sections, we present the results of four studies with 550 experienced marketing managers to test our hypotheses. Across these studies, we measure and manipulate both preference suppression and preference certainty to provide evidence for the robustness of our findings (Mittal, Ross, and Tsiros 2002; Wegener et al. 1995; Wenzlaff and Wegner 2000).

Study 1 is an observational study in which we provide initial evidence that preference suppression decreases (increases) the FCE for high (low) levels of preference certainty (H_1 , H_2). Moreover, the observational nature of Study 1 enables us to show that (1) numerous managers with a high level of preference certainty do not choose to suppress their personal preferences (even though this would reduce the FCE; see H_1), and (2) numerous managers with a low level of preference certainty choose to suppress their personal preferences, which increases the FCE (see H_2). In other words, marketing managers' "natural suppression behavior" is suboptimal. In Study 2, we revalidate the findings of Study 1 in a more controlled setting by manipulating preference suppression (Wenzlaff and Wegner 2000). Moreover, the results of Studies 1 and 2 suggest that for high (low) levels of preference certainty, preference suppression increases (decreases) the accuracy of managers' consumer predictions (Flynn and Wiltermuth 2010).

Study 3 replicates the proposed interaction pattern by manipulating both preference suppression and preference certainty. Importantly, Study 3 also shows that the backfire effect of preference suppression for low certainty levels is mediated by a Wegner-type ironic monitoring process (see H_3). Finally, Study 4 indicates that the intervention proposed in H_4 eliminates this backfire effect. Notably, Study 4 also shows that the proposed intervention makes managers completely immune to the FCE—regardless of their level of preference certainty. Table 2 summarizes the main results of all studies (including pilot and supplementary studies).

Study 1

Study 1 has two main purposes. First, it was designed to provide initial evidence that preference suppression decreases (increases) the FCE for high (low) levels of preference certainty (H_1 , H_2). Second, we used Study 1 to test whether managers are more likely to suppress their preferences for high (vs. low) certainty levels (which would be an effective strategy to reduce their susceptibility to the FCE if H_1 and H_2 are correct).

Method

Eighty-six marketing managers ($M_{\text{age}} = 48.1$ years; 14.6% female) were recruited from a European marketing association to take part in a case study on new product development. The general study setup was similar to the investigations conducted by Hattula et al. (2015); furthermore, to ensure that the materials and measures were realistic and comprehensible, we

consulted experienced executives and modified the case study according to their suggestions.

The study consisted of three parts. In the first part of the study, participants were shown a picture of a new product, a virtual reality headset, and they were asked to familiarize themselves with its value propositions and technical specifications, which were shown in a table. For example, they were informed that the virtual reality headset projects images directly onto the user's retina and could be used to watch movies or to play games. Next, we measured participants' personal evaluation of the product (1 = "don't like it at all," and 7 = "like it very much"; 1 = "find it very negative," and 7 = "find it very positive"; $\alpha = .926$). Following Holland, Verplanken, and Van Knippenberg (2003), we used two items to measure participants' preference certainty ("How certain are you of your evaluation of the Virtual Reality Headset?" and "How certain are you that you would not change your opinion of the Virtual Reality Headset if we showed you the same product information three weeks from now and asked you to evaluate the product again?"; 1 = "very uncertain," and 7 = "very certain"; $\alpha = .707$).

In the second part of the study, participants were informed that the company "HHTech" had developed the product, and detailed information about the company was presented. Furthermore, they were asked to assume the role of the company's segment manager responsible for the target group "university students." Then, participants were shown the picture of the new product and its technical specifications again. Finally, they were asked to predict how target consumers (i.e., university students) would evaluate the product (1 = "don't like it at all," and 7 = "like it very much"; 1 = "find it very negative," and 7 = "find it very positive"; $\alpha = .880$).

In the final part of the study, participants were asked to provide general information about themselves and their experiences during the case study. Most importantly, we asked them to indicate whether they attempted to suppress their personal preferences when predicting the target consumers' preferences ("Have you tried to suppress your own opinion of the new product when predicting target consumers' opinion of the Virtual Reality Headset?" and "Have you tried to suppress your own opinion of the new product when predicting the success of the Virtual Reality Headset in the target market?"; 1 = "not at all," and 7 = "very much"; $\alpha = .756$).

Results

We analyzed the data by regressing participants' consumer preference predictions on their personal preferences, their preference certainty, and their level of preference suppression. All predictor variables were mean-centered, and the regression model also included all possible linear interaction terms. Web Appendix D reports the full regression results.

Test of H_1 and H_2 . Three parameters of the regression model are of particular importance for our investigation: First, the analysis revealed a highly significant main effect of participants'

Table 2. Summary of Studies.

Overview of Pilot Studies						
Study Sample	Context	Certainty	Suppression	Hypothesis	Qualitative Finding	Test Statistic Effect Size ^a
<p><i>Main results of Pilot Study 1 (N = 100 marketing executives):</i> 82.0% of the participating marketing executives had an intuitive knowledge of the FCE. Of these participants, 92.7% witnessed that other managers in their company had been recently affected by the FCE, 86.6% admitted that they themselves had been recently affected by the FCE, and 79.3% indicated that marketing managers should avoid the FCE. Importantly, 79.3% stated that they attempt to avoid the FCE when predicting preferences of their target consumers. Of these participants, 75.4% use a very natural and straightforward approach to avoid the FCE—that is, they simply try to ignore or “suppress” their personal preferences when making consumer predictions (Petty et al. 2007; Strack and Mussweiler 2001; Wegner 2009).</p> <p><i>Main results of Pilot Study 2 (N = 64 marketing executives):</i> The participating marketing executives were asked to complete Moorman’s (2020) scale on the nine most important abilities of marketing managers. Notably, we added one additional ability to this scale: the ability to separate personal preferences from the preferences of target consumers (i.e., the ability to avoid the FCE). The ability to separate personal from consumer preferences was rated as more important than the nine original abilities and thus tops the list of the most important skills of marketing managers.</p>						
Overview of Main Studies						
1	86 marketing executives	Perception of a new product	Measured	H ₁	Preference suppression decreased the FCE for high preference certainty	b = -.121, t = -2.59, p < .05
				H ₂	Preference suppression increased the FCE for low preference certainty	b = .170, t = 2.93, p < .01
2	133 marketing executives	Perception of a new product	Measured	H ₁	Preference suppression decreased the FCE for high preference certainty	b = -.282, t = -2.68, p < .01
			Manipulated (control condition, preference suppression condition)	H ₂	Preference suppression increased the FCE for low preference certainty	b = .341, t = 3.04, p < .01
3	206 marketing executives	Perception of artificial intelligence	Manipulated (low certainty condition, high certainty condition)	H ₁	Preference suppression decreased the FCE for high preference certainty	b = -.451, t = -3.05, p < .01
			Manipulated (control condition, preference suppression condition)	H ₂	Preference suppression increased the FCE for low preference certainty	b = .430, t = 2.62, p < .01
				H ₃	Ironic monitoring mediated the effect of preference suppression on the FCE for low preference certainty	95% CI of indirect effect: [.014, .271]

(continued)

Table 2. (continued)

Overview of Main Studies								
Study Sample	Context	Certainty	Suppression	Hypothesis	Qualitative Finding	Test Statistic Effect Size ^a		
4	125 marketing executives	Perception of an advertising campaign	Measured	Manipulated (control condition, preference suppression condition, modified preference suppression condition)	H ₁ H ₂ H ₄	Preference suppression decreased the FCE for high preference certainty Preference suppression increased the FCE for low preference certainty (a) Modified preference suppression decreased the FCE for high preference certainty (b) Modified preference suppression did not affect the FCE for low preference certainty	b = -.487, t = -2.39, p < .05 b = .676, t = 2.09, p < .05 b = -.501, t = -2.86, p < .01 b = .166, t = .51, p > .61	-.22 .20 -.27 .05

Meta-Analysis (N = 550 Marketing Executives)

The average size of the effect of preference suppression on the FCE is -0.234 ($z = -5.29, p < .0001$) for high preference certainty and $.229$ ($z = 5.19, p < .0001$) for low preference certainty. These results are in line with H₁ and H₂. Furthermore, the analysis reveals that the effect sizes are not heterogeneous across studies (high preference certainty: $Q(3) = .32, p > .95$; low preference certainty: $Q(3) = 1.44, p > .69$). Both effect sizes are similar to the average effect sizes reported in marketing studies (i.e., $\pm .24$; Eisend 2015). For further details, see Web Appendix K.

Supplementary Analyses of Prediction Accuracy

For Studies 1 and 2, we ran a series of supplementary analyses and collected data on target consumers' actual preferences to test whether preference suppression also affects the accuracy of managers' preference predictions (see Web Appendix J). The results suggest that for low preference certainty, preference suppression increases the FCE, which in turn decreases managers' prediction accuracy; for high preference certainty, preference suppression decreases the FCE, which in turn increases prediction accuracy (see also Flynn and Wiltermuth 2010; Krueger and Clement 1994).

^aWe computed partial correlation coefficients as effect size measures (e.g., Han et al. 2019).

personal preferences on their predictions of consumer preferences ($b = .427$, $t = 5.88$, $p < .001$). That is, on average, participants displayed a strong FCE. Second, there was no significant interaction effect between participants' personal preferences and preference suppression on their preference predictions ($b = .025$, $t = .67$, $p > .50$). This indicates that, on average, preference suppression did not reduce the FCE. Third and most importantly, the latter finding is qualified by a strong and highly significant three-way interaction ($b = -.111$, $t = -3.89$, $p < .001$). This result implies that the effect of preference suppression on the FCE (i.e., the interaction effect between personal preferences and preference suppression on preference predictions) heavily depends on managers' preference certainty (H_1 and H_2).

We explored the nature of the identified three-way interaction by running a series of simple slope and simple interaction analyses, which are visually summarized in Figure 1. Specifically, the figure shows the simple effects of managers' personal preferences on their preference predictions (i.e., the FCE) as a function of both their level of preference suppression (at one standard deviation unit below and above the mean) and preference certainty (at one standard deviation unit below and above the mean). The right-hand side of the figure shows that for high preference certainty, the FCE drops from .605 ($t = 4.64$, $p < .001$) to .218 ($t = 2.35$, $p < .05$) as a manager's level of preference suppression increases. Consistent with this result, a simple interaction analysis reveals that preference suppression significantly reduces the FCE for a high level of preference certainty ($b = -.121$, $t = -2.59$, $p < .05$). This finding is in line with H_1 .

The left-hand side of Figure 1 shows that for low preference certainty, the FCE increases from .170 ($t = 1.33$, $p > .18$) to .716 ($t = 4.62$, $p < .001$) as a manager's level of preference suppression increases. In line with this pattern, a simple interaction analysis indicates that preference suppression significantly increases the FCE for a low level of preference certainty ($b = .170$, $t = 2.93$, $p < .01$). This result is in line with H_2 and implies that managers' well-intentioned efforts to suppress their personal preferences backfire and increase the FCE.

Moreover, it is noteworthy that managers with a low level of preference certainty are *not* susceptible to the FCE (.170, $p > .18$) unless they start suppressing their personal preferences (.716, $p < .001$; Figure 1). In other words, we find that managers' attempts to avoid the FCE actually *create* this inference bias for low certainty levels. We are not aware of previous empirical evidence for such "bias creation effects" in the existing literature on debiasing.

The effect of preference certainty on the FCE conditional on preference suppression. For a low level of preference suppression, the FCE increases with preference certainty ($b = .166$, $t = 2.54$, $p < .05$). This finding is in line with the extant literature on preference certainty. Specifically, prior research has demonstrated that high levels of preference certainty are associated with the perception that one's preferences are objectively "correct" and should thus be shared by other individuals

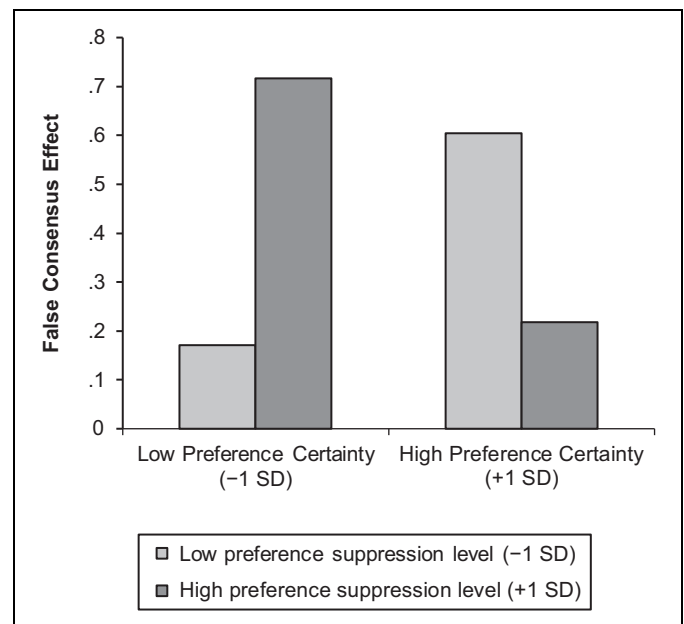


Figure 1. Simple effects analysis (Study 1).

Notes: The figure shows the false consensus effect (i.e., the effect of managers' personal preferences on predicted consumer preferences) as a function of the managers' preference certainty and their level of preference suppression.

(Marks and Miller 1985). Consequently, preference certainty can increase the FCE (Ross and Ward 1996). However, the effects postulated in H_1 and H_2 are strong enough to reverse this pattern of results once managers start suppressing their preferences (Figure 1). In this case, preference certainty decreases the FCE ($b = -.190$, $t = -2.78$, $p < .01$).

Managers' natural suppression behavior. Given the empirical evidence for H_1 and H_2 , decision makers should only suppress their personal preferences for high (but not for low) certainty levels. If the managers in our sample had intuitively adopted this "optimal suppression strategy," we would have observed a positive correlation between preference certainty and the level of preference suppression. However, our analysis reveals a correlation coefficient close to zero ($r = .015$, $t = .13$, $p > .89$). This result implies that (1) numerous managers with a high level of preference certainty do not attempt to suppress their preferences (even though this would reduce the FCE), and (2) numerous managers with a low level of preference certainty engage in counterproductive suppression attempts, thereby increasing the FCE.

Discussion

Two main findings emerge from Study 1. First, the results provide support for our main hypotheses, H_1 and H_2 . Specifically, our analysis suggests that preference suppression is an effective remedy for the FCE for high levels of preference certainty. In contrast, preference suppression backfires and increases the FCE for low levels of preference certainty. Second, our analysis indicates that numerous managers with a high

level of preference certainty do not suppress their preferences (even though this would reduce the FCE; H_1). At the same time, we find that numerous managers with a low level of preference certainty engage in counterproductive suppression attempts, thereby increasing the FCE (H_2).

Moreover, we ran a series of supplementary analyses and collected data on target consumers' actual preferences to test whether preference suppression also affects the accuracy of managers' preference predictions (see Web Appendix J). Overall, the results suggest that for low certainty levels, preference suppression increases the FCE, which in turn decreases managers' prediction accuracy; for high certainty levels, preference suppression decreases the FCE, which in turn increases prediction accuracy. These findings imply a strong link between the FCE and prediction (in)accuracy (see also Flynn and Wiltermuth 2010; Krueger and Clement 1994).

Study 2

Although the results of Study 1 are consistent with our hypotheses, it is important to note that preference suppression was not randomized in this study. We chose this study design to analyze managers' natural suppression behavior (Wenzlaff and Wegner 2000). In Study 2, we manipulated preference suppression to increase the causal validity of our results.

Method

One hundred thirty-three marketing managers ($M_{\text{age}} = 45.6$ years; 14.4% female) were recruited from a large alumni pool of a European management school to take part in the study. The design of Study 2 was similar to Study 1, and we used the same materials and measures. As in Study 1, participants were first asked to familiarize themselves with the target product (i.e., the virtual reality headset). Subsequently, we measured their personal preferences for the product ($\alpha = .927$), as well as their preference certainty ($\alpha = .843$).

In contrast to Study 1, however, we then manipulated participants' level of preference suppression when estimating the product preferences of the target consumers (i.e., university students). Similar to the procedures used by Duguid and Thomas-Hunt (2015) and Macrae et al. (1994), we randomly assigned participants to either a preference suppression condition or a control condition. In the preference suppression condition, participants read the following instructions:

A number of scientific studies indicate that decision makers tend to project their personal attitudes onto target customers. In other words, managers tend to equate their own evaluation of a product with the target customer's evaluation of the product. This effect can result in severe decision errors as well as biased perceptions of the market success of new products. Therefore, please try to suppress your personal evaluation of the Virtual Reality Headset in order to estimate customers' perceptions of the product in an objective way.

In the control condition, participants did not receive such instructions. Subsequently, all participants estimated the target consumers' evaluations of the product ($\alpha = .763$).

Results

We created a dummy variable indicating experimental group membership (1 = preference suppression condition, 0 = control condition) and used the same regression method as in Study 1 to test our hypotheses. Specifically, we regressed participants' consumer preference predictions on their personal preferences (mean-centered), their preference certainty (mean-centered), the dummy variable indicating group membership (i.e., the preference suppression manipulation), and the interactions among these variables (see Web Appendix E).

Test of H_1 and H_2 . Consistent with Study 1, we find a strong and highly significant three-way interaction ($b = -.257$, $t = -4.40$, $p < .001$). Thus, the impact of preference suppression on the FCE depends on a manager's level of preference certainty (H_1 and H_2). In a next analysis step, we investigated the three-way interaction more closely by running a series of simple slope and simple interaction analyses. The results of these analyses are visualized in Figure 2.

For a high level of preference certainty, the FCE is .354 ($t = 4.13$, $p < .001$) in the control condition, and it decreases to .072 ($t = 1.19$, $p > .23$) in the preference suppression condition. A simple interaction analysis confirms that preference suppression significantly reduces the FCE if managers have a high level of preference certainty ($b = -.282$, $t = -2.68$, $p < .01$). This finding is consistent with the results of Study 1 and provides further evidence for H_1 .




Figure 2. Simple effects analysis (Study 2).

Notes: The figure shows the false consensus effect (i.e., the effect of managers' personal preferences on predicted consumer preferences) as a function of the managers' preference certainty and experimental condition (control condition and preference suppression condition).

For a low level of preference certainty, the FCE is $-.005$ ($t = -.07, p > .94$) in the control condition, and it increases to $.336$ ($t = 4.17, p < .001$) in the preference suppression condition. A simple interaction analysis indicates that preference suppression significantly increases the FCE if managers have a low level of preference certainty ($b = .341, t = 3.04, p < .01$). This finding is in line with H_2 . Furthermore, we find that for low levels of preference certainty, the FCE does not exist unless managers attempt to suppress their preferences. In other words, the FCE is “generated” by the attempt to avoid it.

The effect of preference certainty on the FCE conditional on preference suppression condition. We again find that the postulated effects of preference suppression are strong enough to reverse the positive effect of preference certainty on the FCE (Marks and Miller 1985): the effect of preference certainty on the FCE is positive in the control condition ($b = .148, t = 3.52, p < .001$) and negative in the preference suppression condition ($b = -.109, t = -2.69, p < .01$).⁶

Discussion

In Study 2, we used an experimental manipulation of preference suppression to increase the causal validity of our findings. Overall, the results of Study 2 are in line with those of Study 1. Specifically, preference suppression decreases the FCE for high levels of preference certainty (H_1); in contrast, preference suppression backfires for low levels of preference certainty (H_2).

Parallel to Study 1, we ran a series of supplementary analyses to test whether preference suppression also affects the accuracy of managers' preference predictions (see Web Appendix J). We again find that preference suppression increases (decreases) the FCE for low (high) preference certainty, which in turn decreases (increases) prediction accuracy.

Study 3

The primary aim of Study 3 is to explore the mechanism behind the effects of preference suppression on the FCE (H_3). For low levels of preference certainty, we expect that (1) suppression attempts trigger a Wegner-type ironic monitoring process (i.e., an effortful but futile attempt to infer whether personal preferences affect consumer predictions) and (2) the ironic monitoring process increases the FCE (Higgins 1996; Wegner 1994). For high levels of preference certainty, we expect that suppression attempts do not trigger an ironic monitoring process, and thus, they should be effective according to Wegner's (1994, 2009) framework.

Furthermore, Study 3 was designed to provide additional evidence for the robustness of our results. First, and in contrast

to the previous studies, we manipulated preference certainty in Study 3 (Wegener et al. 1995). Second, Studies 1 and 2 mainly focused on the effect of managers' personal preferences on their *predictions* of consumer preferences. However, the two studies provided no evidence that managers *act* in accordance with their self-referential predictions. In Study 3, we therefore investigated whether managers actually try to influence important decisions of a company in a way that is consistent with their personal preferences. Specifically, Study 3 examines how managers' personal opinion about a new technology—artificial intelligence (AI)—affects their decision to invest in AI-based products.⁷ Third, in Studies 1 and 2 we used the same items for measuring personal and predicted preferences. This may have caused inflated estimates of the FCE. By focusing on the effect of managers' personal preferences on their decision making, we also hope to overcome this issue.

Method

Two hundred six marketing managers ($M_{\text{age}} = 41.0$ years; 39.8% female) were recruited from a large alumni pool of a European management school and from a European marketing association to participate in the study. Similar to the previous studies, we ensured that the materials and measures were realistic and comprehensible by asking experienced executives for their feedback and by modifying the study design according to their suggestions.

In the first part of the study, participants were asked to watch a short video (3:02 minutes) on AI-based products. After participants watched the video, we measured their personal opinion about AI-based products using the same items as in Studies 1 and 2 (1 = “don't like it at all,” and 7 = “like it very much”; 1 = “find it very negative,” and 7 = “find it very positive”; $\alpha = .858$). Next, participants were asked a series of general questions related to AI (e.g., “IBM supercomputer Watson helps physicians to diagnose illnesses”; right/wrong).

We then manipulated participants' preference certainty by means of a classic false feedback intervention (Fazio and Zanna 1978; Wegener et al. 1995; for details, see Web Appendix F). Specifically, we stated that we had recorded a wide range of objective data on how participants responded to the previous AI-related attitude and general questions (e.g., the time they needed to respond to each question). Furthermore, we allegedly used this information to compute a preference certainty score that ranged from 1 (“low certainty”) to 7 (“high certainty”). Participants in the high (low) certainty condition were then told that their certainty score was 5.9 (2.1). Next, we stated that we were also interested in participants' subjective assessment of their preference certainty (Fazio and Zanna

⁶ We find evidence for this effect reversal in all remaining studies, that is, negative effects of preference certainty on the FCE in the preference suppression conditions of Study 3 ($b = -.448, t = -2.58, p < .05$) and Study 4 ($b = -.249, t = -2.43, p < .05$). We do not discuss these results in Studies 3 and 4 further to focus on our central hypothesis tests.

⁷ The study materials are based on a communication with a marketing manager. The manager was recently asked to predict the acceptance of AI-based products in the general consumer population. The marketing manager also encouraged us to study predictions about the general consumer population (vs. specific consumer segments; Studies 1 and 2) as such “high level” predictions are frequently used for strategic decisions in his company.

1978; Marks and Miller 1985). We used the same items as in the previous studies to measure this construct ($\alpha = .744$). The results indicate that our manipulation was successful ($M_{\text{low certainty}} = 4.13$, $M_{\text{high certainty}} = 4.81$, $t = 4.20$, $p < .001$).⁸

In the second part of the study, participants were asked to complete a management case study on consumers' reactions to AI-based products. However, before participants began working on the case study, we manipulated preference suppression in a manner parallel to Study 2. In the preference suppression condition, participants were warned against the FCE and asked to suppress their personal preferences; participants assigned to the control condition received no specific instructions before working on the case study.

All participants were then asked to assume the role of a marketing expert at a holding company that has traditionally invested in a highly diverse portfolio of firms and start-ups to minimize financial risk. Participants were told that recent advances in AI have triggered questions as to whether the company should pursue a riskier strategy and predominantly invest in AI-oriented ventures, starting with two specific investment targets: (1) an AI-based app that helps users learn new languages by directly communicating with them and correcting their pronunciation and grammar and (2) an AI-based app that gradually learns how the user responds to different types of humor and provides humorous content in line with the user's specific preferences. Participants were then informed that this strategic move has garnered mixed reactions; in particular, some executive board members were unsure whether the time is right to heavily invest in AI-based products, as consumers are increasingly concerned, even fearful, of this new technology (Pegasystems 2018). Thus, participants—in their role as marketing experts—were asked to evaluate the suggested strategic move from a consumer-centric perspective and report back to the board with a recommendation (for a similar task, see Hattula et al. [2015, Study 1]). We used a four-item scale to measure this variable (“How much would you try to convince the executive board members that consumers like AI-based products?,” “How much would you try to address critical comments against the AI-based investment strategy?,” “How much would you criticize a market research study indicating that consumers might not like AI-based products?,” and “How much would you try to show that consumers like AI-based products even if some board members are against the AI-based investment strategy?”; 1 = “not at all,” and 7 = “very much”; $\alpha = .753$).

We then measured whether participants experienced an ironic monitoring process (Wegner 1994, 2009). Specifically, we indicated that sometimes decision makers think a lot about whether their personal preferences affect their consumer predictions, but ultimately, they do not find an answer to this

question. We then assessed whether they focused on this question without coming to a conclusion (“I have spent a lot of time thinking about this question and still was not able to answer it,” “I have put a lot of thought into this question and still was not able to answer it,” “I have spent a lot of energy thinking about this question and still was not able to answer it,” “My thoughts circled around this question to no avail,” and “I ‘got stuck’ thinking about this question”; 1 = “not at all,” and 7 = “very much so”; $\alpha = .754$).

Results

We created a dummy variable for the preference certainty manipulation (1 = high preference certainty condition, and 0 = low preference certainty condition) and a dummy variable for the preference suppression manipulation (1 = preference suppression condition, and 0 = control condition). We then used the same regression method as in Studies 1 and 2 to test our hypotheses. Specifically, we regressed participants' intentions to advocate the AI-based investment strategy on their personal opinion about AI (mean-centered), the two dummy variables, and all linear interactions among these variables. Web Appendix G shows the full regression results.

Test of H_1 and H_2 . As in Studies 1 and 2, the analysis reveals a three-way interaction ($b = -.881$, $t = -3.99$, $p < .001$). We ran a series of simple slope and simple interaction analyses to characterize the three-way interaction. Figure 3 provides a visual summary of these analyses. Specifically, it shows the extent of egocentric decision making (i.e., the effect of managers' personal opinion about AI on their intentions to support

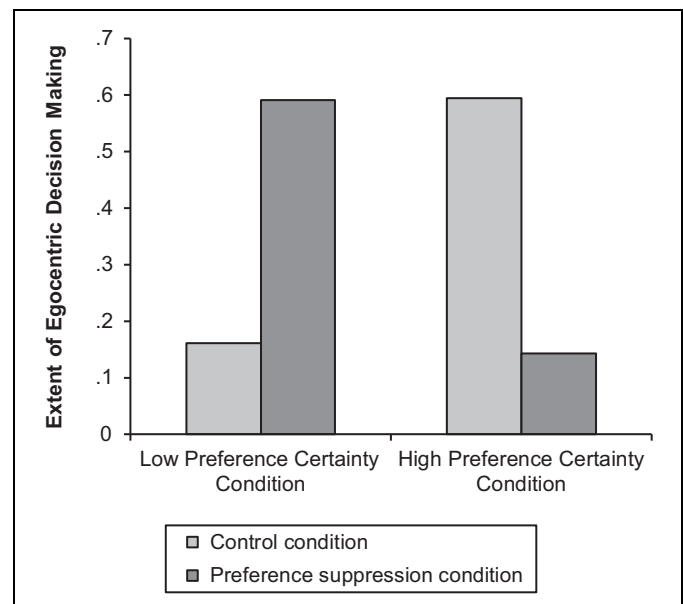


Figure 3. Simple effects analysis (Study 3).

Notes: The figure shows the extent of egocentric decision making (i.e., the effect of managers' personal opinion of AI on their intentions to support the AI-based strategy of the company) as a function of preference certainty (manipulated) and preference suppression (manipulated).

⁸ At the end of the study, participants completed an open-ended suspicion probe on the purpose of the study. None of the participants guessed the true purpose of the study or expressed suspicion regarding the false feedback intervention.

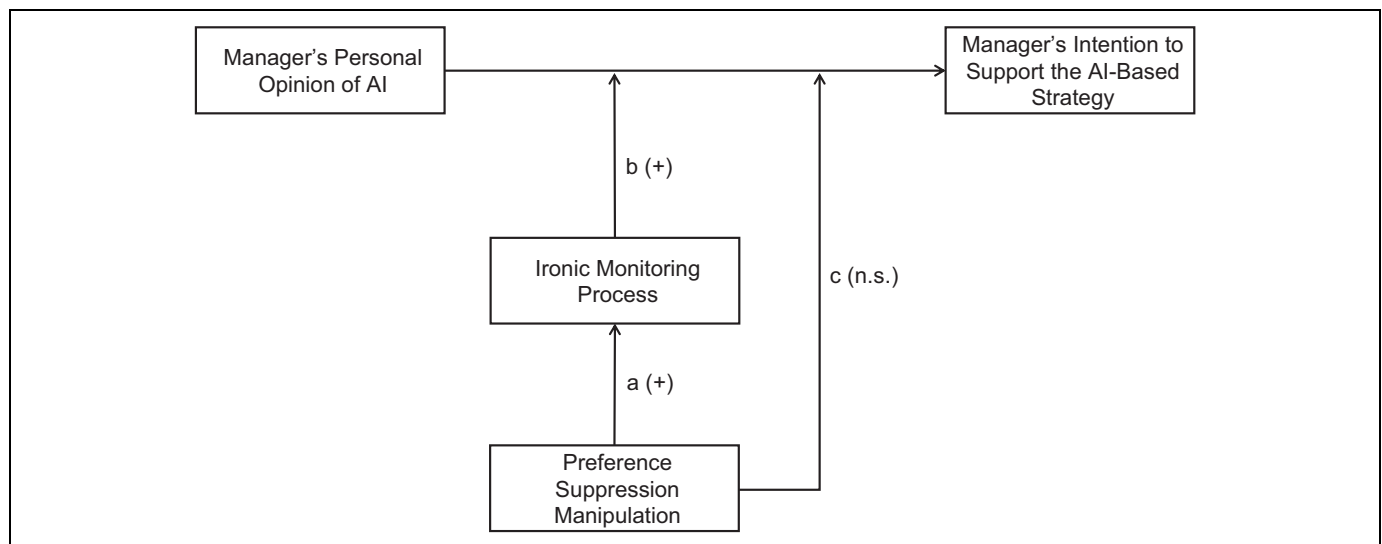


Figure 4. Mediation analysis within the low preference certainty condition (Study 3).

Notes: The mediation model shows the indirect effect of the preference suppression manipulation on egocentric decision making (i.e., the effect of managers' personal opinion of AI on their intentions to support the AI-based strategy of the company) via the postulated "ironic monitoring process" (Wegner 1994, 2009) in the low preference certainty condition. The model setup for estimating the indirect effect and the corresponding mediation figure is adopted from Hattula et al. (2015, Study 3).

an AI-based investment strategy) as a function of the preference certainty manipulation and the preference suppression manipulation.

Simple slope analysis indicates that in the high preference certainty condition, preference suppression decreases the extent of egocentric decision making from a high level ($b = .594, t = 5.93, p < .001$) to a low level ($b = .143, t = 1.32, p > .18$). Simple interaction analysis shows that this difference is highly significant ($b = -.451, t = -3.05, p < .01$). This result is in line with H_1 .

In the low preference certainty condition, preference suppression increases the extent of egocentric decision making from a low level ($b = .161, t = 1.76, p < .10$) to a high level ($b = .591, t = 4.35, p < .001$). Simple interaction analysis shows that this difference is significant ($b = .430, t = 2.62, p < .01$). This result provides further evidence for H_2 . Moreover, and consistent with the previous studies, we find that in the low preference certainty condition, there is only modest evidence for egocentric decision making unless managers attempt to suppress their personal preferences; again, egocentric processing is "generated" by the attempt to avoid it.

Mechanism analysis (H_3). Our theorizing implies that in the low preference certainty condition, preference suppression triggers an ironic monitoring process, which in turn increases egocentric decision making. The corresponding mediation model is depicted in Figure 4 (based on Hattula et al. [2015]). To test this mediation hypothesis, we ran a series of regressions in the low preference certainty condition (Hattula et al. 2015). As predicted, preference suppression increased the extent of ironic monitoring ("a path" in the model: $.560, t = 2.87, p < .01$). Furthermore, ironic monitoring increased egocentric decision making; that is, it increased the effect of managers' personal

opinion about AI on their advocacy of the AI-based investment strategy ("b path" in the model: $.200, t = 2.70, p < .01$). A bootstrap analysis based on 5,000 resamples shows that the indirect effect $a \times b$ is significant (95% confidence interval: $[.014, .271]$). Moreover, the direct effect of preference suppression on egocentric decision making is not significant ("c path" in the model: $.240, t = 1.31, p > .19$). Thus, ironic monitoring fully mediates the backfire effect of preference suppression on egocentric decision making.

Next, we ran the same mediation analysis in the high preference certainty condition. Consistent with our theorizing, preference suppression did *not* trigger an ironic monitoring process in the high preference certainty condition ($a = -.235, t = -1.19, p > .23$) and thus, the indirect effect was also not significant (95% confidence interval: $[-.019, .095]$). Due to the absence of an ironic monitoring process, Wegner's (2009) model suggests that preference suppression is effective in this condition—a prediction in line with our data (Figure 3).

Overall, our mediation analyses are in line with H_3 and provide support for Wegner's (1994, 2009) model: suppression attempts are effective unless they trigger an ironic monitoring process. If they do, suppression attempts are not only ineffective but counterproductive. We extend Wegner's model by showing that the former situation is likely to occur for high levels of preference certainty, whereas the latter situation occurs for low levels of preference certainty.

Discussion

Study 3 replicates the predicted effects of preference suppression and extends the findings of the previous studies in two important ways. First, it shows that our results are robust to three changes in the experimental and measurement procedures

(manipulation of preference certainty, measurement of ego-centric decision making, and AI as a different study context). Second, the study suggests that a Wegner-type ironic monitoring process triggers the backfire effect of preference suppression for low preference certainty (H_3).

The latter finding also implies a simple way to eliminate the backfire effect of suppression for low preference certainty. Specifically, advising managers to abandon their suppression efforts at the first sign of ironic monitoring should eliminate this effect (H_4). At the same time, this advice should have *no consequences* for high levels of preference certainty, because in this case, suppression does *not* trigger an ironic monitoring process (see our mediation results for high preference certainty). We test these predictions in Study 4.

Study 4

In Study 4, we test the effectiveness of the aforementioned remedy for the backfire effect of suppression attempts (H_4). Further, Study 4 tests our predictions in another decision-making context (i.e., communication management).

Method

One hundred twenty-five marketing managers ($M_{\text{age}} = 45.9$ years; 19.2% female) were recruited from a large alumni pool of a European management school and from a European marketing association to participate in the study. Similar to the previous studies, we ensured that the materials and measures were realistic and comprehensible by asking experienced executives for their feedback and by modifying the study design according to their suggestions.

In the first part of the study, participants were asked to watch a short video⁹ of a popular TV show, *The Big Bang Theory*, and to read a description of the show. We then measured their personal evaluation of the TV show using the same items as in the previous studies (1 = “don’t like it at all,” and 7 = “like it very much”; 1 = “find it very negative,” and 7 = “find it very positive”; $\alpha = .900$). Furthermore, we measured preference certainty using a four-item scale ($\alpha = .909$).¹⁰

In the second part of the study, participants were asked to complete a management case study related to *The Big Bang Theory*. Specifically, we informed the participants that they

were going to estimate whether the target consumers of a particular company would like this TV show and its characters. However, before participants started to work on the case study, they were randomly assigned to one of three experimental conditions. In the preference suppression condition, participants read the same instruction as in Studies 2 and 3; specifically, they were warned against the FCE and asked to suppress their personal preferences. Participants assigned to the control condition received no specific instruction. To test H_4 , we also created a modified preference suppression condition, in which participants read the same instruction as in the preference suppression condition—however, in addition, they were also recommended to abandon their suppression efforts at the first sign of ironic monitoring (i.e., when they find it difficult to tell whether they are affected by the FCE; for details, see Web Appendix H).

All participants were then asked to assume the role of a marketing manager at the company “H&H Fitness.” Furthermore, they were provided with some background information on the firm. In particular, they were informed that H&H Fitness pursues a low-price strategy and targets price-conscious consumers. Moreover, participants learned that in a recent meeting, the executive board decided to launch a large-scale advertisement campaign. Specifically, a partnership with the stars of *The Big Bang Theory* was suggested. However, there were mixed reactions to this idea in the meeting; in particular, some board members were unsure whether the firm’s target consumers would identify with the show and its characters. Participants—in their role as marketers representing the “voice of the customer” in the company—were asked to determine whether such a partnership would be desirable from a customer-oriented point of view. Thus, and consistent with Study 3, the dependent variable in this study was participants’ intentions to support the idea of a partnership with the stars of *The Big Bang Theory* and to defend it against counterarguments. We adapted the scale from Study 3 to measure this variable (“How much would you try to convince the executive board of the cooperation?,” “How much would you try to address critical comments against the cooperation?,” “How much would you criticize a market research study suggesting that the target customers do not like the TV show?,” and “How much would you try to push the cooperation through even if it is more costly than expected?”; 1 = “not at all,” and 7 = “very much”; $\alpha = .808$).

Results

We created two dummy variables to indicate experimental group membership. The effect of the first dummy variable (1 for participants in the preference suppression condition, 0 otherwise) represents the difference between the preference suppression condition and the control condition; the effect of the second dummy variable (1 for participants in the modified preference suppression condition, 0 otherwise) represents the difference between the modified preference suppression condition and the control condition. We then used the same

⁹ Participants were randomly assigned to watch one of two videos (2:54 vs. 2:43 minutes). Our results did not vary across the two groups; thus, this variable is not discussed further.

¹⁰ Following Tormala, Clarkson, and Petty (2006, Study 4), we added two additional items to the certainty scale from Studies 1 and 2 to increase the reliability of the scale (“How certain are you of your evaluation of the TV show?” [1 = “very uncertain,” and 7 = “very certain”], “How certain are you that you will not change your opinion of the TV show if we asked you to evaluate it again in three weeks from now?” [1 = “very uncertain,” and 7 = “very certain”], “How convinced are you that your evaluation of the TV show is correct?” [1 = “not convinced at all,” and 7 = “very convinced”], and “How much confidence do you have in your evaluation of the TV show?” [1 = “no confidence at all,” and 7 = “very high confidence”]).

regression method as in the previous studies to test our hypotheses. Specifically, we regressed participants' intentions to advocate the advertising campaign with the stars of *The Big Bang Theory* on their personal preference for the TV show (mean-centered), their preference certainty (mean-centered), the two dummy variables indicating experimental group membership, and all possible linear interaction terms. Web Appendix I shows the full results of the regression analysis.

As in the previous studies, the regression analysis reveals a strong three-way interaction effect between managers' personal preferences, their preference certainty, and preference suppression (i.e., the first dummy variable) on their advocacy of the advertising campaign ($b = -.435$, $t = -2.94$, $p < .01$). As expected, the three-way interaction involving modified preference suppression (i.e., the second dummy variable) was weaker and only marginally significant ($b = -.249$, $t = -1.71$, $p < .10$). In a next step, we ran a series of simple slope and simple interaction analyses to characterize these three-way interactions (for a visual summary, see Figure 5). Figure 5 shows the extent of participants' egocentric decision making (i.e., the effect of managers' personal opinion about *The Big Bang Theory* on their intentions to support an advertising campaign with the stars of the show) as a function of both managers' preference certainty and the three experimental conditions.

Results for high preference certainty. Simple slope analysis reveals that for high preference certainty, the extent of egocentric decision making is high in the control condition ($b = .344$, $t = 2.64$, $p < .01$) and close to zero in the preference suppression condition ($b = -.143$, $t = -.91$, $p > .36$). Simple interaction analysis shows that preference suppression significantly reduces egocentric decision making for high preference certainty ($b = -.487$, $t = -2.39$, $p < .05$). This finding is consistent with H_1 and all previous studies.

Importantly, our analysis also implies that managers in the modified suppression condition did not engage in egocentric decision making ($b = -.157$, $t = -1.34$, $p > .18$). Furthermore, simple interaction analysis shows that relative to the control condition, the modified suppression strategy significantly reduced participants' egocentric decision making ($b = -.501$, $t = -2.86$, $p < .01$). Overall, the modified suppression strategy was as effective as "conventional" preference suppression in terms of reducing egocentric decision making.

Results for low preference certainty. For low preference certainty, the extent of egocentric decision making is close to zero in the control condition ($b = -.152$, $t = -.63$, $p > .53$) and high in the preference suppression condition ($b = .524$, $t = 2.42$, $p < .05$). Simple interaction analysis shows that preference suppression significantly increases egocentric decision making for low preference certainty ($b = .676$, $t = 2.09$, $p < .05$). This result is in line with H_2 . Furthermore, and consistent with our previous studies, we find that for low preference certainty there is no evidence for egocentric decision making unless managers attempt to suppress their personal preferences;

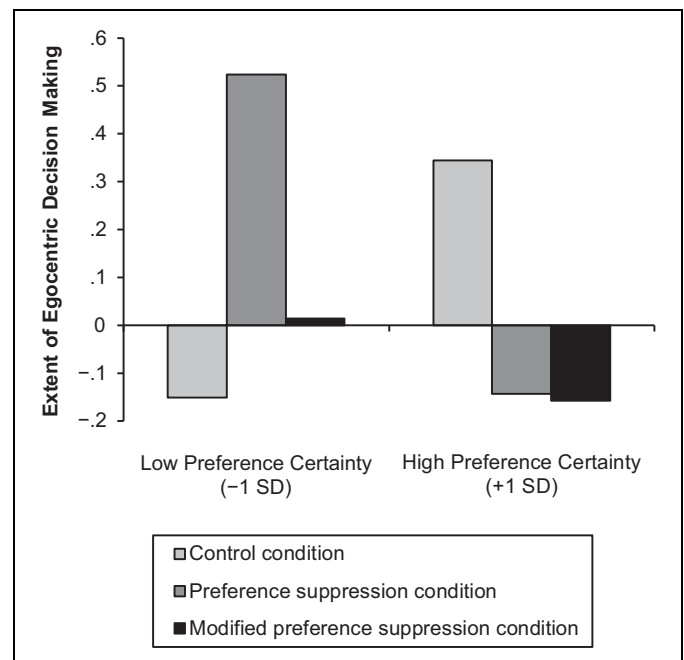


Figure 5. Simple effects analysis (Study 4).

Notes: The figure shows the extent of egocentric decision making (i.e., the effect of managers' personal opinion of *The Big Bang Theory* on their intentions to support an advertising campaign with the stars of the show) as a function of preference certainty and experimental condition (control condition, preference suppression condition, and modified preference suppression condition).

again, egocentric processing is "generated" by the attempt to avoid it.

Importantly, our analysis also shows that in the modified suppression condition, managers did *not* engage in egocentric decision making ($b = .014$, $t = .07$, $p > .94$). Furthermore, simple interaction analysis indicates that relative to the control condition, the modified suppression strategy did *not* increase managers' egocentric decision making ($b = .166$, $t = .51$, $p > .61$). Overall, it seems that our modification of "conventional" preference suppression eliminated the backfire effect for low preference certainty. This finding is consistent with H_4 .

Discussion

Study 4 provides further evidence for the robustness of the postulated effects in a different decision-making context (communication management). Moreover, the study reveals a simple and easy-to-implement intervention to avoid the backfire effect of suppression attempts for low preference certainty. Specifically, instructing managers to abandon their suppression efforts at the first sign of ironic monitoring eliminates the backfire effect (H_4). This result also provides experimental evidence for the postulated process behind the backfire effect (i.e., ironic monitoring). Finally, it is noteworthy that managers receiving this intervention were completely immune to egocentric processing—regardless of their level of preference certainty (Figure 5, modified preference suppression condition).

General Discussion

The findings of the present investigation contribute to a broad range of research streams in marketing and psychology. Most importantly, our research contributes to the existing literature on managerial biases in marketing (Blattberg and Hoch 1990; Hattula et al. 2015; Mittal, Ross, and Tsiros 2002; Perkins and Rao 1990; Puto 1987; Qualls and Puto 1989; Ross 1991) by exploring the role of the FCE from a managerial perspective. Specifically, our results reveal that most marketing executives (1) have a surprisingly strong lay intuition about the existence of the FCE, (2) admit that they are frequently affected by the FCE, and (3) attempt to avoid the FCE when predicting consumer preferences. Further, marketing managers consider the ability to separate their personal preferences from the preferences of target consumers (i.e., the ability to avoid the FCE) to be one of the most important skills in the marketing profession.

A natural question arising from these observations is whether and, if so, how marketers can avoid the FCE. We find that, in practice, marketers typically attempt to avoid the FCE by trying to ignore or “suppress” their personal preferences when predicting consumer preferences (Pilot Study 1). Yet existing research on the FCE has not explored whether this simple but frequently used tactic is an effective remedy for the FCE. We address this gap in the existing literature and explore whether preference suppression effectively reduces the FCE. To do so, we build on Wegner’s (1994, 2009) ironic process theory and show that the effect of preference suppression on the FCE is moderated by managers’ level of preference certainty (Rucker et al. 2014). Specifically, our studies suggest that suppression attempts reduce managers’ susceptibility to the FCE for high levels of preference certainty. However, for low levels of preference certainty, suppression attempts backfire and increase managers’ susceptibility to the FCE.

The latter finding is particularly important from a managerial perspective, because marketers frequently predict consumer reactions to novel stimuli (e.g., product innovations, new technologies; Bolton 2003; Hattula et al. 2015)—a situation known to result in lower levels of preference certainty (Rucker et al. 2014). Thus, our results suggest that the way marketing managers attempt to avoid the FCE in practice (see Pilot Study 1) may frequently backfire.

In addition, it is noteworthy that “weak preferences” (i.e., preferences that are held with low certainty levels) are easily affected by situational factors and can change quickly (Rucker et al. 2014). Therefore, it seems particularly undesirable that managers’ well-intentioned attempts to inhibit such “volatile” preferences backfire and increase their impact on consumer predictions.

Moreover, our studies provide support for the postulated process behind the backfire effects. For low preference certainty, suppression attempts trigger an ironic monitoring process (i.e., an effortful but futile attempt to infer whether personal preferences affect consumer predictions), which in turn increases the FCE (Study 3). This result provides direct support for Wegner’s theory and the notion that ironic

monitoring processes cause backfire effects (Dalton and Huang 2014; Wegner 1994). Consistent with our process hypothesis, we find that instructing managers to abandon their suppression efforts at the first sign of ironic monitoring (i.e., when they find it difficult to tell whether they are affected by the FCE) eliminates the backfire effect for low certainty levels (Study 4). Wegner (2009) notes that effective remedies for backfire effects of suppression attempts “remain largely unexplored” (p. 50) and calls for further research on this topic. Thus, the results of Study 4 provide a valuable contribution to this research area.

Implications for Marketing Managers

In recent years, many organizations have implemented learning concepts with the aim of reducing biases among decision makers (e.g., management training and educational programs; Duguid and Thomas-Hunt 2015). Our results have direct implications for the content of such initiatives. First, Pilot Study 1 suggests that many, but not all, marketers hold correct lay theories regarding the impact of personal preferences on preference predictions. Thus, management training should be used to adjust managers’ lay theories and create general awareness for the FCE. These training sessions could consist of multiple elements, such as (1) illustrative cases on the impact of the FCE on consumer predictions (e.g., Gershoff and Johnson 2003), (2) management summaries of scientific studies on the FCE (e.g., Lee, Acito, and Day 1987), and/or (3) quantitative assessments of each participant’s susceptibility to the FCE to show that most managers are affected by this inference bias (e.g., based on the methods by Krueger and Zeiger 1993). To stay aware of all biases that are typically discussed in such training sessions, marketers should also be encouraged to create a “checklist” of the most important biases—including the FCE—that they can consult before making important decisions (Beshears and Gino 2015). Such checklists are based on short descriptions of biases and examples of how they affect decisions. Checklists of this kind are a powerful tool to increase managers’ attention to biases, and their effectiveness has been demonstrated in many decision contexts (Gawande 2009).

Second, we find that many, but not all, managers who are aware of the FCE are also motivated to avoid this behavioral tendency (Pilot Study 1). Thus, it seems essential that training and education programs emphasize that the FCE is not a desirable tendency. For example, such programs should clarify that (1) decision makers tend to project their personal preferences independent of their actual similarity to the target group (Krueger and Clement 1994), (2) the FCE increases with factors that are unrelated to actual similarity (e.g., perceived influence over the target group; Overbeck and Drouman 2013), and (3) the FCE frequently leads to prediction errors (e.g., Flynn and Wiltermuth 2010). However, interventions of this kind can easily be perceived as “lecturing” and may therefore cause reactance (Wilson and Brekke 1994). A subtler approach may be found in setting the norm that it is common practice to avoid the FCE (Duguid and Thomas-Hunt 2015). For example, the results of

our pilot studies indicate that many marketers attempt to avoid the FCE (a descriptive norm) and also think that marketers should try to avoid the FCE (an injunctive norm). Thus, merely sharing these results with participants in training sessions may help them “internalize” the norm that it is good practice to avoid the FCE (Duguid and Thomas-Hunt 2015; Larrick 2004).

Third, Pilot Study 1 shows that if properly motivated, managers will most likely resort to preference suppression as their preferred strategy to avoid the FCE. Moreover, Study 1 shows that they will apply this strategy independent of their level of preference certainty. In other words, marketers are not aware of the fact that preference suppression is only effective for high levels of preference certainty and backfires for low preference certainty (Studies 1–4). Thus, training and education programs need to emphasize this aspect and encourage marketers to use this strategy for high levels of preference certainty only. Decision makers are quite capable of reporting their preference certainty in a variety of decision contexts (Petty et al. 2007), and it should thus be straightforward for them to implement our recommendation. Our studies clearly demonstrate that marketers adhering to this advice are less susceptible to the FCE (Studies 1–4) and their predictions of consumer preferences are more accurate. For example, in Study 1, managers’ prediction errors decreased by 60% when they decided to suppress (vs. not to suppress) their preferences for high certainty levels and by 54% when they decided not to suppress (vs. suppress) their preferences for low certainty levels (Web Appendix J, Figure W2).

Another alternative for managing the FCE can be found in the results of Study 4. In this study, we encouraged marketers to avoid the FCE but also recommended that they abandon their efforts in cases where they find it difficult to tell whether they are affected by the FCE (H_4). This simple intervention made marketing managers entirely immune to the FCE (Figure 5, modified preference suppression condition). While this method belies a more active approach to training and manager empowerment, it might be an attractive strategy for organizations that simply want to mitigate the issues that can potentially surround the FCE in a parsimonious way.

Future Research Opportunities

Our pilot studies suggest that the ability to differentiate between personal preferences and those of consumers (i.e., the ability to avoid the FCE) is one of the most important abilities in the marketing profession. The current research effort was the first in-depth exploration of this key ability, and it is our hope that it stimulates further research on this topic. For example, future research could examine the effectiveness of other strategies to avoid the FCE (e.g., clearly articulating personal preferences before predicting consumer preferences; Steffel, Williams, and Pogacar 2016). Moreover, future work could explore additional mechanisms that might underlie the effects of preference suppression on the FCE, as many psychological phenomena are multiply determined. Further research efforts could also be directed toward exploring whether our findings

generalize to other decision contexts (e.g., negotiations; Barasz, Kim, and John 2016).

Moreover, our studies suggest that managers with a low level of preference certainty are hardly susceptible to the FCE unless they attempt to avoid this inference bias—in other words, managers’ debiasing efforts seem to “generate” the FCE in this case. We are not aware of prior empirical evidence for such extreme cases of “counterintentional errors” (Wegner 1994, p. 34) in the general literature on ironic effects (e.g., Wegner 2009) or in the literature on managerial biases and debiasing methods (e.g., DeRosia and Elder 2019). Further research could explore whether similar effects can also be observed for other biases and debiasing methods.

Finally, our investigation was the first to measure marketers’ lay intuitions regarding biases, as well as the strategies they use to counter biases (Pilot Study 1). We believe that this approach may prove useful as a starting point for future research on debiasing strategies.

Acknowledgment

The authors would like to thank Christopher Beisecker, Ravi Dhar, Thomas Mussweiler, Zakary Tormala, Leaf Van Boven, Tillmann Wagner, Luk Warlop, Nathan Yang as well as seminar participants at BI Norwegian School of Management, McGill University, and University of Colorado Boulder for their valuable comments. Furthermore, the authors are grateful to the *JMR* review team for their extremely helpful and constructive advice.

Associate Editor

Randolph Bucklin

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

References

- Abramowitz, Jonathan S., David F. Tolin, and Gordon P. Street (2001), “Paradoxical Effects of Thought Suppression: A Meta-Analysis,” *Clinical Psychology Review*, 21 (5), 683–703.
- Barasz, Kate, Tami Kim, and Leslie John (2016), “The Role of (Dis)similarity in (Mis)predicting Others’ Preferences,” *Journal of Marketing Research*, 53 (4), 597–607.
- Baumeister, Roy F., Ellen Bratslavsky, Mark Muraven, and Dianne M. Tice (1998), “Ego Depletion: Is the Active Self a Limited Resource?” *Journal of Personality and Social Psychology*, 74 (5), 1252–65.
- Beshears, John and Francesca Gino (2015), “Leaders as Decision Architects,” *Harvard Business Review*, 93 (5), 51–62.
- Biyalogorsky, Eyal, William Boulding, and Richard Staelin (2006), “Stuck in the Past: Why Managers Persist with New Product Failures,” *Journal of Marketing*, 70 (2), 108–21.

- Blattberg, Robert C. and Stephen J. Hoch (1990), "Database Models and Managerial Intuition: 50% Model + 50% Manager," *Management Science*, 36 (8), 887–99.
- Bolton, Lisa E. (2003), "Stickier Priors: The Effects of Nonanalytic Versus Analytic Thinking in New Product Forecasting," *Journal of Marketing Research*, 40 (1), 65–79.
- Dalton, Amy N. and Li Huang (2014), "Motivated Forgetting in Response to Social Identity Threat," *Journal of Consumer Research*, 40 (6), 1017–38.
- David, Meredith E. (2018), "I Love the Product but Will You? The Role of Interpersonal Attachment Styles in Social Projection," *Psychology & Marketing*, 35 (3), 197–209.
- DeRosia, Eric D. and Ryan S. Elder (2019), "Harmful Effects of Mental Imagery and Customer Orientation During New Product Screening," *Journal of Marketing Research*, 56 (4), 637–51.
- Duguid, Michelle M. and Melissa C. Thomas-Hunt (2015), "Condoning Stereotyping? How Awareness of Stereotyping Prevalence Impacts Expression of Stereotypes," *Journal of Applied Psychology*, 100 (2), 343–59.
- Dunning, David and Andrew F. Hayes (1996), "Evidence for Egocentric Comparison in Social Judgment," *Journal of Personality and Social Psychology*, 71 (2), 213–29.
- Eisend, Martin (2015), "Have We Progressed Marketing Knowledge? A Meta-Analysis of Effect Sizes in Marketing Research," *Journal of Marketing*, 79 (3), 23–40.
- Fazio, Russell H. and Mark P. Zanna (1978), "On the Predictive Validity of Attitudes: The Roles of Direct Experience and Confidence," *Journal of Personality*, 46 (2), 228–43.
- Flynn, Francis J. and Scott S. Wiltermuth (2010), "Who's with Me? False Consensus, Brokerage, and Ethical Decision Making in Organizations," *Academy of Management Journal*, 53 (5), 1074–89.
- Gawande, Atul (2009), *The Checklist Manifesto*. New York: Metropolitan Books.
- Gershoff, Andrew and Eric Johnson (2003), "Avoid the Trap of Thinking Everyone Is Just Like You," *Financial Times* (August 29), 7.
- Gilovich, Thomas (1990), "Differential Construal and the False Consensus Effect," *Journal of Personality and Social Psychology*, 59 (4), 623–34.
- Han, Kyuhong, Jihye Jung, Vikas Mittal, Jinyong Daniel Zyung, and Hajo Adam (2019), "Political Identity and Financial Risk Taking: Insights from Social Dominance Orientation," *Journal of Marketing Research*, 56 (4), 581–601.
- Hattula, Johannes D., Walter Herzog, Darren W. Dahl, and Sven Reinecke (2015), "Managerial Empathy Facilitates Egocentric Predictions of Consumer Preferences," *Journal of Marketing Research*, 52 (2), 235–52.
- Higgins, E. Tory (1996), "Knowledge Activation: Accessibility, Applicability, and Salience," in *Social Psychology: Handbook of Basic Principles*, E. Tory Higgins and Arie W. Kruglanski, eds. New York: Guilford, 133–68.
- Holland, Rob W., Bas Verplanken, and Ad van Knippenberg (2003), "From Repetition to Conviction: Attitude Accessibility as a Determinant of Attitude Certainty," *Journal of Experimental Social Psychology*, 39 (6), 594–601.
- Kahn, Barbara E., Mary F. Luce, and Stephen M. Nowlis (2006), "Debiasing Insights from Process Tests," *Journal of Consumer Research*, 33 (1), 131–38.
- Krueger, Joachim and Joanna S. Zeiger (1993), "Social Categorization and the Truly False Consensus Effect," *Journal of Personality and Social Psychology*, 65 (4), 670–80.
- Krueger, Joachim and Russell W. Clement (1994), "The Truly False Consensus Effect: An Ineradicable and Egocentric Bias in Social Perception," *Journal of Personality and Social Psychology*, 67 (4), 596–610.
- Kunda, Ziva (1990), "The Case for Motivated Reasoning," *Psychological Bulletin*, 108, 480–98.
- Larrick, Richard P. (2004), "Debiasing," in *Blackwell Handbook of Judgment and Decision Making*, Derek J. Koehler and Nigel Harvey, eds. Malden, MA: Blackwell, 316–38.
- Lee, Hanjoon, Frank Acito, and Ralph L. Day (1987), "Evaluation and Use of Marketing Research by Decision Makers," *Journal of Marketing Research*, 24 (2), 187–96.
- Liu, Peggy J., Brent McFerran, and Kelly L. Haws (2020), "Mindful Matching: Ordinal Versus Nominal Attributes," *Journal of Marketing Research*, 57 (1), 134–55.
- Macrae, C. Neil, Galen V. Bodenhausen, Alan B. Milne, and Jolanda Jetten (1994), "Out of Mind but Back in Sight: Stereotypes on the Rebound," *Journal of Personality and Social Psychology*, 67 (5), 808–17.
- Marks, Gary and Norman Miller (1985), "The Effect of Certainty on Consensus Judgments," *Personality and Social Psychology Bulletin*, 11 (2), 165–77.
- Mittal, Vikas, William T. Ross Jr., and Michael Tsiros (2002), "The Role of Issue Valence and Issue Capability in Determining Effort Investment," *Journal of Marketing Research*, 39 (4), 455–68.
- Moorman, Christine (2020), "The CMO Survey: Topline Report, Special Edition," The CMO Survey (June), https://cmosurvey.org/wp-content/uploads/2020/06/The_CMO_Survey-Topline_Report-June-2020.pdf.
- Oostrom, Janneke K., Nils C. Köbis, Richard Ronay, and Myckel Cremers (2017), "False Consensus in Situational Judgment Tests," *Journal of Research in Personality*, 71, 33–45.
- Overbeck, Jennifer R. and Vitaliya Drouman (2013), "One for All: Social Power Increases Self-Anchoring of Traits, Attitudes, and Emotions," *Psychological Science*, 24 (8), 1466–76.
- Pegasystems (2018), "What Consumers Really Think About AI: A Global Study," (accessed May 18, 2019), <https://www.pegasystems.com/ai-survey>.
- Perkins, W. Steven and Ram C. Rao (1990), "The Role of Experience in Information Use and Decision Making by Marketing Managers," *Journal of Marketing Research*, 27 (1), 1–10.
- Petty, Richard E., Pablo Briñol, Zakary L. Tormala, and Duane T. Wegener (2007), "The Role of Metacognition in Social Judgment," in *Social Psychology: Handbook of Basic Principles*, Arie W. Kruglanski and E. Tory Higgins, eds. New York: Guilford Press, 254–84.
- Puto, Christopher P. (1987), "The Framing of Buying Decisions," *Journal of Consumer Research*, 14 (3), 301–15.
- Qualls, William J. and Christopher P. Puto (1989), "Organizational Climate and Decision Framing: An Integrated Approach to

- Analyzing Industrial Buying Decisions,” *Journal of Marketing Research*, 26 (2), 179–92.
- Roggeveen, Anne L. and Gita Venkataramani Johar (2004), “Integration of Discrepant Sales Forecasts: The Influence of Plausibility Inferences Based on an Evoked Range,” *Journal of Marketing Research*, 41 (1), 19–30.
- Ross, Lee and Andrew Ward (1996), “Naive Realism in Everyday Life: Implications for Social Conflict and Misunderstanding,” in *Values and Knowledge*, Edward S. Reed, Elliot Turiel, and Terrance Brown, eds. Mahwah, NJ: Lawrence Erlbaum Associates, 103–36.
- Ross, William T. (1991), “Performance Against Quota and the Call Selection Decision,” *Journal of Marketing Research*, 28 (3), 296–306.
- Rucker, Derek D., Zakary L. Tormala, Richard E. Petty, and Pablo Briñol (2014), “Consumer Conviction and Commitment: An Appraisal-Based Framework for Attitude Certainty,” *Journal of Consumer Psychology*, 24 (1), 119–36.
- Steffel, Mary, Elanor F. Williams, and Ruth Pogacar (2016), “Ethically Deployed Defaults: Transparency and Consumer Protection Through Disclosure and Preference Articulation,” *Journal of Marketing Research*, 53 (5), 865–80.
- Sternberg, Saul (1966), “High-Speed Scanning in Human Memory,” *Science*, 153 (3736), 652–54.
- Strack, Fritz and Thomas Mussweiler (2001), “Resisting Influence: Judgmental Correction and its Goals,” in *Social Influence Processes: Direct and Indirect Influences*, Joseph P. Forgas and Kipling D. Williams, eds. Philadelphia: Psychology Press, 199–212.
- Tormala, Zakary L., Joshua J. Clarkson, and Richard E. Petty (2006), “Resisting Persuasion by the Skin of One’s Teeth: The Hidden Success of Resisted Persuasive Messages,” *Journal of Personality and Social Psychology*, 91 (3), 423–35.
- Wegener, Duane T., John Downing, Jon A. Krosnick, and Richard E. Petty (1995), “Measures and Manipulations of Strength-Related Properties of Attitudes,” in *Attitude Strength*, Richard E. Petty and Jon A. Krosnick, eds. New York: Lawrence Erlbaum Associates, 455–87.
- Wegner, Daniel M. (1994), “Ironic Processes of Mental Control,” *Psychological Review*, 101 (1), 34–52.
- Wegner, Daniel M. (2009), “How to Think, Say, or Do Precisely the Worst Thing for Any Occasion,” *Science*, 325 (5936), 48–50.
- Wenzlaff, Richard M. and Daniel M. Wegner (2000), “Thought Suppression,” *Annual Review of Psychology*, 51 (1), 59–91.
- Wilson, Timothy D. and Nancy Brekke (1994), “Mental Contamination and Mental Correction: Unwanted Influences on Judgments and Evaluations,” *Psychological Bulletin*, 116 (1), 117–42.