

Virtue and Vice Attributions in the Business Context: An Experimental Investigation

Brian Robinson · Paul Stey · Mark Alfano

Received: 19 March 2012 / Accepted: 6 March 2013
© Springer Science+Business Media Dordrecht 2013

Abstract Recent findings in experimental philosophy have revealed that people attribute intentionality, belief, desire, knowledge, and blame asymmetrically to side-effects depending on whether the agent who produces the side-effect violates or adheres to a norm. Although the original (and still common) test for this effect involved a chairman helping or harming the environment, hardly any of these findings have been applied to business ethics. We review what little exploration of the implications for business ethics has been done. Then, we present new experimental results that expand the attribution asymmetry to virtue and vice. We also examine whether it matters to people that an effect was produced as a primary or side-effect, as well as how consumer habits might be affected by this phenomenon. These results lead to the conclusion that it appears to be in a businessperson's self-interest to be virtuous.

Keywords Attributions · Knobe effect · Side-effect · Side-effect effect · Vice · Virtue

Introduction

In the popular media, business leaders are more often portrayed as vicious masterminds rather than as benevolent giants. The financial crisis of 2007–2008 generated a spate of hardnosed investigative reportage on the excess and greed of the boom years (Sorkin's *Too Big to Fail*, Lewis's *The Big Short* and *Liar's Poker*, and McLean and Nocera's *All the Devils are Here*, to name a few examples). Corporations were the supervillains in the four of the last five James Bond movies, while the greatly hyped film adaptation of Ayn Rand's *Atlas Shrugged* was a box office flop. The Occupy Wall Street movement dominated the news cycle for much of the second half of 2011, branding the mantra "We are the 99 %" deep in the zeitgeist.

It would seem, then, that businesses and businesspeople have an uphill climb ahead of them if they want to be regarded as virtuous, or at least as not vicious. One might be tempted to respond by bemoaning this sorry state of affairs or sermonizing about the proper scope of business activity. We suggest an alternative path. Based on recent findings in the growing field of experimental philosophy, which attempts to probe, systematize, and explain ordinary people's intuitions about philosophical problems and concepts, we make some suggestions about what it might take to be considered virtuous in a business context.

We center our discussion on the *side-effect effect* (also known as the Knobe effect, after its discoverer), which crops up when people make attributions of psychological states to agents. In a nutshell, the effect is this: people are more willing to attribute a variety of attitudes (intention, belief, knowledge, desire, advocacy, etc.) to someone who brings about a norm-violating side-effect than to someone who brings about a norm-conforming side-effect. For instance, people are more inclined to say that the chairman of a company intentionally

B. Robinson (✉)
Grand Valley State University, Allendale, MI, USA
e-mail: robibri@gvsu.edu

P. Stey
University of Notre Dame, Notre Dame, IN, USA
e-mail: pstey@nd.edu

M. Alfano
Princeton University Center for Human Values, Princeton,
NJ, USA
e-mail: mark.alfano@gmail.com

harms the environment in his single-minded pursuit of profit than that he intentionally helps it, even if the help is a side-effect of profit seeking as well.

Because the side-effect effect came as such a surprise to philosophers of mind, most of the relevant subsequent research has been an attempt to explain it, rather than to draw out its normative implications. Two exceptions to this generalization are Nadelhoffer (2004), who worries that the side-effect effect spells trouble for juror impartiality, and Wible (2009), who explores some potential consequences for business ethics. Thus, although the experiment in which this effect was first detected involved a business decision, surprisingly little work has been done to develop the implications of the side-effect effect for business ethics, and none of it has connected the effect with the concepts of virtue and vice.

In this article, we attempt to fill that gap. We begin by briefly reviewing the evidence for the side-effect effect, as well as Alfano et al.'s (2012) interpretation of that evidence, which we take to be the most plausible explanation of the side-effect effect in the literature. Next, we critique Wible's exploration of the side-effect effect in the business ethics context, arguing that his distinctive contribution is based on a misunderstanding of its nature. After clearing the ground in this way, we discuss three new experimental findings. First, all extant studies of the side-effect effect have dealt with attributions of propositional attitudes, such as intention, belief, knowledge, desire, and advocacy; we expand the scope of the effect to include the character traits of compassion and callousness. Second, we report the surprising result that people are more inclined to attribute intentionality to a norm-violating *side-effect* than to a norm-violating *primary effect*.¹ That is to say, people seem to be more willing to say that someone intentionally brings about a norm-incongruent side-effect than that he intentionally brings about the same effect when he explicitly announces that that is what he means to do. This surprising result suggests that business leaders should be especially wary of producing negative externalities as side-effects. Third, we report results suggesting that consumers and other counterparties are willing to punish vicious businesspeople by not doing business with them and to pay a premium to deal with virtuous businesspeople.² Finally, we

conclude by exploring various applications of these results to the context of business ethics. We argue against the view that in the business context, morality requires only a bare minimum of not violating any norms. Instead, businesspeople are generally expected to be virtuous. Merely producing norm-conforming side-effects is insufficient to be regarded as virtuous. Instead, the norm-conforming effects must be primary and intended. One of the most effective means of appearing virtuous, we contend, is actually being virtuous. Further, since our results indicate that consumers will reward those who appear virtuous and punish those who appear vicious, virtue appears to be in a businessperson's self-interest. Finally, we note that people's notion of virtue that emerges from our data is consonant with a Neo-Aristotelian conception of virtue.

The Side-Effect Effect

Although Knobe (2003) did not set out to examine issues in business ethics, he used an executive decision scenario in his seminal study of the side-effect effect. In this experiment, participants read one of the following two vignettes (emphasis ours):

Harm: The vice-president of a company went to the chairman of the board and said, "We are thinking of starting a new program. It will help us increase profits, but it will also **harm** the environment." The chairman of the board answered, "I don't care at all about **harming** the environment. I just want to make as much profit as I can. Let's start the new program." They started the new program. Sure enough, the environment was **harmed**.

Help: The vice-president of a company went to the chairman of the board and said, "We are thinking of starting a new program. It will help us increase profits, and it will also **help** the environment." The chairman of the board answered, "I don't care at all about **helping** the environment. I just want to make as much profit as I can. Let's start the new program." They started the new program. Sure enough, the environment was **helped**.

Participants were asked to what extent they agreed with the claim that the chairman intentionally harmed or helped the environment. As you can see, the two vignettes are identical in all respects except the valence of the predicted side-effect. Despite this structural similarity, participants were more inclined to say that the chairman intentionally harmed the environment than that he intentionally helped it. This is an example of the side-effect effect.

Since 2003, there has been a lively debate among philosophers and psychologists about the best explanation of

¹ Although the convention in philosophy of mind would designate the opposite of a side-effect a "main effect," we use the expression "primary effect" in order to avoid confusion with the statistical expression "main effect," which will be used to describe the results below.

² Note that we are discussing here ordinary people's *attributions* of mental attitudes, virtues, and vices. There is of course a question whether the people's patterns of attributions are right or even roughly correct. Nevertheless, from the point of view of business leaders, people's attributions matter even when they are mistaken. As we show in the "New Experimental Results on Virtue and Vice" section, people are less inclined to do business with a company whose leaders they regard as vicious.

the effect. Knobe (2003) initially proposed that it indicated that the concept of intentional action was essentially moralized, and that normativity was somehow built into the very concept of intentionality. However, in the ensuing years, it has become clear that the scope of the side-effect effect is much broader than initially assumed, both in terms of the types of states it applies to and in terms of the types of norm-violations that trigger it. It is not just about attributions of intentionality. The effect crops up for attributions of cognitive attitudes such as belief (Beebe and Jensen 2012, Beebe, in preparation), knowledge (Beebe and Buckwalter 2010), and remembering (Alfano et al. 2012). It also shows up for attributions of motivational attitudes such as deciding, being in favor of, and advocacy (Pettit and Knobe 2009). Moreover, it is not just about attributions following moral norm violations. The effect can be triggered by having the protagonist violate a prudential norm (Knobe and Mendlow 2004), an esthetic norm (Knobe 2004), a legal norm (Knobe 2007), a conventional norm, or even purely a descriptive norm (Alfano et al. 2012).

The most reasonable interpretation left standing after considering the diverse circumstances in which the side-effect effect arises is the heuristics approach of Alfano et al. (2012). According to this view, people tend to employ a pair of interlocking heuristics. First, they are disposed to notice and form beliefs about others' norm violations. When somebody does something contrary to a norm, people take note. This is why attributions of all kinds of mental states increase when the target of the attribution violates a norm. We can reconstruct the attributor's reasoning as follows: "He did x which goes against norm y , which he should have foreseen. So he must have formed the intention to violate the norm (formed the desire to violate the norm/believed that he would violate the norm/known that he would violate the norm)." Second, when deliberating about what to do, people are sensitive to potential norm violations, and are therefore more disposed to form intentions, desires, beliefs, and so on, in exactly the situations where others are more inclined to attribute those states to them.

As Alfano et al. emphasize, this explanation of the side-effect effect not only saves the phenomena, but also makes the pattern of people's attributions of attitudes understandable, even rational. It makes sense to pause and deliberate when the action you are about to take would violate a salient norm. Deliberation in turn leads to the formation of various mental attitudes, such as beliefs, desires, and intentions. Since we have limited cognitive resources, it is sensible to curtail our deliberative engagement to those cases where ignorance (particularly of whether a norm is being violated) would be most deleterious. Such cases typically involve the violation of norms, and so it would be reasonable to deliberate more about potential norm-violation than about potential norm-conformity.

If this is right, then one prima facie plausible upshot of the side-effect effect turns out to be mistaken. It might seem immediately obvious in light of this effect that profit-driven businesses and businesspeople risk being blamed for bad effects they produce (whether as primary effects or as side-effects), but are not in a position to receive credit for good effects they produce (again, whether as primary or side-effects). From this, it would seem to follow that businesses should concentrate only on avoiding negative effects and not on producing positive externalities. For instance: do not pollute, but also do not bother about improving environmental conditions. This seems to be the conclusion that Wible (2009) draws. The heuristics story suggests otherwise. In fact, we will argue that if the heuristics interpretation of the side-effect effect is correct, businesses have good reason to avoid negative externalities (both as primary and as side-effects) and to produce positive ones (especially as primary effects). In the next section, we critique Wible's application of the side-effect effect to business ethics, then, in the subsequent sections, present our own application.

Wible's Interpretation

Little work has been done to explore the ramifications of the side-effect effect for business ethics. As we mentioned above, so far only Wible (2009) has done so. While we appreciate his expansion of the side-effect effect into the realm of business ethics, we believe that some of his conclusions deserve new examination. Since the publication of Wible's article, additional research has both expanded the scope of the effect and forced a reinterpretation of its nature, and the application of the effect to business ethics must be updated in light of these subsequent developments. In particular, we want to explain why some of Wible's conclusions will no longer work, and then use his conclusions as a stalking horse; once we see what the implications of the side-effect effect are *not*, we can start to reconsider what they really *are*.

Originally, intentionality was the only psychological state thought to be subjected to the side-effect effect. So, not surprisingly, Wible's analysis and application of the side-effect effect is based on this view. In particular, he assumes that knowledge and belief are not subjected to the effect (2009, p. 175). In reality, however, the phenomenon has a much wider scope than originally thought. As we pointed out above, people are more inclined to attribute knowledge and belief (as well as many other mental attitudes) to protagonists in norm-incongruent conditions than in norm-congruent conditions. Furthermore, Wible seems to subscribe to a mixture of Knobe's conceptual competence model and Nadelhoffer's (2004) distortion model of the side-effect effect, claiming that moral considerations as well as praise and blame mediate the effect (2009, p. 176).

As Alfano et al. (2012) have shown, though neither of those theories is consistent with the full range of data on the side-effect effect, some studies have found the side-effect effect in cases without a moral norm at play, such as Knobe (2004) with an aesthetic norm and Alfano et al. (2012) with conventional and descriptive norms. If praise and blame mediated the effect, it would not appear in such non-moral cases.

Wible's first conclusion is that praise and blame do not attach directly to the valence of the side-effect (2009, p. 176). He contends that people do not blame the protagonist for producing a bad effect (or praise him for producing a good effect) as such. Instead, they praise or blame the actor's motives, using both what he says and the effects he produces as evidence of those motives. So the chairman is morally blameworthy not for what he does as such, but for why he does what he does, namely only caring about profits. The moral failing of the chairman is that he lacks the proper motive of also caring for the environment. Wible surmises that this moral failing of the chairman is responsible for the side-effect effect. If, in the **harm** condition, the chairman had not been so immoral in his singular fixation on profit, then he would not have ended up harming the environment. Thus, people have good reason to blame him for producing this deleterious side-effect. People also regard the side-effect as intentional because it is a direct result of his explicitly immoral, profit-seeking behavior, according to Wible. Conversely, in the **help** condition, people find nothing meritorious or intentional in the chairman's helping the environment because of his prior moral failing of only caring for profits. The chairman should care about profits *and* the environment.

This is a tempting conclusion to draw, especially for the purpose of business ethics, as it lets us condemn single-minded pursuit of profit. While initially plausible, new evidence appears to challenge the claim. This explanation appears to work for the original 2003 vignettes, but falls apart when applied to subsequent side-effect effect studies. As we explained in the previous section, the effect also crops up in attributions of motivational attitudes such as desire, decision, and advocacy. People are more inclined to say that someone wanted to produce a norm-incongruent side-effect than a norm-congruent one. Furthermore, as Alfano et al. (2012) have shown, the side-effect effect appears when the norm violated is purely descriptive, with no moral content whatsoever.³ In such a case, praise and blame cannot be invoked

as mediating factors. For instance, they tested a case of a non-moral, descriptive norm with the following vignette:

Jessica lives in a neighborhood where everyone (including Jessica herself) happens to own a dog. One afternoon, she is planning to go for a walk and decides *not to/to* take her dog. Her friend Aaron says, "Jessica, if you go out like that, you *will/won't* be doing what everyone else is doing." Jessica responds, "I don't care at all what everyone else is doing. I just want to go for a walk *without/with* my dog." She goes ahead with her plan, and sure enough, she ends up doing what *everyone/no one* else is doing.

The side-effect effect was detected in participants' attributions of intentionality to Jessica. They were more inclined to say that she intentionally did what no one else was doing (walking without a dog) than that she intentionally did what everyone else was doing (walking with a dog). For the models that Wible employs to explain this effect, people would have to *blame* Jessica for her lacking the allegedly proper motive of caring about what everyone else happens to be doing. The doxastic heuristics model of Alfano et al. however, easily accounts for this finding. Jessica violated a norm—a merely descriptive norm, yes—but a norm nonetheless. That is the sort of thing people instinctively take note of, and about which they are therefore more inclined to attribute propositional attitudes. Morality, praise, and blame do not enter into it. This is not to say that the side-effect effect supplies no lessons on the role of profit seeking by businesses, a point on which we will present our own conclusions later.

Following on the previous conclusion, Wible presents a second understandable but mistaken conclusion based on Knobe's (2003) study, namely that a profit-focused businessperson can be credited with *intentionally* producing a good side-effect if she or he is a moral agent with other motives in addition to profit-seeking. Wible describes a variant on the original vignette in which the chairman responds to the news that the new program will increase profits and help the environment by saying, "Great! I care about helping the environment. I am happy that we can help the environment and make a profit at the same time. Let us start the program." In this version of the vignette, the chairman does not rank the profit motive above the environmental impact. In fact, it might even seem that he is more concerned with helping the environment, since that is what he mentions first. The conclusion based on this vignette therefore contains a subtle elision of the difference between primary effects and side-effects. If an agent has only one goal (making a profit) and in achieving that goal (implementing the new program) produces a side-effect, whether or not the side-effect is generally considered intentional depends on whether or not a norm was violated

³ Knobe (2004) also found the side-effect effect in cases involving an aesthetic norm. It is preferable to adopt an explanation of the side-effect effect that covers all cases of the phenomena rather than to have several theories; each account for only some of the data. Wible's explanation then seems inadequate on this criterion, while Alfano and company's heuristics account succeeds. So far, Alfano et al. (2012) is the only attempt we know of that attempts to be consistent with all the published side-effect effect studies.

when the agent acted. If, however, an agent has two goals—such as making a profit *and* helping the environment—then by definition, there is no side-effect; there are two primary effects, both of which are intentional. When a businessperson who strives to increase profits and help the environment enacts a program that achieves both goals, neither effect is a side-effect. In the next section, we discuss new experimental data that suggest exactly this: it is possible to produce a good effect while making a profit and to be praised for that good effect, but only when it is a second primary effect, not a side-effect.

Wible bases his final conclusion on the obligations of businesspeople on Thomson's (1971) notion of a "Minimally Good Samaritan." Thomson contrasts the Maximally Good Samaritan—who will aim for good results (like helping the environment) to the point of sacrificing his own life—with the minimally good Samaritan—who will do some good, but not when the costs are high. Wible applies this distinction to business ethics, arguing that businesspeople have a moral obligation only to live up to the standard of the minimally good Samaritan, not the maximal one. In terms of Knobe's chairman, a businessperson is obligated to produce the good side-effect of helping the environment, but only if profits also increase. We can call this view this **Minimally Good Businessperson Theory**.

We not only take Wible to be on the right path with this conclusion, but also contend that his view can be improved upon. First, as we noted regarding the previous conclusion, Wible thinks a minimally good Samaritan can intentionally produce good side-effects, and it is his obligation to do so when possible. But, a closer examination of the case in which this happens reveals that what is produced is a second main effect, not a side-effect. Hence, this conclusion needs to be recast in terms of two primary effects, instead of one primary effect and one side-effect. Second, virtue theory is in a better position to express the moral obligations of the businessperson than competing theories. Wible will have to claim that there is a myriad of good and proper motives that a businessperson ought to have in addition to seeking profits; helping the environment is but one. Alternatively, we could require that businesspeople have the motives of increasing profits and being virtuous, or—if you prefer a slightly more nuanced view—of having (and so demonstrating when appropriate) enough of the virtues. Third, while Wible only speaks of what businesspeople's motives ought to be, we can offer a reason why; being virtuous appears also to be in their self-interest. To develop these points, we must now turn to our new experimental findings.

New Experimental Results on Virtue and Vice

We turn now from our critique of Wible to the articulation of the real upshot of the side-effect effect for business

ethics. To this end, we think that the ethical implications of the side-effect effect will be clearer by expanding our focus beyond the mere intentionality of the side-effect to also include attributions of character traits, especially virtues and vices. Before turning to new empirical evidence on this score, let us begin with a thumbnail sketch of the nature of virtue and vice. Next, we argue that the side-effect effect extends beyond the attribution of propositional attitudes (belief, knowledge, intention, desire, etc.) to character traits. If this is right, and if businesses and business leaders want to be thought of as virtuous rather than vicious, then certain normative and prescriptive conclusions follow.

The notion of a virtue is not univocal, but one attractive conception is the Neo-Aristotelian theory according to which a virtue is a complex disposition comprising a cluster of sub-dispositions, including a perceptual sensitivity, a tendency to construe ambiguous situations in particular ways, a motivation, deliberative excellence, and the ability to carry out intentions reliably and successfully (Alfano 2013; Zagzebski 1996, p. 137). For example, generosity is not just the disposition to give resources to people, or even to give resources to people who need them. Such a disposition seems insufficiently reason-guided to count as a virtue. Rather, the generous person is disposed to *notice* opportunities for giving, to *construe* ambiguous cues charitably, to *want* to help, to *deliberate soundly* about what would help in each particular circumstance, and to *act* reliably and successfully when she intends to help. Note that many of these sub-dispositions refer to attitudes that are subject to the side-effect effect. Noticing and construal lead to belief, knowledge, and remembering; wanting is a form of desire; and acting involves intentions. It seems plausible, then, that people will be more inclined to attribute a virtue to someone if he or she violates a relevant norm while acting in accordance with a virtue than if he or she acts in accordance with a virtue without violating a norm. Of course, the norm violated cannot be a moral norm (since then he or she would not be acting virtuously), but as we explained above, the side-effect effect is triggered by all sorts of violations, including violations of self-regarding norms and even descriptive norms.

A vice is not merely the absence of some virtue; it is the presence of a morally bad disposition that is structurally similar to a virtue. Thus, a vice is also a complex disposition comprising a cluster of sub-dispositions, including a perceptual sensitivity, a tendency to construe ambiguous situations in particular ways, a motivation, deliberative excellence, and the ability to carry out intentions reliably and successfully. While the compassionate person notices suffering, construes ambiguous cues correctly, wants to alleviate suffering, deliberates well about how to do so in the particular circumstances, and reliably acts on the intention to do so, someone with a correlative vice inverts

at least one of the sub-dispositions. For instance, a sadistic person notices suffering and construes ambiguous cues correctly, but what he wants is not to alleviate suffering but to cause or at least observe it. A callous person tends to notice suffering and construe ambiguous cues correctly; he just does not care about others' well being. This tendency is in contrast with the merely careless person who may want to help alleviate suffering but is too morally obtuse to notice it, as well as the weak-willed person who notices, construes, desires, and deliberates well but who—when it finally comes to acting—fails to carry out his good intentions. Since the component parts of a vice are structurally similar to those of a virtue, it might seem by parity of reasoning that people will be more inclined to attribute a vice to someone if she violates a relevant norm while acting in accordance with a vice than if she acts in accordance with a vice without violating a norm. The parity breaks down, however, because there is no way to act in accordance with a vice without violating a norm, namely the norm not to be vicious. Thus, all vicious behavior is likely to trigger the side-effect effect, making vice attributions more prevalent than virtue attributions. This is a worrisome implication for business leaders.

In order to more systematically investigate the nature of the relationship between norm violations and the attribution of vice and virtue, the above points were tested experimentally. Participants ($N = 81$, $M_{\text{age}} = 34.4$, 40 women) were recruited and compensated using Amazon.com's Mechanical Turk. The sample was restricted to adults living in the United States.

Participants were presented with one of four vignettes describing the CEO of a corporation who is approached by his vice-president. In all the four vignettes, the vice-president presents the CEO with a new policy that will produce an increase in profits. Two of them involve increasing the workforce by 10 %, while the other two involve the firing of 10 % of the current employees. The vignettes also varied with respect to the CEO's reaction to the vice-president's proposal. Sometimes the CEO only cares for increasing profits, making the change in the number of employees (either up or down) a side-effect. At times, the CEO has two goals, increasing profits *and* hiring or firing employees (depending on the experimental condition). In these cases, the hiring or firing is no longer a side-effect, but a primary effect. Consequently, the outcome (**hire** or **fire**) is brought about in one of two ways: as a primary effect or as a side-effect. The result is a 2 (outcome) \times 2 (effect-type) factorial design. Each participant viewed only one of the four vignettes, and therefore, was only included in one of the four experimental conditions.⁴

After reading the vignette, participants were asked to rate their level of agreement with four statements: (1) *The CEO intentionally harmed/helped the workers.* (2) *The CEO was callous.* (3) *The CEO was compassionate.* (4) *I would avoid purchasing products produced by the company, even if I had to pay more.* These questions were presented in random order, and were meant to assess participants' attributions of intentionality, vice, virtue, and the degree to which the CEO's conduct would influence their own future purchasing behaviors. Ratings of agreement were on a Likert scale ranging from 1 (*Strongly Disagree*) to 7 (*Strongly Agree*).

Let us begin with participants' ratings of intentionality across the four vignettes (see Fig. 1). With respect to the intentionality ratings, the results of a two-way analysis of variance (ANOVA) indicated significant main effects of both the outcome, and the effect-type conditions. These main effects indicate that both the outcome and the effect-type variables influence participants' ratings of the CEO's intentionality. Specifically, when the outcome is **fire**, the CEO is more likely to be perceived as having acted intentionally than when the outcome is **hire**: $F(1, 77) = 59.27$, $p < .001$, $\eta_p^2 = .44$. Likewise, when the effect-type is a primary effect, the CEO is also more likely to be perceived as having acted intentionally than when the effect-type is a side-effect: $F(1, 77) = 9.04$, $p < .01$, $\eta_p^2 = .11$. However, we should be cautious in drawing conclusions from these main effects, since a significant interaction of outcome and effect-type on intentionality ratings was also observed: $F(1, 77) = 33.42$, $p < .001$, $\eta_p^2 = .30$. This interaction indicates that the effect-type (i.e., primary effect vs. side-effect) influences participants' intentionality ratings differently depending on the outcome condition (i.e., **hire** vs. **fire**). We see this interaction in Fig. 1. When **fire** is the outcome, we see that participants in the side-effect condition gave a slightly higher mean intentionality rating than participants in the primary-effect condition. But when **hire** is the outcome condition, participants in the side-effect condition gave much lower ratings of intentionality than participants in the primary-effect condition. Essentially, the effect-type variable influences intentionality ratings differently depending on whether the outcome is hiring or firing.

Follow-up simple effects tests revealed that within the **hire** condition, participants in the primary-effect condition provided significantly higher intentionality ratings compared to those in the side-effect condition: $F(1, 77) = 38.24$, $p < .001$, $\eta_p^2 = .33$. This result indicates that people are not inclined to say that a business leader intentionally brings about a beneficial side-effect, but that they are inclined to say that he intentionally brings about a beneficial, second primary effect.

It is worth mentioning two additional points regarding Fig. 1. First, note that intentionality ratings are stable

⁴ See Appendix for the full text of each of the vignettes.




Fig. 1 Summary of mean intentionality ratings by condition. Error bars indicate ± 1 standard error of the mean

across the **hire** and **fire** conditions when the outcome is a primary effect. Essentially, when the CEO asserts that hiring or firing is one of his principal concerns, the outcome does not impact his perceived intentionality. In contrast, the intentionality ratings differ greatly between the **hire** and **fire** conditions when the outcome is a side-effect. This difference proved significant in follow-up tests: $F(1, 77) = 87.68, p < .001, \eta_p^2 = .53$.

Second, and somewhat surprisingly, within the **fire** outcome condition participants were more likely to attribute intentionality when the CEO fired workers as a side-effect than when firing workers was the primary effect. This difference proved statistically significant in follow-up tests: $F(1, 77) = 3.89, p = .05, \eta_p^2 = .05$. Thus, when the CEO fired his employees, participants were more likely to say that he did so intentionally *when it was a side-effect* than when it was a primary effect. This is a surprising result, and one that—to our knowledge—has not yet been reported in the literature. We think it might be possible to interpret it with the Alfano et al. heuristics theory, but doing so goes beyond the scope of this article; at the very least, this is a finding that merits further investigation.⁵

Next, we asked participants to rate the callousness of the CEO (see Fig. 2). A two-way ANOVA again revealed a main effect of both the outcome condition and the effect-type condition. Just as above, these main effects indicate that both the outcome and the effect-type variables influence participants' ratings of the CEO's callousness. Not surprisingly, when the outcome is **fire**, the CEO is more likely to be perceived as being callous than in the **hire** outcome condition: $F(1, 77) = 29.92, p < .001, \eta_p^2 = .28$. Likewise, when the effect-type is a side-effect, the CEO is

⁵ The closest thing to a study contrasting attributions in side-effect and primary-effect cases is Cova et al. (2012). They contrasted side-effect cases with means-to-an-end cases, and found that the side-effect cropped up for both, but that means were overall judged to be more intentional than side-effects. We are encouraged by this finding, since it seems plausible to think of means as falling between primary effects and side-effects on the scale of intentionality.

also more likely to be perceived as being callous than when the effect is a primary one: $F(1, 77) = 14.85, p < .001, \eta_p^2 = .16$. We should again be cautious in drawing conclusions from these main effects: however, since we find a statistical interaction of outcome and effect-type on ratings of callousness: $F(1, 77) = 15.55, p < .001, \eta_p^2 = .17$. This interaction indicates that the effect-type influences ratings of callousness differently depending on the outcome condition. As we see in Fig. 2, within the **fire** condition, participants in the side-effect condition and the primary-effect condition gave callousness ratings that were quite similar (and high). Yet within the **hire** condition, callousness ratings seem to differ quite substantially depending on the effect-type condition. Follow-up tests confirmed a statistically significant difference indicating that participants were more likely to view the CEO as callous if the hiring of more workers was a side-effect rather than a primary effect, $F(1, 77) = 30.11, p < .001, \eta_p^2 = .28$.

Notice in Fig. 2 that in three of the four experimental conditions the CEO is viewed as quite callous. It is only when the CEO's primary goal was hiring that he was viewed as less callous. Furthermore, within the **fire** condition, ratings of callousness do *not* differ significantly depending on whether the firing was brought about as a primary effect or a side-effect, $F(1, 77) = 0.00, ns$. A CEO who fires his employees as a side-effect is just as likely to be perceived as callous as one who does so as a primary effect.

If we turn our attention to ratings of compassion, we find essentially the reverse of the pattern reported above (see Fig. 3). Using a two-way ANOVA we observed significant main effects of both the outcome and the effect-type conditions. Once again, this indicates that the effect-type and outcome variables are significant predictors of participants' ratings of the CEO's level of compassion. Not surprisingly, with respect to the outcome condition, when the CEO hires employees: he is more likely to be viewed as compassionate than when he fires employees, $F(1, 77) = 31.22, p < .001, \eta_p^2 = .29$. Regarding the effect-type condition, the CEO is more likely to be viewed as compassionate when the outcome is a primary effect than when the outcome is a side-effect: $F(1, 77) = 14.35, p < .001, \eta_p^2 = .16$. But again, we should be careful in our interpretation of these main effects since we also find a significant statistical interaction of outcome and effect-type on ratings of compassion: $F(1, 77) = 18.13, p < .001, \eta_p^2 = .19$. This interaction indicates that effect-type influences compassion ratings differently depending on the outcome condition. This finding is apparent in Fig. 3, where we see that, within the **fire** condition, ratings of the CEO's level of compassion from the primary-effect and the side-effect conditions are fairly similar. However, within the **hire** condition, we can see that participants who read about the CEO who hired




Fig. 2 Summary of mean ratings of the CEO's callousness by condition. Error bars indicate ± 1 standard error of the mean

employees as a primary effect were more likely to rate him as compassionate compared to those participants who read about the CEO who hired employees as a side-effect. Follow-up tests confirmed that this was a significant difference. Within the **hire** condition, participants gave significantly higher compassion ratings when hiring was a primary effect compared to when it was a side-effect: $F(1, 77) = 32.06$, $p < .001$, $\eta_p^2 = .29$. This result is to be expected. When the CEO asserts that hiring employees is one of his principal concerns, he is more likely to be viewed as compassionate. Within the **fire** condition, however, follow-up tests find no significant difference between compassion ratings when the CEO fired employees as a primary effect versus when he fired employees as a side-effect, $F(1, 77) = 0.11$, *ns*. If the CEO fires his employees, then he is unlikely to be perceived as compassionate regardless of whether the firing was his primary objective or a side-effect.

Finally, we examined participants' self-reported inclinations to avoid purchasing from the CEO's company (see Fig. 4). A two-way ANOVA revealed main effects for both the outcome condition and the effect-type condition. This indicates once again that the effect-type and the outcome conditions influence participants' inclinations to avoid purchasing from the CEO's company in the future. More specifically, when the outcome is the firing of employees, participants are more likely to report that they would avoid purchasing from this company rather than when the outcome is the hiring of employees: $F(1, 77) = 23.39$, $p < .001$, $\eta_p^2 = .23$. And similarly, when the CEO brings about the outcome as a side-effect, participants are also more likely to avoid purchasing from his company rather than when the outcome is a primary effect: $F(1, 77) = 6.45$, $p < .05$, $\eta_p^2 = .08$. However, in our analysis of correlations, we found that the inclination to avoid purchasing from the company was highly correlated with participants' ratings of the CEO's callousness, $r = .50$, $p < .001$. Therefore, in order to control for the influence of perceived callousness on the inclination to avoid purchasing, we included

callousness as a covariate in a two-way analysis of covariance (ANCOVA). Callousness proved to be a significant predictor in the model: $F(1, 77) = 9.00$, $p < .01$, $\eta_p^2 = .11$, and the analysis revealed a significant main effect of outcome: $F(1, 77) = 7.43$, $p < .01$, $\eta_p^2 = .09$. But interestingly, when the data were analyzed in this more nuanced way, there was no longer a significant main effect for the effect-type condition on the inclination to avoid purchasing: $F(1, 77) = 1.53$, *ns*. Thus, after we controlled for the influence of callousness ratings, the only significant predictor of participants' inclination to avoid purchasing was whether or not the CEO fired employees. The difference between this ANCOVA model and the original ANOVA model seems to indicate that the influence of the effect-type condition on the inclination to avoid purchasing is at least partly explained by participants' ratings of the CEO's callousness. This mediational hypothesis was tested empirically using a bootstrapping procedure (Bollen and Stine 1990) with 5,000 draws in the `lme4` package of R (Zhang and Wang 2009). Results confirmed that the influence of the effect-type condition on the inclination to avoid purchasing is mediated by participants' ratings of the CEO's callousness. We observed a significant indirect effect with callousness ratings mediating the relationship between the effect-type variable and the inclination to avoid purchasing: $B = -.52$, 95 % CI $[-1.11, -.013]$.^{6,7} This suggests perceived callousness is one reason the effect-type influences the inclination to avoid purchasing, which in turn indicates that business leaders have especially strong reasons not to appear callous or otherwise vicious. Their reputations will suffer and people will be less inclined to do business with them in the future, a point we develop further in the next section.

One final note before we turn from these data to a sustained discussion of their upshot for business ethics. As we pointed out in the [introduction](#) section in this article, the original experiment documenting the side-effect effect involved a business context in which the business leader makes decisions that impact the environment. Environmental help and harm are only some of the externalities that business leaders need to take into account if they care about how their actions are interpreted and their moral character assessed. In this section, we described studies related to one of those other externalities: contributing to (un)employment. There are certainly others. On the

⁶ We are grateful to Joshua Knobe for suggesting we test this hypothesis.

⁷ This mediational model was also tested using not only the more common—but also more conservative—Sobel test (MacKinnon et al. 1995). The results mirrored those of the bootstrapping test of mediation; there was a significant indirect effect with callousness mediating the relationship between the effect-type variable and the self-reported inclination to avoid purchasing, $Z = -2.47$, $p < .05$.




Fig. 3 Summary of mean ratings of the CEO's compassion by condition. *Error bars* indicate ± 1 standard error of the mean

positive side, businesses can provide useful services and products, pay taxes, and afford opportunities for the cultivation of skills. On the negative side, they can provide hazardous services and products, shirk their tax responsibilities, endanger workers, and alienate workers from the means of production. We have not yet investigated whether the side-effect effect will crop up in relation to all of these potential externalities, but given the effect's robustness so far, we have good reason to suspect that it will. As such there is good reason for businesspeople to apply the normative lessons from our study generally to a wide range of externalities, and not just to matters of contributing to (un)employment. So let us turn to those lessons now.

Application to Business Ethics

In interpreting our experimental results and applying them to business ethics, we will begin by extrapolating specific lessons from our data, points immediately entailed by our findings. From there, we will expand to some larger conclusions that the data suggest. The first and most basic lesson here is that businesspeople must be aware of and conform to norms relevant to their behavior in the workplace. If the side-effect of a businessperson's action—be she a CEO or a mailroom worker—violates a norm that is (or should be) salient to her, then she will be regarded as having intentionally violated the norm and blamed for it. Even worse, she will be seen as vicious! That means that if your side-effect violates a norm, others will see you as having a general disposition to behave that way. They'll generally expect you to violate at least that norm again in the future, and perhaps other norms as well. Your very reputation as a businessperson is at stake.

However, we should not draw the further conclusion that one can only be blamed (and thought vicious) for violating a norm, but never praised (and thought virtuous) for adhering to it. Such a conclusion would advocate a




Fig. 4 Summary of means for the inclination to avoid purchasing from the CEO's company by condition. *Error bars* indicate ± 1 standard error of the mean

bare-minimum theory of business ethics, where businesspeople should *only* concern themselves with a basic decency of not violating norms, while they need not worry about being especially virtuous either. Wible considered and rejected this view as inadequate. On this point we concur and roundly reject such a theory of business ethics. In place of the bare-minimum theory, Wible offered his minimally good businessperson theory. While we regard Wible to have been on the right track, our results indicate that his theory does not go far enough. We can now explain why more is required of businesspeople. In their stead, we propose our **Compassion-Expectation Theory**, which—briefly put—claims that businesspeople are expected to and ought to be compassionate (at least when doing so is consistent with profitability),⁸ and people will tend to reward them for being compassionate and punish them for being callous.

The simple explanation for the insufficiency of both the bare minimum and minimally good theories is that they violate people's expectations. Our findings indicate that people generally expect more of a CEO than merely not to violate norms. All three theories predict that the CEO in both fire conditions to be callous. What the compassion-expectation theory predicts and finds as well, though, is that the CEO who increased hiring a side-effect would also be regarded as callous. The bare minimum view neither predicts nor explains this result, since the CEO did not violate a norm. As it turns out, people typically expect more of businesspeople than mere norm conformity. Increased hiring as a side-effect should also be sufficient to meet the requirements of Wible's minimally good businessperson

⁸ We cannot provide any conclusions on the basis of our data here as to what degree virtuous behavior is expected or required of businesspeople at the cost of profitability. Despite being a legitimate and important question, it goes beyond the scope of our study which focused solely on cases in which profitability was not a variable.

view. People, however, apparently expect businesspeople to help intentionally.

One might object at this point that intuitions about intentionality, compassion, and callousness are not appropriate sources for positing or rejecting a normative theory of business ethics. Not only are the intuitions we probed highly fallible, but they are also the intuitions of ordinary people and hence theoretically uninformed or un-illuminating. This concern would have merit were it not for two points. First, the view expressed by our participants does not stand alone; it conforms to the Neo-Aristotelian notion of virtue (and vice). In particular, our results are evidence for the theoretical link between intentionality and virtue. Second, satisfying people's moral expectations is in the CEO's self-interest since doing so appears likely to increase profits. As our mediational results indicate, participants' assessments of virtue or vice mediated the impact of effect-type on the desire to do business with a company.

On the first point, the intuitions of ordinary people expressed in our findings conform to the Neo-Aristotelian view: virtue is not accidental. According to this view, one must be aiming at and striving for virtuous ends to achieve them. Accidentally realizing a good effect (i.e., an effect that not only does not violate norms, but also achieves some social good) does not make one virtuous. Rather, the Neo-Aristotelian view holds that in order for an effect to be a product of virtue, it must be intentionally produced. That is because virtue itself cannot be accidentally achieved; it must be intended and strived for result. For instance, to act courageously, not only one must overcome a threat, but one must also perceive it *as* a threat. Someone cannot perform an action and discover afterward that it was courageous because (unbeknown to her at the time) her life was in danger when performing the action. Likewise, performing an action that just so happens to end up helping someone in need is insufficient to establish the virtues of compassion or charity. The agent must recognize the action at the time as helpful to those in need.

This theoretical view closely matches our data. Our results (along with a host of others on the side-effect effect) indicate that participants are significantly less willing to attribute intentionality to a side-effect that merely does not violate a norm. Participants were much less willing to attribute the virtue of compassion to the CEO who hired but only as a side-effect. Rather they only considered the CEO compassionate when he hired more employees as an *additional* primary effect. These results are not odd; they are consonant with the Neo-Aristotelian picture of virtue. If the moral intuitions people reported were unsystematic and couldn't be given a theoretical underpinning, then it would be justifiable to ignore when formulating a normative theory of business ethics. Fortunately, that is not the case here. Our results are evidence for the requirement that

virtues (in this case compassion) require intentionality, as claimed for a long time by the Aristotelian tradition of virtue ethics. That point alone provides sufficient reason to consider people's intuitions in normative theory construction. Luckily, in the case of our compassion-expectation theory, we also have a second reason.

The second reason we have for looking to people's intuitions as having implications for a normative theory of business ethics goes back to our finding that people's perceptions of callousness help us explain why the effect-type impacted their stated inclination to avoid purchasing from the CEO's company. This mediation indicates how people are responding to a problematic case. It is reasonable to consider why the type of effect (primary or secondary) impacts people's inclination to do business with a company, as we observed. As it turns out, in this case, their perceptions of compassion and callousness explain this relationship.⁹ The more callous they found the CEO, the less inclined they were to do business with him in the future. Three of our four CEOs were regarded as more callous, and participants were generally more willing to punish him by avoiding future purchases. On the other hand, they were less willing to punish the CEO whom they regarded as more compassionate.¹⁰ Appearing compassionate (and not callous) seems to be profitable.

One might worry at this point that all we have found are self-reports of whether or not participants will avoid purchasing from a businessperson, and they may be very biased. People might want to report the intention to financially punish vicious businesspeople and reward virtuous ones because it would make them feel better about themselves as good and virtuous people. But these biased self-reports may diverge from how they actually behave in the marketplace. Consumers may *say* they will only buy from virtuous businesspeople, but that might not match how they actually spend their money.

Though this concern has some merit, there is empirical evidence that suggests consumers are not completely deceiving themselves. People at least sometimes actually do reward companies they find to be honest and hard working and punish those they see as vicious. To offer a couple of recent examples, consumers punished BP after the recent environmental disaster in the Gulf of Mexico (Swint 2010). On the other hand, when Ford did not take bailout money during the Great Recession (while GM and Chrysler did) consumers rewarded Ford by buying more of their cars than those manufactured by other American

⁹ See Sripada (2011) for a similar point about mediation.

¹⁰ Presumably, participants who report themselves to be less willing to punish the compassionate CEO will be more likely to praise him and reward him with future business. Despite being intuitive, this move is theoretical since our study only asked about participants' willingness to punish.

carmakers (Williams 2010). While consumers may not *always* punish the vicious and reward the virtuous, there is plenty of evidence indicating that they often behave this way.

Admittedly, perceptions of compassion and callousness are not the only factors influencing one's purchases. People do not just buy what the virtuous have to sell; they buy what they want and need. But the point remains that these perceptions do in part mediate how we intend or want to behave toward the compassionate or callous. For instance, how many of us wanted to punish BP in 2010 after their environment disaster in the Gulf of Mexico? But how many also still ended up, at least once, stopping at BP gas station to fill up, perhaps because it was the cheapest nearby alternative? In that case, our perception of callousness caused us to want to punish BP, but other factors sometimes superseded that desire. Companies compete with one another on many dimensions. A company can acquire a competitive advantage by producing more efficiently, selling at a lower price point, having higher quality products, or guaranteeing products with warranties and implicit branding. If our results are correct, then a company can also acquire a competitive advantage by appearing more virtuous than its peers. Presumably, in many cases, the cost associated with appearing virtuous is less than the cost of, for instance, reducing prices. Executives should therefore think of their reputation for virtue as one of several ways to differentiate their business. The lesson is that businesspeople have a self-interested reason to go beyond the bare minimum morality or Wible's minimally good theory and appear compassionate. Doing so will tend to increase profits.

Note that so far we have only spoken of *appearing* compassionate. Just because the CEO is thought to be compassionate by our participants based on this one action does not mean that he actually is. All we know is that he appears compassionate, which tells us that businesspeople ought to *appear* compassionate. Our theory goes further, claiming that businesspeople ought to *be* compassionate because, we want to suggest, being compassionate is one of the most efficient and reliable ways to appear compassionate.¹¹

¹¹ One might worry that the is-ought problem arises at this point, thinking that we have moved from factual statements about how participants responded to the normative conclusion about how businesspeople ought to behave. Luckily, this concern is easily allayed. One commits the is-ought fallacy by inferring a normative conclusion (an "ought" as it were) from a set of only factual premises ("is" statements). As long as there is at least one normative premise as well, one has not committed the is-ought fallacy. Such is the case here. Our premises include not only the fact that people report themselves willing to punish the callous, but also the normative assertion that business people ought be profitable. Thus, our normative conclusion is legitimate.

At this point, we move more into the realm of suggested conclusions that are at least consistent with our data. So long as there in fact is any correlation between people's reports of a willingness to punish callous businesspeople, our results suggest a **Profit-Compassion Hypothesis**: *ceteris paribus*, being a compassionate business person will tend to increase profits and being callous will tend to decrease profits. However, our results only found a correlation between *attributions* of a virtue and likelihood of avoiding purchases with the CEO's company. So we need a link between attributions of virtue and actually being virtuous to render the profit-virtue hypothesis plausible. To that end, we propose the **Appearance of Compassion Strategy**: *ceteris paribus*, one of the most efficient and reliable means for having a virtue widely attributed to you is to behave regularly in a manner consistent with that virtue. How is that accomplished? Well, one could put up a systematic pretense of virtue, perhaps abetted by a public relations campaign. Or one could just be virtuous. Presumably, both would lead to reliable behavior in accordance with virtue. The question is therefore which is more efficient: Does it make more sense just to be virtuous, and so be perceived so, or to pretend to be virtuous, all the while attempting to convince others that you are virtuous? We do not have empirical evidence that relates directly to this claim, but we think that it is plausible to suggest that for most people most of the time, it is easier just to be virtuous than to pretend to be so. Much of the time, faking virtue is a lot harder and less reliable than actually being virtuous.

Note that the appearance of compassion strategy does not claim that being compassionate is the *only* way to appear compassionate. It surely is the case that a CEO could hire a savvy public relations firm to separate appearances from reality in the minds of ordinary people. Our claim is instead that being compassionate is a reliable and efficient means for appearing compassionate. The CEO with the PR firm will need to employ their services regularly to keep up the fallacious appearance of compassion. Not only will that get expensive, there is no guarantee it will work; appearances might well catch up with him at some point.

Finally, it is worth noting that so far, we have only tested for the virtue of compassion and the vice of callousness, and found compassion to appear to be in one's self-interest. While our data cannot speak to other virtues and vices, some theoretical extensions of our results are possible. We take it that there is not nothing particularly special about compassion and callousness (in general or in a business context) such that our findings—that the side-effect effect and the mediation between trait attributions and reporting of a willingness to punish—will not show up for other virtues and vices. This expectation does not

commit us to Aristotle's unity of virtue thesis (that to have a virtue, one must possess them all). Our point here is the much weaker claim that most virtues exhibit a similar structure. Compassion involves a sensitivity to need, a disposition to construe, a motivation, deliberative excellence, and reliability in acting on the deliverances of deliberation. The sub-dispositions of compassion produce mental states the attribution of which is subject to the side-effect. Most other virtues involve similar sub-dispositions and hence produce mental states of the same type. It is therefore plausible to hypothesize that their attribution will be subject to the Knobe effect as well.

We therefore recommend further testing for other virtues, such as wisdom and prudence. We predict, first, that people will not attribute other virtues to businesspeople when they merely produce norm-conforming side-effects, but rather only when the norm-conforming effects are primary, and so intended. Second, we predict that there will be a similar mediation between participants' perception of the virtue and vice in question and their self-reported intentions to avoid doing business with those they regard as possessing the corresponding vice. However, it is not unreasonable to expect a difference in the strength of the mediation. Some virtues and vices may play a greater or smaller role in determining participants' inclination to avoid doing business with an agent. For instance, people presumably would want to punish callous businesses more than profligate businesses, but less than sadistic businesses.

To sum up, in this article, we have argued that the most attractive interpretation of the side-effect effect, along with the new experimental evidence we report, suggests that businesspeople are both expected to behave and prudentially ought to behave in accordance with compassion specifically (and likely virtue in general). Since it is plausible that the best way to achieve this end is actually to have or cultivate compassion, it follows that the best business strategy is to cultivate compassion (and likely many other virtues as well).

Appendix

Fire/Side-Effect

The vice-president of a manufacturing company was talking with the CEO. The vice-president said, "We are thinking of implementing a new policy. If we implement the policy, it will increase profits for our corporation. It will also mean that we have to fire 10 % of our employees, many of whom will have difficulty finding other work." The CEO said, "I don't care at all about the employees. I just want to make as much money as possible. Let's implement the policy."

They implemented the policy. Just as the vice-president had predicted, profits increased, 10 % of the employees were fired, and many of them were unable to find other work.

Hire/Side-Effect

The vice-president of a manufacturing company was talking with the CEO. The vice-president said, "We are thinking of implementing a new policy. If we implement the policy, it will increase profits for our corporation. It will also mean that we have to increase our workforce by 10 %, hiring many people who would have difficulty finding other work." The CEO said, "I don't care at all about the employees. I just want to make as much money as possible. Let's implement the policy."

They implemented the policy. Just as the vice-president had predicted, profits increased, the workforce was increased by 10 %, and many of the new workers would have been unable to find other work.

Fire/Main Effect

The vice-president of a manufacturing company was talking with the CEO. The vice-president said, "We are thinking of implementing a new policy. If we implement the policy, it will increase profits for our corporation. It will also mean that we have to fire 10 % of our employees, many of whom will have difficulty finding other work." The CEO said, "I've been looking for ways to fire some of our employees, and of course I always want to increase profits. Let's implement the policy."

They implemented the policy. Just as the vice-president had predicted, profits increased, 10 % of the employees were fired, and many of them were unable to find other work.

Hire/Main Effect

The vice-president of a manufacturing company was talking with the CEO. The vice-president said, "We are thinking of implementing a new policy. If we implement the policy, it will increase profits for our corporation. It will also mean that we have to increase our workforce by 10 %, hiring many people who would find it difficult to find other work." The CEO said, "I've been looking for ways to hire some more employees, and of course I always want to increase profits. Let's implement the policy."

They implemented the policy. Just as the vice-president had predicted, profits increased, the workforce was increased by 10 %, and many of the new workers would have been unable to find other work.

References

- Alfano, M. (2013). *Character as moral fiction*. Cambridge: Cambridge University Press.
- Alfano, M., Beebe, J., & Robinson, B. (2012). The centrality of belief and reflection in Knobe effect cases: A unified account of the data. *The Monist*, 95(2), 246–289.
- Beebe, J. R., & Buckwalter, W. (2010). The epistemic side-effect effect. *Mind & Language*, 25(4), 474–498.
- Beebe, J. R., & Jensen, M. (2012). Surprising connections between knowledge and action: The robustness of the epistemic side-effect effect. *Philosophical Psychology*, 25(5), 689–715.
- Bollen, K. A., & Stine, R. (1990). Direct and indirect effects: Classical and bootstrap estimates of variability. *Sociological Methodology*, 20, 115–140.
- Cova, F., & Naar, H. (2012). Side-effect effect without side-effects: The pervasive impact of moral considerations on judgments of intentionality. *Philosophical Psychology*, 25(6), 837–854.
- Knobe, J. (2003). Intentional action and side-effects in ordinary language. *Analysis*, 63(3), 190–193.
- Knobe, J. (2004). Folk psychology and folk morality: Response to critics. *Journal of Theoretical and Philosophical Psychology*, 24(2), 270–279.
- Knobe, J. (2007). Reason explanation in folk psychology. *Midwest Studies in Philosophy*, 31(1), 90–107.
- Knobe, J., & Mendlow, G. (2004). The good, the bad and the blameworthy: Understanding the role of evaluative reasoning in folk psychology. *Journal of Theoretical and Philosophical Psychology*, 24(2), 252–258.
- MacKinnon, D. P., Warsi, G., & Dwyer, J. H. (1995). A simulation study of mediated effect measures. *Multivariate Behavioral Research*, 30(1), 41–62.
- Nadelhoffer, T. (2004). On praise, side-effects, and folk ascriptions of intentionality. *Journal of Theoretical and Philosophical Psychology*, 24(2), 196–213.
- Pettit, D., & Knobe, J. (2009). The pervasive impact of moral judgment. *Mind & Language*, 24(5), 586–604.
- Sripada, C. S. (2011). What makes a manipulated agent unfree? *Philosophy and Phenomenological Research*. doi:10.1111/j.1933-1592.2011.00527.x.
- Swint, B. (2010, November 2). BP profit drops after taking further charge on gulf spill. *Bloomberg*. Retrieved from <http://www.bloomberg.com/>.
- Thomson, J. J. (1971). A defense of abortion. *Philosophy & Public Affairs*, 1(1), 47–66.
- Wible, A. (2009). Knobe, side-effects, and the morally good business. *Journal of Business Ethics*, 85(Supplement 1), 173–178.
- Williams, B. (2010). *Nightly news with Brian Williams*. Retrieved from <http://www.msnbc.msn.com/id/3032619/#39858528>.
- Zagzebski, L. (1996). *Virtues of mind*. Cambridge: Cambridge University Press.
- Zhang, Z., & Wang, L. (2009). BMEM: bootstrap mediation analysis using EM algorithm, Version 4.0. Retrieved from <http://bmem.psychstat.org>.