

Exploring the Factors That Encourage the Illusions of Control

The Case of Preventive Illusions

Fernando Blanco and Helena Matute

Departamento de Fundamentos y Métodos de la Psicología, Universidad de Deusto, Bilbao, Spain

Abstract. Most previous research on illusions of control focused on generative scenarios, in which participants' actions aim to produce a desired outcome. By contrast, the illusions that may appear in preventive scenarios, in which actions aim to prevent an undesired outcome before it occurs, are less known. In this experiment, we studied two variables that modulate generative illusions of control, the probability with which the action takes place, $P(A)$, and the probability of the outcome, $P(O)$, in two different scenarios: generative and preventive. We found that $P(O)$ affects the illusion in symmetrical, opposite directions in each scenario, while $P(A)$ is positively related to the magnitude of the illusion. Our conclusion is that, in what concerns the illusions of control, the occurrence of a desired outcome is equivalent to the nonoccurrence of an undesired outcome, which explains why the $P(O)$ effect is reversed depending on the scenario.

Keywords: illusion of control, contingency judgments

In a broad sense, behavior is the means by which animals try to cause relevant changes in their environments. Every action performed by an animal is aimed at either producing a desired outcome (e.g., as when foraging to gather food, or taking a painkiller to stop a headache) or preventing an undesired one (e.g., as when spraying insect repellents to keep mosquitoes away).¹ However, needless to say, certain actions will fail to yield the outcome they aimed at. For instance, using a lucky charm will hardly affect tomorrow's weather in any meaningful way. Moreover, even when the potential of a given action to produce an outcome does actually exist, it may eventually wear off (e.g., flipping a light switch would normally work to switch a light bulb on, unless the electricity supply is interrupted by a storm). The existence of ineffective actions and, even more importantly, the fact that the potential of an action to produce an effect may change due to external factors force animals to engage in a constant, dynamical coupling between their behavior and the environment, with the goal of persisting in the effective actions that yield the desired outcome, and quitting on the useless actions that fail to produce it.

The first issue an animal must address to successfully adjust its behavior to the environment comprises the task of telling the difference between effective and ineffective

actions. According to a vast literature, this question is closely related to contingency learning. Normatively, the degree of contingency between two events is given by the ΔP rule (Allan, 1980), which amounts to the difference between two conditional probabilities: Namely, the probability of the outcome if the action is performed, $P(O|A)$, minus the probability of the outcome if the action is not performed, $P(O|\neg A)$. To the extent to which these two conditional probabilities differ, ΔP departs from zero, and consequently the two events, action and outcome, are contingent on each other. A positive contingency (i.e., $\Delta P > 0$) means that the probability of the outcome occurrence (whether desired or undesired) is higher when the action is performed than when it is not. This would happen when the action effectively causes (i.e., generates) the outcome. A negative contingency (i.e., $\Delta P < 0$) reflects the opposite situation, typical of a preventive scenario (i.e., the action prevents the occurrence of the outcome, which could either be desired or undesired). Whenever ΔP equals zero, the null contingency suggests that the action fails to affect the state of the outcome, just as the weather changes are not contingent on using any lucky charms.

A considerable amount of evidence suggests that people and other animals are able to match some aspects of their

¹ There are also actions that produce undesired outcomes (e.g., touching a hot stove and burn oneself) or prevent the occurrence of a desired outcome (e.g., insisting too much on asking for a date to a potential couple, and being rejected). In these cases, people (and animals, more generally) will presumably learn to stop performing such harmful actions. They are not the focus of the present paper (for further discussion on the illusion of control in such scenarios, see Matute & Blanco, 2014).

behavior to the actual contingency between their actions and relevant outcomes (Blanco, Matute, & Vadillo, 2010; Rescorla, 1968; Shanks & Dickinson, 1987; Wasserman, 1990). However, it has been shown that systematic errors appear under certain circumstances. Of special interest are those situations in which there is no contingency between action and outcome (i.e., $\Delta P = 0$), but still people believe that their actions affect the outcome. We call these phenomena illusions of control (see Langer, 1975). Two conditions that are known to facilitate the illusion of control are a high probability of the outcome occurrence, $P(O)$, and a high probability of performing the action, $P(A)$. That is, even when the actual action-outcome contingency is null, exposing participants to very frequent outcome occurrences, that is, high $P(O)$, often leads to an illusion of control (this is sometimes called outcome-density bias; Allan & Jenkins, 1983; Alloy & Abramson, 1979; Buehner, Cheng, & Clifford, 2003; Musca, Vadillo, Blanco, & Matute, 2010; Vallée-Tourangeau, Murphy, & Baker, 2005; Wasserman, Kao, Van Hamme, Katagiri, & Young, 1996). Likewise, participants who perform the action very often, that is, high $P(A)$, are likely to develop an illusion of control, particularly in high $P(O)$ conditions (Allan & Jenkins, 1983; Blanco, Matute, & Vadillo, 2011, 2012; Hannah & Beneteau, 2009; Matute, 1996).

As the mentioned research illustrates, illusions of control in which people believe in their actions as effective generative causes of desired events have been extensively studied, and some variables that modulate the illusion have been identified. But we know less about the illusions that may arise in the preventive scenario, when actions are thought to prevent the occurrence of undesired, but still uncontrollable, outcomes (Bloom, Venard, Harden, & Seetharaman, 2007).

Studying preventive illusions is relevant for various reasons. First, the available evidence suggests that the occurrence of undesired outcomes is able to induce even stronger illusions than the occurrence of desired ones (e.g., Aeschleman, Rosen, & Williams, 2003; Bloom et al., 2007). Second, most everyday superstitions involve the avoidance of undesired outcomes, such as omens (e.g., knocking on wood to avoid bad luck), rather than obtaining desired outcomes (e.g., using an amulet to bring about good luck) (Blum & Blum, 1974; Wiseman & Watt, 2004). Most importantly, the preventive illusion turns out to be of interest for clinical and abnormal psychology: Whereas generative illusions have been associated to optimistic and adaptive biases (Taylor & Brown, 1988), preventive illusions could be instead related to maladaptive behaviors. For instance, patients suffering from Obsessive-Compulsive Disorder (OCD) are known to develop strange superstitious habits that often resemble preventive illusions of control (e.g., Dettore, & O'Connor, 2013; Reuven-Magril, Dar, & Liberman, 2008; Zebb & Moore, 2003). Typically, the patient engages in a repetitive behavioral pattern aimed at preventing catastrophic events from happening, despite this behavior being evidently useless. Note that, in this pathological condition, the action happens *before* the undesired outcome takes place (i.e., with the goal to prevent it). Furthermore, the action can happen even

despite the aversive outcome never taking place. Thus, a patient suffering from OCD would typically act in a certain systematic way that is only supported by the thought that a terrible outcome (such as the death of a close relative) would occur if the patient stopped engaging in that particular behavior, despite the fact that the fearsome outcome never actually happened. It is the absence of the outcome that reinforces the action, and this is an important difference as compared to the more thoroughly studied generative scenario, where a positive reinforcement schedule takes place (i.e., the occurrence of a desired outcome after emitting a response increases the likelihood of further responding). A parallel to the situation of OCD patients can be found in Aeschleman et al. (2003, Experiment 2, negative reinforcement condition), in which no outcome was presented and still a reliable illusion of control was reported.

To sum up, the illusions of control have been thoroughly studied in the case of generative scenarios in which a certain behavior is repeatedly followed (and reinforced) by a desired outcome that occurs frequently (being it the occurrence of an appetitive stimulus or the termination of an aversive one). However, this rationale cannot be easily applied to the preventive scenario in which an undesired (to-be-avoided) outcome never (or hardly ever) takes place. In fact, we could make the prediction that the manipulations of $P(O)$ would actually have the opposite effect when conducted in a preventive scenario as compared to when they are conducted in a generative scenario. In line with this prediction, OCD patients display apparently strong preventive illusions, yet the $P(O)$ they are exposed to is extremely low (i.e., the fearsome outcome they are trying to avoid is unlikely to ever occur). The $P(A)$, on the other hand, is expected to increase the magnitude of the illusion, regardless of whether the outcome is desired or undesired. The reason is that the more often one acts, the greater the chances that the action is accidentally followed by the desired outcome (or the absence of the undesired outcome), thus strengthening the association between both. The effect of $P(A)$ was not studied in previous experiments that used undesired outcomes (e.g., Aeschleman et al., 2003).

The goal of the current experiment is to explore the effects of $P(A)$ and $P(O)$, two variables that affect generative illusions (i.e., produce/terminate an outcome), on the magnitude of illusions developed in a preventive scenario. Our prediction is that, in the preventive condition, the manipulation of $P(O)$ should yield an effect of opposite sign to that typically observed in the generative condition, whereas $P(A)$ should affect the illusion in the same way as in the generative scenario: That is, a higher $P(A)$ should produce stronger illusions of any sign.

Method

Participants and Apparatus

Seventy-nine first year students from the University of Leuven took part in the study, in exchange for course credit.




Figure 1. Diagram showing the flow of events in each trial. First, the inter-trial interval (ITI), represented as Event 1 (left), consisted of showing the light bulb off for 2 s before the trial started. Event 2 (middle section) consisted of the response button appearing below the light bulb. During Event 2, participants had the opportunity to press the spacebar or not. Finally, during Event 3 (right) one of two things could happen: either the light bulb went on (3a) or stayed off (3b). In either case, the event lasted for 2 s, after which a new ITI started.

We excluded the data of two participants because they responded on every trial. Therefore, they did not expose themselves to the outcome base-rate information, $P(O|A)$, which is essential to give an accurate judgment of control. In this particular experiment in which the outcome occurrence has two different interpretations depending on the condition, including these participants would have been a matter of concern because there were no participants who failed to respond in every trial to compensate. The final sample consisted of 77 students: 21 in the Produce-High group; 19 in the Produce-Low group; 19 in the Prevent-High group; and 18 in the Prevent-Low group.

The experimental task was programmed in E-Prime for Windows. Participants were tested individually in a sound-proof room at the Department of Psychology of the University of Leuven. The experimental session took approximately 15 min.

Procedure

The “light bulb task” described by Msetfi, Murphy, Simpson, and Kornbrot (2005) was adapted for this experiment. Before the session, participants were given written instructions on the computer screen (they are available in the Appendix). The sequence of events in every trial is depicted in Figure 1. The inter-trial interval (ITI), depicted as Event 1 in Figure 1, lasted 2 s, during which the

participants were presented with the picture of a light bulb that was off against a white background in the computer screen. During this interval that aimed to separate consecutive trials, participants could only wait until the next trial started. All trials had a fixed duration of 4 s, comprising a sequence of two events, each one lasting for 2 s. First, a red button appeared below the bulb, together with a text-box stating “*You may press the button now*” (i.e., Event 2 in Figure 1). This indicated to the participants that they had the opportunity to do something: If they wanted to press the button, they should press the spacebar on the keyboard immediately; if they decided not to press the button, they should do nothing and wait. The button remained available on the screen for 2 s before it disappeared. Any response given while the button was not present on the screen was not recorded. Subsequently, the light bulb either came on for 2 s (i.e., outcome-present trial) or, alternatively, stayed off for 2 s (i.e., outcome-absent trial). This corresponds to events 3a and 3b in Figure 1. Then, the ITI was presented again, followed by a new trial.

After a series of 50 trials, the participants were asked to rate how much control they exerted over the light (i.e., a judgment of control: “*To what extent did you control the switching on of the light bulb?*”) by clicking on a scale ranging from -100 to $+100$. The use of the scale was described as follows:

“Please, answer by **CLICKING THE MOUSE** on the scale. -100 means: Pressing the button **ALWAYS PREVENTED** the light bulb from switching on; 0 means: Pressing the button **HAD NO EFFECT AT ALL** on the switching on of the bulb; $+100$ means: Pressing the button **ALWAYS MADE** the light bulb switch on; Intermediate numbers mean **INTERMEDIATE LEVELS OF CONTROL**, either to prevent the light from switching on (negative values), or to make it switch on (positive values).”

Two between-participants manipulations were conducted. The first one concerned the participants’ goal. Before the experiment, two different sets of instructions specified different goals for the participants (see Appendix). In the “Produce” condition, the light coming on was described as the desired outcome that participants should try to produce. By contrast, in the “Prevent” condition, participants were instructed to try to prevent the light from coming on. That is, by means of instructional manipulations, the very same event (i.e., the switching on of the bulb) was a desired outcome for half of the participants, and an undesired outcome for the other half.

The second manipulation was the probability of the light coming on, $P(O)$. There are several ways in which $P(O)$ can be manipulated in this type of experiment. The way in which we programmed it was as follows: At the beginning of the experiment, the computer program generated a list of 50 items. A number of them (either 10 or 40, depending on the condition) were labeled as “outcome,” and the rest as “no outcome.” In each trial, the program chose a random item from the list without

Table 1. Descriptive statistics of the raw judgments of control given in each group. LL and UL are the lower and upper limits of the 95% confidence intervals for the mean. Note that the zero value lies within the confidence interval in the Produce-Low and Prevent-High groups, suggesting that they developed little or no illusion of control. By contrast, judgments were positive in the Produce-High group, and negative in the Prevent-Low group, suggesting a generative and a preventive illusion, respectively

Group	Mean	LL	UL	SD
Produce-High	25.24	7.98	42.49	40.34
Produce-Low	-0.84	-10.84	9.15	22.23
Prevent-High	9.89	-2.35	22.14	27.24
Prevent-Low	-34.11	-51.88	-16.34	38.46

replacement, and used it to determine whether or not the light would come on in that particular trial. In the High P(O) condition, the light came on in 40 out of 50 trials, that is, $P(O) = .80$, whereas in the Low P(O) condition, it came on in 10 out of 50 trials, that is, $P(O) = .20$. Since the sequences of trials were randomly determined by the program regardless of the participant's decisions, the outcome (i.e., the light coming on) was uncontrollable.

Results

Judgments of Control

Table 1 reports the descriptive statistics for the judgments of control obtained in each group. Additionally, the information is complemented by the histograms in Figure 2. As the figure suggests, there was some variability in the individual participants' judgments. Several participants, in all groups, realized that their actions were useless and gave a judgment of zero (eight in the Produce-High group, nine in the Produce-Low group, seven in the Prevent-High group, and six in the Prevent-Low group). This is not uncommon in experiments where the programmed contingency between the action and the outcome is null. Nonetheless, not all participants became aware that the schedule was not contingent, and what matters for our present purposes is how the distributions of these judgments varied between groups. In the groups Produce-Low and Prevent-High, in which either the desired outcome was scarce or the undesired one was frequent, the participants' judgments did not differ significantly from zero, as expected: $t(18) = 0.16$, $p = .87$, and $t(18) = 1.58$, $p = .13$ (respectively). Moreover, in the two remaining groups, the judgments were either significantly above zero, $t(20) = 2.87$, $p = .01$ (Produce-High group) or significantly below zero, $t(17) = 3.97$, $p = .001$ (Prevent-Low group). These two groups are the ones that provide the most rewarding

experience to the participant; hence they promote the illusion (be it generative or preventive).

Given that the judgment scale was bidirectional (-100 to $+100$), it could mask the illusions by compensating positive and negative departures from zero (the normatively correct value of control exerted over the outcome). Therefore, we conducted the following analysis taking the absolute values of the judgments as dependent variable (see Figure 3). A 2×2 ANOVA with P(O) (high vs. low) and Goal (produce vs. prevent) as factors revealed a $P(O) \times \text{Goal}$ interaction, $F(1, 73) = 8.45$, $p = .005$, $\eta_p^2 = .10$, whereas neither the main effect of P(O) nor of Goal were significant (both $F_s < 1$). This means that, as suggested by Figure 3, illusions of control appeared more prominently either when the outcome was desired and frequent or when it was undesired and scarce. As Table 1 indicated, the illusion was generative (positive judgments) in the former case and preventive (negative judgments) in the latter.

Effect of the Probability of the Action on the Judgments of Control

The probability of the action, P(A), was calculated as the number of trials in which the participant decided to press the button over the total number of trials (i.e., 50). We were mainly interested in the effect of the P(A) on judgments of control, and how this effect could be modulated by Goal and P(O). The scatter plots in Figure 4 depict the raw judgments as a function of P(A) for each group. The figure suggests that the slope of the linear relation between P(A) and judgments was highly dependent on the group. Thus, we regressed the raw judgments onto the P(A) within each of the four groups. The simple linear regression analysis yielded a significant positive relation in the Produce-High group, $\beta = .59$, $t(19) = 3.18$, $p = .005$, $R^2 = .35$. This is consistent with previous reports conducted in similar conditions (e.g., Blanco et al., 2011; Matute, 1996), in which the higher the P(A) was, the more positive the judgment of control (i.e., stronger generative illusion). In addition, we found a significant negative relation in the Prevent-Low group, $\beta = -.52$, $t(17) = 2.44$, $p = .027$, $R^2 = .27$. As we predicted in this condition, higher P(A) led to more negative judgments of control (i.e., stronger preventive illusion). Contrasting with these two groups, the slopes were not significantly different from zero in the Produce-Low group, $\beta = -.23$, $t(18) = 0.99$, $p = .33$, $R^2 = .05$, and in the Prevent-High group, $\beta = .23$, $t(18) = 0.96$, $p = .35$, $R^2 = .05$ (Prevent-High group), as could be expected.

Together with the previous analyses, these data indicate that two latent types of group underlie our design. On the one hand, the two groups in which either the desired outcome occurred frequently or the undesired outcome occurred scarcely (Produce-High and Prevent-Low) reproduce similar highly rewarding situations. On the other hand, the two groups in which either the desired outcome occurred seldom or the undesired outcome occurred very often (Produce-Low and Prevent-High) result in a situation

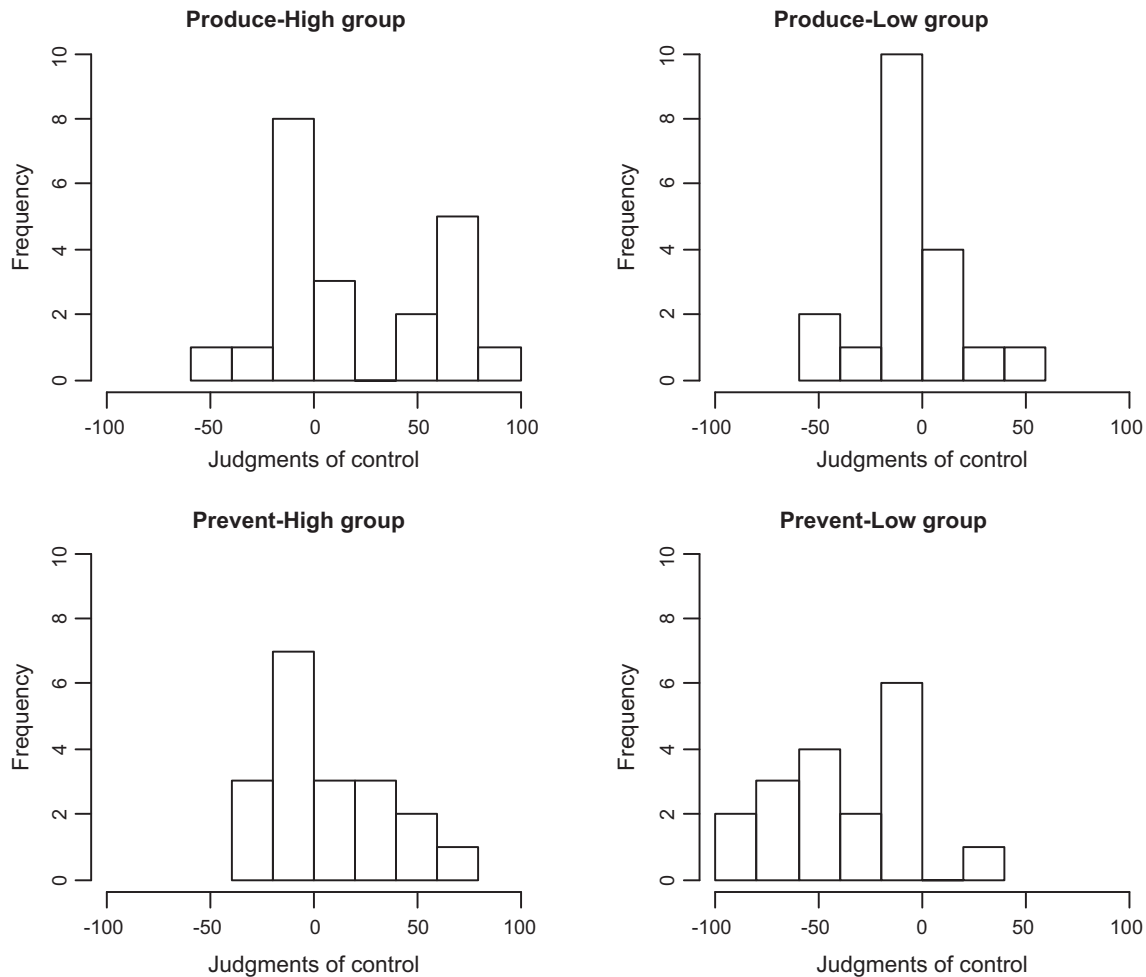


Figure 2. Histograms depicting the distributions of the raw judgments in each group. As can be seen in the figure, many participants gave a judgment of zero (this was the modal value in all groups). However, the distribution of the rest of the judgments varied between groups. It is sensible that the mode of the distributions was zero for at least two reasons: First, it was the normatively correct value. Consequently, many participants might have realized that the outcome was noncontingent on the actions. Second, whereas there is only one value to choose when the participant realized that there was no control over the outcome (zero), there were plenty of other values to choose when the participant developed an illusion, therefore resulting in a spread distribution instead of a narrow one.

where failure, not reward, prevails. Accordingly, the two former groups showed a stronger illusion of control than the two latter ones, as detailed above.

Additional Analyses

Although the main results of the experiment have been reported in the previous section, we now provide additional information concerning two potential confounds of instrumental tasks in which the decision of whether or not to act is left to the participant. First, we needed to check that all groups were comparable in their $P(A)$ level. Thus, we conducted a 2×2 ANOVA with $P(O)$ and Goal as factors on the $P(A)$. As Figure 5 suggests, neither the interaction, $F(1, 73) = 3.23$, $p = .077$, $\eta_p^2 = .04$, nor any of the main effects (minimum $p = .20$) was found significant.

Therefore, we can conclude that there were no significant differences in $P(A)$ between groups. This renders the effects of $P(A)$ on the judgments of control that we reported in the previous section easier to interpret: They could not be attributable to the groups differing in their levels of $P(A)$. In fact, there were reasons to expect our participants to press the button more often in those groups where they were more frequently rewarded, either because the desired outcome occurred often (i.e., Produce-High group), or because the undesired outcome was absent in most of the trials (i.e., Prevent-Low group). On the other hand, according to the instructions given to participants (see Appendix), the outcome could reinforce either the action or the absence of the action. If a participant decided not to act on a given trial and the desired event occurred, then refraining from acting was probably reinforced. This may explain why we found no between-groups differences in $P(A)$.




Figure 3. Mean absolute values for the judgments of control in the four experimental groups. Illusions of control (i.e., departures from zero) appeared prominently in the Produce-High and the Prevent-Low groups. In addition, note that, as Table 1 indicates, the illusion was positive in the Produce-High group and negative in the Prevent-Low group. Whiskers depict 95% confidence intervals for the means.

The second potential confound in these procedures is the actual contingency. Even when the programmed contingency is set to zero, as in our experiment, participants may end up exposing themselves to slightly different levels of actual contingency (Hannah, Allan, & Siegel, 2007). Therefore, we proceeded to analyze the actual contingency to which our participants exposed themselves during the session. Figure 6 includes the actual contingency tables experienced by each group, as well as the actual contingency values (measured as the ΔP index) computed from the total of 50 trials. A 2×2 ANOVA with P(O) and Goal as factors was conducted on the actual contingency values. The analysis yielded no significant results: For the main effect of P(O) and the interaction, both F s were < 1 ; and the main effect of Goal did not reach the significance threshold, $F(1, 73) = 2.67$, $p = .11$, $\eta_p^2 = .03$. That is, no significant differences between the four groups were observed in the actual contingency values. Consequently, it seems unlikely that this variable can explain the between-groups differences in the judgments that we reported in the previous section. Very similar results were obtained when we repeated the analyses on the actual contingency computed from the last block of ten trials, which indicates the consistency of the result.

Discussion

Although animals strive to successfully discriminate between those outcomes that are under their control and

those that remain uncontrollable, sometimes their sense of control is misled by situational variables. Three of these variables are of interest for the current paper: The probability of the outcome, P(O), the probability of the action, P(A), and the causal polarity of the scenario (preventive vs. generative). The effect of P(O) has been reported as the observation that delivering desired outcomes with high frequency creates the illusion that they are under the participant's control (Allan & Jenkins, 1983; Alloy & Abramson, 1979; Buehner et al., 2003; Musca et al., 2010; Vallée-Tourangeau et al., 2005; Wasserman et al., 1996). Likewise, the effect of P(A) entails a similar illusion when the participant's action is performed very frequently (Blanco et al., 2011, 2012; Hannah & Beneteau, 2009; Matute, 1996). These two effects have been thoroughly studied in generative scenarios where participants attempt to produce a desired outcome, and not yet so in preventive scenarios where they try to prevent the occurrence of an undesired one. Thus, the participant's goal is a third factor of crucial interest for this research. If we re-interpret the nonoccurrence of an undesired outcome (e.g., mosquitoes not appearing after spraying a repellent) as an actually appetitive outcome, then we could hypothesize that the effect of P(O) should be present in preventive scenarios too, but only its sign would be reversed. In addition, we could also hypothesize that the effect of P(A) should be similar in preventive and generative scenarios. Therefore, the conjunction of low P(O) and high P(A) should lead to stronger preventive illusions.

This is precisely what we found in our experiment. First, the illusion of control over an actually uncontrollable outcome (a light onset) was sensitive to P(O), but dependent on the participant's goal (to prevent the occurrence of an undesired outcome vs. to produce the occurrence of a desired one). We found a strong illusion of positive sign in the generative scenario as long as the desired outcome was frequent (in line with many previous reports). The negative illusion (control judgments below zero) appeared in the preventive condition, but only when the undesired outcome was delivered with low probability. We found no illusion in the other two conditions (i.e., desired outcome occurring seldom and undesired outcome occurring often). Second, P(A) also affected the control judgments, but its effect depended on both P(O) and the participant's goal. In the generative scenario with high P(O), P(A) showed a positive relationship with the judgments (i.e., the more often the participant performed the action, the stronger the generative illusion, as found in many previous reports). In the completely opposite situation, with an undesired, to-be-prevented outcome occurring with low P(O), the relationship between P(A) and judgments was negative (i.e., the more frequent the action, the stronger the preventive illusion). Thus, high values of P(A) always facilitated the illusion in these two groups, but the sign of the illusion depended on the participant's goal. By contrast, P(A) did not significantly predict the judgments in the remaining two conditions (i.e., desired outcomes occurring seldom, and undesired outcomes occurring frequently), suggesting that the P(A) effect is subject to the high P(O) level, as recent evidence has pointed out (Blanco, Matute, &

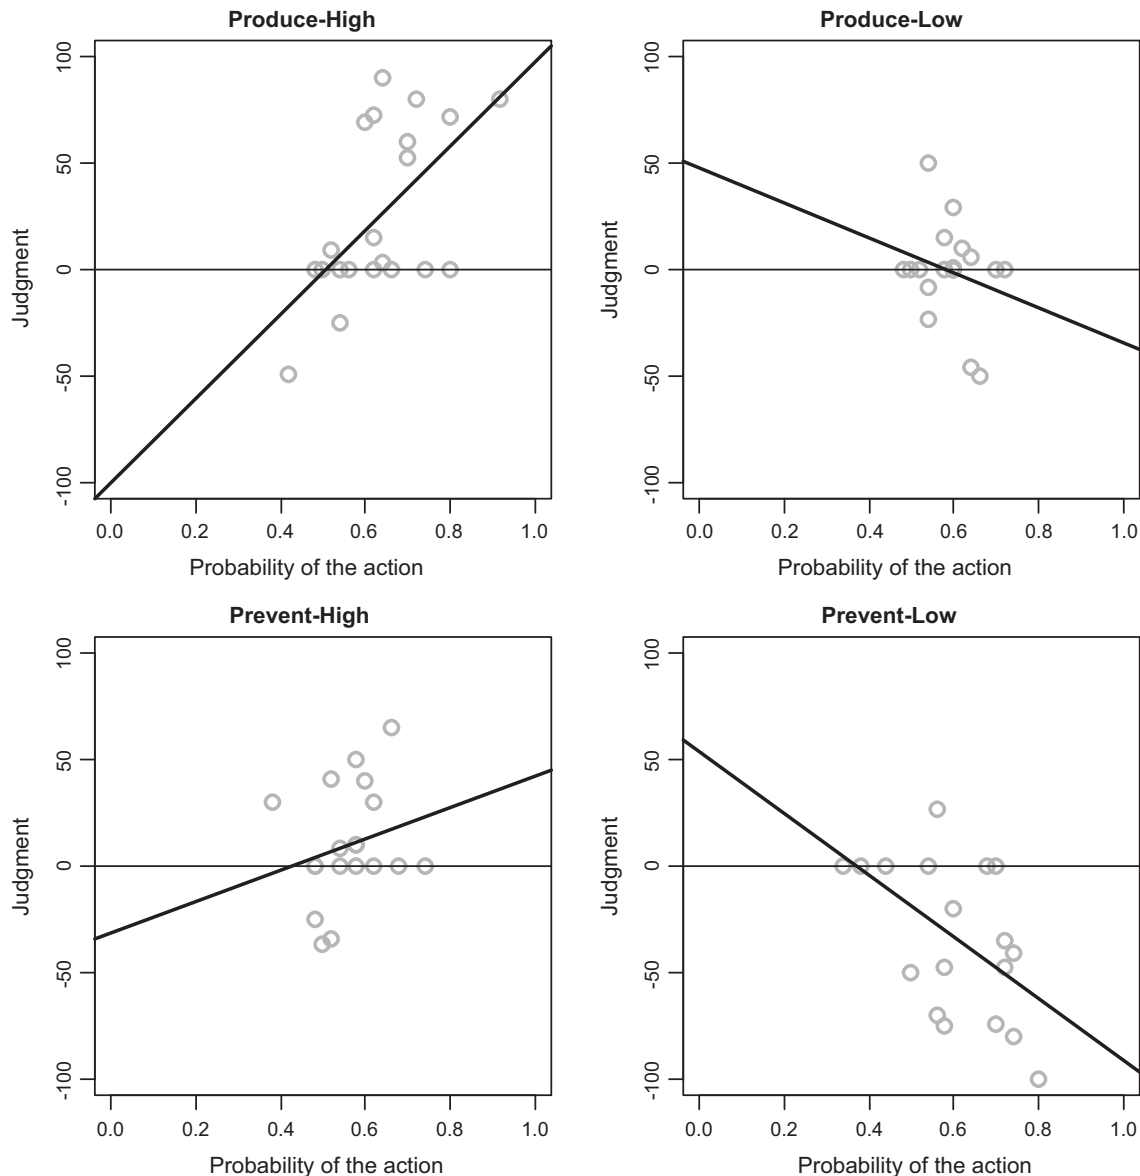


Figure 4. Scatter plots depicting the participants' judgments (vertical axes) as a function of their $P(A)$ (horizontal axes), by group. Simple regression lines are fitted to the data points. Only the slopes in groups Produce-High and Prevent-Low were found significantly different from zero: positive in the former case, and negative in the latter one (see main text).

Vadillo, 2013, Experiment 1; the present report extends this finding to the preventive scenario and to an instrumental learning paradigm).

Overall, we found that the Produce-High and Prevent-Low groups, on the one hand, and the Produce-Low and Prevent-High groups, on the other hand, exhibited similar patterns of results. In other words, we had two grand types of situation in this experiment: First, we had two groups in which participants were frequently rewarded, and then another two groups in which they failed to obtain the goal they were pursuing. This may be taken as evidence that the absence of undesired outcomes can be safely treated as equivalent to the occurrence of desired outcomes. At least, both events yield the same illusion (in terms of magnitude,

not sign) concerning the manipulations and dependent variables used in this study. It is interesting that this happened in an experimental setting involving buttons and light bulbs, in which people would be, presumably, more familiar with the generative scenario. Indeed, in everyday life, most causal relations between buttons and light bulbs are generative. This suggests that participants' judgments were guided by the experimental instructions and contingencies experienced during the task, and not solely by their previous interactions with similar situations.

We admit that the conclusions given above must be interpreted in the appropriate context: As the histograms indicate (Figure 2), many participants realized that the contingency they were being exposed to was null, and

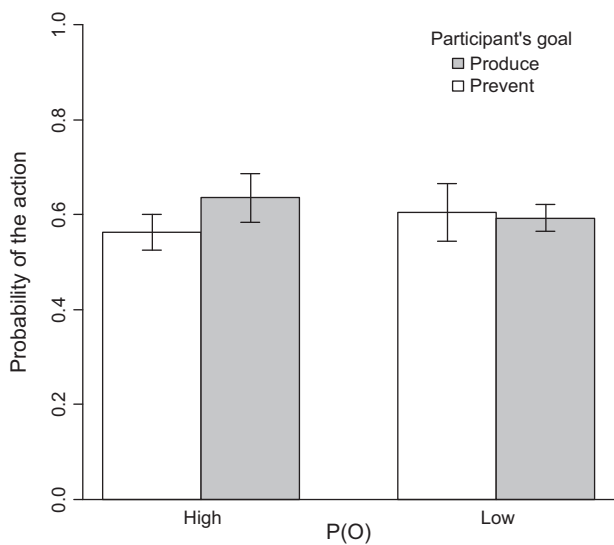


Figure 5. Mean probability of the action, P(A), in the four experimental groups (computed for the whole training phase). Whiskers depict 95% confidence intervals for the means.

consequently gave a judgment of zero, while others showed judgments indicative of an illusion of control (either generative or preventive). That is, there was some degree of inter-individual variability. This is not an uncommon finding in experiments where judgments of control are collected in null contingency settings: In fact, this variability allowed us to detect the effect of P(A) on judgments in one-group

designs (e.g., Blanco et al., 2011). Notably, in the current experiment, the amount of participants who gave a judgment of zero was fairly similar in all groups (as mentioned in the Results section). Thus, the significant differences between groups in the judgments were more likely due to the participants who showed a degree of illusion in either direction (generative or preventive).

As we argued elsewhere (Matute, Vadillo, Blanco, & Musca, 2007), associative theories are good candidates to model illusions of control and related phenomena. According to the influential Rescorla-Wagner model (Rescorla & Wagner, 1972), an individual's judgment of control would be given by the strength of the association between the representation of the action and the representation of the outcome, V_A . This associative strength is updated every time the action is performed according to the following equation:

$$\Delta V_A = \alpha_A \beta (\lambda - V_{TOTAL}) \tag{1}$$

In Equation 1, ΔV_A is the change in the associative strength of the action in the current trial; λ represents the asymptote of learning possible with the outcome; V_{TOTAL} is the sum of the associative strength of the action, V_A , and the associative strength of a constant background stimulus, the context, V_{Ck} . Then, the difference $(\lambda - V_{TOTAL})$ represents the surprisingness of the outcome: Once an organism has learnt to predict the outcome from other stimuli, including their own actions, the outcome occurrence will not be surprising and little additional learning will occur. Additionally, there are two learning rate parameters, α and β , which represent the saliences of the events. When the associative strength of the action is being updated, α_A is used. When the associative strength of the context is being

Produce-High

	Outcome	¬Outcome
Action	23.86 (7.60)	7.57 (4.46)
¬Action	13.29 (5.92)	5.29 (5.45)
		$\Delta P = 0.017$ (0.15)

Produce-Low

	Outcome	¬Outcome
Action	6.84 (3.25)	22.47 (4.92)
¬Action	4.73 (4.17)	15.94 (3.66)
		$\Delta P = 0.019$ (0.097)

Prevent-High

	Outcome	¬Outcome
Action	21.68 (4.85)	6.63 (4.14)
¬Action	16.74 (4.32)	4.95 (3.64)
		$\Delta P = -0.01$ (0.101)

Prevent-Low

	Outcome	¬Outcome
Action	15.34 (9.69)	14.63 (9.69)
¬Action	10.25 (6.86)	9.68 (6.99)
		$\Delta P = -0.005$ (0.129)

Figure 6. Actual contingency matrixes averaged for each group. The numbers inside each cell correspond to the mean number of occurrences of each type of trial and the mean actual contingency as computed by the ΔP index. Standard deviations are provided between brackets.




Figure 7. Simulations of the four experimental groups using the Rescorla-Wagner model (Rescorla & Wagner, 1972). One simulation was run for each of our participants, using the same trial sequences that they generated, and then they were averaged per group and trial to produce the lines in this figure. The parameter values were taken from a previous study (as detailed in the top-left corner). Note that the asymptotic pattern observed here at the end of the training resembles the results of the experiment as presented in Figure 3.

updated, an analogous parameter α_{Ctx} is used. Finally, β_{O} is the salience of the outcome in those trials in which the outcome occurs, while $\beta_{-\text{O}}$ represents the salience of the outcome absence in those trials in which the outcome is absent.

We used the original Rescorla-Wagner model to simulate our data with the trial sequences produced by each participant. The parameter values were taken from a previous publication in which the effects of $P(\text{O})$ and $P(\text{A})$ on control estimations were successfully reproduced (Matute et al., 2007): $\alpha_{\text{A}} = 0.6$, $\alpha_{\text{Ctx}} = 0.2$, $\beta_{\text{O}} = 0.5$, and $\beta_{-\text{O}} = 0.5$. It is a usual assumption that the salience of the action (or the target stimulus) is greater than that of the context, hence $\alpha_{\text{A}} > \alpha_{\text{Ctx}}$. To conduct the simulation, we assumed that λ equals 1 whenever the desired event occurs (i.e., the light coming on in the Produce condition, and the light staying off in the Prevent condition), and it equals 0 otherwise. Figure 7 depicts the results of the simulation with the Rescorla-Wagner model. Not surprisingly, the ordinal pattern of results at the end of the simulation is almost identical to the one we obtained with the judgments of control at the end of our experiment (see Figure 3). It must be noted, however, that the choice of another set of parameters could lead to different predictions.

It is interesting to mention at least yet another model, namely Cheng's (1997) power PC theory. Contrary to ΔP ,

Table 2. Descriptive statistics of the power PC index computed for each participant, and then aggregated by group. The generative-cause version of the power PC index was used in the Produce condition, and the preventive-cause version was used in the Prevent condition. LL and UL are the lower and upper limits of the 95% confidence intervals for the mean

Group	Mean	LL	UL	SD
Produce-High	-0.328	-0.903	0.248	1.313
Produce-Low	0.012	-0.044	0.067	0.123
Prevent-High	-0.006	-0.074	0.062	0.151
Prevent-Low	0.028	-0.267	0.324	0.640

power PC is formulated as an index to assess causality rather than covariation. To this end, it isolates the causal strength of the action from any other potential cause operating in the background, which implies taking into account the base rate of the outcome occurrence. One of the features of the power PC model is that the causal polarity (generative and preventive) plays an important role, as the integration rule embedded in the index is different for generative and for preventive causes. It has been a matter of debate how participants would choose the appropriate rule before any information has been collected (Lober & Shanks, 2000). In our experiment, the expected polarity was made clear to the participant via instructions. Thus, the appropriate rule (for generative or for preventive causes) could be chosen from the beginning. Because of this, our experiment offers an opportunity to examine the predictions made by Cheng's model in generative and preventive noncontingent settings.

We computed the power PC index for each of our participants,² and then they were averaged for each group (Table 2). In principle, power PC predicts no deviations from zero in null contingency settings (like in our experiment). However, since participants were free to choose when to act, their actual contingencies departed slightly from zero (see Figure 6), hence we found some small variability in the power PC indexes of our sample. On the other hand, this variability was not attributable to the two between-group manipulations, $P(\text{O})$ and Goal (the 2×2 ANOVA on the power PC actual indexes yielded all F s < 1). Moreover, the pattern of predictions by the power PC model in Table 2 did not mirror the findings of the experiment: First, these predictions did not differ significantly from zero in any group (minimum $p = .28$). Second, the order of the groups according to their mean power PC predictions was Produce-High < Prevent-High < Produce-Low < Prevent-Low, which does not coincide with the order found in the participants' judgments. Finally, there were no significant differences between groups in the power PC predictions (minimum $p = .15$). Other research (Lober &

² Data from one participant in the Produce-High group were discarded from this analysis because it yielded an invalid value (i.e., the denominator was 0).

Shanks, 2000) has also found problems when using power PC to model actual effects of P(A) and P(O) effects.

Although one must bear in mind this lack of significant differences between groups, we can observe that the mean predictions made by power PC in our experiment were negative for the groups in which P(O) was high, and positive for the groups in which P(O) was low. Additionally, the Preventive condition produced slightly higher power PC predictions than the Produce condition. The ordinal pattern did not coincide with that yielded by ΔP (Figure 6) either, because the computation of power PC corrects for the base rate of the outcome.

We found a wide range in the power PC values computed for each participant in our sample (from -5.00 to 2.66). In active procedures like this, the actual conditional probabilities of the outcome, $P(O|A)$ and $P(O|\neg A)$, can vary due to chance and to the participants' decisions to act or not on each trial (Hannah et al., 2007). Note that in our experiment we used rather extreme values of P(O) (i.e., .80 and .20), which led some participants to be exposed to very high or very low values of the outcome base rate, $P(O|\neg A)$, even when their P(A) level was medium. When computing ΔP , these extreme values of $P(O|\neg A)$ were usually compensated by the similar values of $P(O|A)$, leading to actual ΔP values that were close to the programmed value, zero, as we have showed. However, the extreme values of the outcome base rate $P(O|\neg A)$ strongly affect the result of the computation of power PC. Just as an example: If the outcome base rate is very high in the noncontingent, generative scenario, then ΔP would be close to zero, whereas the absolute value of power PC would increase without limit. All this suggests that people do not use the type of normative causal induction that power PC describes, at least when judging their control over outcomes that are actually uncontrollable.

The current design might be used to model certain everyday situations. For instance, a football player who wears a lucky charm to bring about good luck and turns out to win most of the games could belong to the Produce-High group (see Figure 4, top-left panel). In this case, performing the superstitious behavior very often (e.g., wearing the amulet in every game) reinforces the illusion. Previous research points in this direction too (Blanco et al., 2011; Matute, 1996). On the other pole, we have OCD patients, engaging in eccentric behaviors to prevent terrible events that in fact never occur. Their situation is mirrored by the Prevent-Low group (see Figure 4, bottom-right panel). As we have shown in our experiment, if the behavior intended to prevent infrequent undesired outcomes takes place very often, then the preventive illusion appears, and it is seemingly as strong in magnitude as its generative counterpart. Thus, it is not surprising that many patients show a solid conviction in that their repetitive behavior is actually preventing some event that never took place. A promising, and testable, prediction derived from our research is that, just as reducing P(A) in the generative scenario attenuates the illusion, preventive illusions should similarly vanish as P(A) is reduced. Thus, OCD patients' problematic beliefs could also be diminished if they were able first

to reduce the frequency of their avoidance behaviors (as shown in Figure 4, bottom-right panel).

To sum up, we provide an exploration of how the illusion of control appears in a situation that has not been typically studied in the literature: Preventive (as opposite to generative) scenarios in which the outcome of interest should be prevented, rather than produced. We found that the illusion appeared in the preventive scenario, but only when the probability of the to-be-prevented outcome was low, which is the opposite of the usual finding in the generative scenario. Another factor that affects the illusion of control, the probability of the action, increased the illusion only in the groups where the situation was highly rewarding (either with frequent desired outcomes, or with infrequent undesired outcomes), although the signs of the illusions were reversed, as we expected (positive judgments in the former case, negative ones in the latter case). This suggests that, to some extent, the mistaken belief that one can prevent an uncontrollable outcome and the belief that one can produce it represent in fact similar illusions, but of opposite sign.

Acknowledgments

This research was supported by Dirección General de Investigación of the Spanish Government (Grant No. PSI2011-26965) and the Departamento de Educación, Universidades e Investigación of the Basque Government (Grant No. IT363-10). During the conduction of this experiment, Fernando Blanco was a postdoctoral researcher at the Center for Learning and Experimental Psychopathology, KU Leuven, Belgium. We thank Frank Baeyens, Dirk Hermans, and the CLEP team for their support and enlightening conversations.

References

- Aeschleman, S. R., Rosen, C. C., & Williams, M. R. (2003). The effect of non-contingent negative and positive reinforcement operations on the acquisition of superstitious behaviors. *Behavioural Processes*, 61, 37–45.
- Allan, L. G. (1980). A note on measurement of contingency between two binary variables in judgment tasks. *Bulletin of the Psychonomic Society*, 15, 147–149.
- Allan, L. G., & Jenkins, H. M. (1983). The effect of representations of binary variables on judgment of influence. *Learning and Motivation*, 14, 381–405.
- Alloy, L. B., & Abramson, L. Y. (1979). Judgement of contingency in depressed and nondepressed students: Sadder but wiser? *Journal of Experimental Psychology: General*, 108, 441–485.
- Blanco, F., Matute, H., & Vadillo, M. A. (2010). Contingency is used to prepare for outcomes: Implications for a functional analysis of learning. *Psychonomic Bulletin and Review*, 17, 117–121.
- Blanco, F., Matute, H., & Vadillo, M. A. (2011). Making the uncontrollable seem controllable: The role of action in the illusion of control. *The Quarterly Journal of Experimental Psychology*, 64, 1290–1304.

- Blanco, F., Matute, H., & Vadillo, M. A. (2012). Mediating role of activity level in the depressive realism effect. *PLoS One*, 7, e46203.
- Blanco, F., Matute, H., & Vadillo, M. A. (2013). Interactive effects of the probability of the cue and the probability of the outcome on the overestimation of null contingency. *Learning & Behavior*, 41, 333–340. doi: 10.3758/s13420-013-0108-8
- Bloom, C. M., Venard, J., Harden, M., & Seetharaman, S. (2007). Non-contingent positive and negative reinforcement schedules of superstitious behaviors. *Behavioural Processes*, 75, 8–13.
- Blum, S. H., & Blum, L. H. (1974). Dos and don'ts: An informal study of some prevailing superstitions. *Psychological Reports*, 35, 567–571.
- Buehner, M. J., Cheng, P. W., & Clifford, D. (2003). From covariation to causation: A test of the assumption of causal power. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 1119–1140.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104, 367–405.
- Dèttore, D., & O'Connor, K. (2013). OCD and cognitive illusions. *Cognitive Therapy Research*, 37, 109–121.
- Hannah, S., Allan, L. G., & Siegel, S. (2007). The consequences of surrendering a degree of freedom to the participant in a contingency assessment task. *Behavioural Processes*, 74, 265–273.
- Hannah, S., & Beneteau, J. L. (2009). Just tell me what to do: Bringing back experimenter control in active contingency tasks with the command-performance procedure and finding cue-density effects along the way. *Canadian Journal of Experimental Psychology*, 63, 59–73.
- Langer, E. J. (1975). The illusion of control. *Journal of Personality and Social Psychology*, 32, 311–328.
- Lober, K., & Shanks, D. R. (2000). Is causal induction based on causal power? Critique of Cheng (1997). *Psychological Review*, 107, 195–212.
- Matute, H. (1996). Illusion of control: Detecting response–outcome independence in analytic but not in naturalistic conditions. *Psychological Science*, 7, 289–293.
- Matute, H., & Blanco, F. (2014). Reducing the illusion of control when the action is followed by undesired outcomes. *Psychonomic Bulletin & Review*, 21, 1087–1093. doi: 10.3758/s13423-014-0584-7
- Matute, H., Vadillo, M. A., Blanco, F., & Musca, S. C. (2007). Either greedy or well informed: The reward maximization – unbiased evaluation trade-off. In S. Vosniadou, D. Kayser, & A. Protopapas (Eds.), *Proceedings of the European Cognitive Science Conference* (pp. 341–346). Hove, UK: Erlbaum.
- Msetfi, R. M., Murphy, R. A., Simpson, J., & Kornbrot, D. E. (2005). Depressive realism and outcome density bias in contingency judgments: The effect of the context and intertrial interval. *Journal of Experimental Psychology: General*, 134, 10–22.
- Musca, S. C., Vadillo, M. A., Blanco, F., & Matute, H. (2010). The role of cue information in the outcome-density effect: Evidence from neural network simulations and a causal learning experiment. *Connection Science*, 22, 177–192.
- Rescorla, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology*, 66, 1–5.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York, NY: Appelton-Century-Crofts.
- Reuven-Magril, O., Dar, R., & Liberman, N. (2008). Illusion of control and behavioral control attempts in obsessive-compulsive disorder. *Journal of Abnormal Psychology*, 117, 334–341.
- Shanks, D. R., & Dickinson, A. (1987). Associative accounts of causality judgment. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 21, pp. 229–261). New York, NY: Academic Press.
- Taylor, S. E., & Brown, J. D. (1988). Illusion and wellbeing: A social psychological perspective on mental health. *Psychological Bulletin*, 103, 193–210.
- Vallée-Tourangeau, F., Murphy, R. A., & Baker, A. G. (2005). Contiguity and the outcome density bias in action–outcome contingency judgements. *The Quarterly Journal of Experimental Psychology*, 58, 177–192.
- Wasserman, E. A. (1990). Detecting response–outcome relations: Toward an understanding of the causal texture of the environment. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 26, pp. 27–82). San Diego, CA: Academic Press.
- Wasserman, E. A., Kao, S.-F., Van Hamme, L. J., Katagiri, M., & Young, M. E. (1996). Causation and association. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *The psychology of learning and motivation* (Vol. 34, pp. 207–264). San Diego, CA: Academic Press.
- Wiseman, R., & Watt, C. (2004). Measuring superstitious belief: Why lucky charms matter. *Personality and Individual Differences*, 37, 1533–1541.
- Zebb, B. J., & Moore, M. C. (2003). Superstitiousness and perceived anxiety control as predictors of psychological stress. *Journal of Anxiety Disorders*, 17, 115–130.

Received November 26, 2013

Revision received July 15, 2014

Accepted June 22, 2014

Published online November 10, 2014

Fernando Blanco

Departamento de Fundamentos y Métodos de la Psicología
Universidad de Deusto
48007 Bilbao
Spain
E-mail fernandoblanco@deusto.es

Appendix

Full Instructions of the Experiment

These are the instructions for the Produce condition. Between square brackets we indicate the words that changed in the Prevent condition.

“At the beginning of this experiment, a light bulb will be presented on the screen. Your task is to learn how to keep this light ON [OFF] because your final score will depend on how many times the light was lit.

From time to time, a button will appear. This button indicates the start of a new trial, the occasion to do something.

While the button is on the screen, you will have the option of either making a button press response or not making a button press response. A button press response consists of pressing the spacebar on the computer keyboard **ONCE AND ONLY ONCE** immediately after the button appears. Not making a button press consists, of course, of doing nothing while the button is on the screen.

From the moment when the button appears, you will have 2 s to decide whether you want to press the button or not, before it disappears again. If you press the spacebar after the button has gone, the trial will count as a not press trial.

So, in this experiment there are only two possibilities as to what you can do on each of the trials:

- a) Press the button (by pressing the spacebar) within 2 s,
- b) Just sit back and observe what happens.

Immediately after the button has disappeared, you will see whether the light switches on or not.

If the light **SWITCHES ON**, you will **EARN [LOSE]** 1 point. Otherwise if the light **REMAINS OFF**, you will

LOSE [EARN] 1 point. Take this into account, because it is your goal to earn as many points as you can.

If the light switches on, it will stay on for 2 s before switching off by itself, thus allowing further tests. Besides, the same button will re-appear after a while when you can press it again. This means that, during the experiment, there will be many opportunities to press the button and see what happens.

You may find that the light will come on, on some percentage of the trials on which you do make a button press response. You may also find that the light will come on, on some percentage of the trials when you do not make a button press response. Alternatively, you may find that the light will not come on, on some percentage of the trials on which you do make a button press response. And, you may find that the light will not come on, on some percentage of the trials when you do not make a button press response.

So, there are four possibilities as to what may happen on any given trial:

- 1) you press and the light does come on;
- 2) you press and the light does not come on;
- 3) you don't press and the light does come on;
- 4) you don't press and the light does not come on.

Since it is your job to earn points by learning how to **TURN ON** the light [**KEEP** the light **OFF**], it is to your advantage to press on some trials and not on others, so you know what happens when you don't press as well as when you do press. That is, try to avoid pressing the button **EVERY TIME** it appears, and avoid also **NOT** pressing it **AT ALL**.

Remember, you must try to earn as many points as you can by keeping the bulb **ON [OFF]** as long as possible with your actions and omissions.

Good luck!”