Journal of Experimental Psychology: General
2007, Vol. 136, No. 2, 184–199

0096-3445/07/$12.00 DOI: 10.1037/0096-3445.136.2.184

Metacognition of Agency

Janet Metcalfe and Matthew Jason Greene
Columbia University

The feeling that we are agents, intentionally making things happen by our own actions, is foundational to our understanding of ourselves as humans. People's metacognitions of agency were investigated in 4 experiments. Participants played a game in which they tried to touch downward scrolling Xs and avoid touching Os. Variables that affected accuracy included speed of the scroll, density of the targets, and feedback. Of central interest were variables directed not only at accuracy but also at people's control: the turbulence of the cursor and how close the cursor had to come to the target for a hit (i.e., "magic"). After each trial, people made judgments of agency or judgments of performance. People were selectively sensitive to the variables to which they should be responsive in agency monitoring—whether the cursor moved in close synchrony to their movements and whether targets disappeared by magic. People knew, separably from their objective or judged performance, when they were in control and when they were not. These results indicate that people can sensitively monitor their own agency.

Keywords: metacognition, agency, motor plan, control, control judgments

The idea that people are agents—that they are in control of themselves—is at the heart of our legal system, is focal for our definition of sanity, and is central for our understanding of ourselves as human beings. According to Bandura (2001), "To be an agent is to intentionally make things happen by one's own actions. . . . The capacity to exercise control over the nature and quality of one's life is the essence of humanness" (p. 1). The profound implications of agency are predicated on the idea that people are, in fact, able to monitor their own agency, that is, to make metacognitive assessments about when and whether they are, themselves, in control. The distinction between making things happen, intentionally, oneself, as opposed to them just happening without one's willful intention, has a long history. For example, although no such distinction can be found in the first version of the Hittite laws, by the third revision, written over 3,000 years ago, the law poses differential penalties on agentic versus nonagentic acts: "If anyone blinds a free man in a quarrel he shall give 1 mina of silver. If (only) his hand does wrong he shall give 20 shekels" (from a clay tablet, 13th century B.C., in the Archeological Museum, Istanbul, Turkey, as translated by the museum).

Despite the importance of agency, the ontological status of this philosophical and psychological construct—that people do, themselves, control or will their own actions—has sparked considerable scientific debate. Most researchers have agreed that the intention, as something separate from the brain, at least, does not and could not occur before some brain manifestation of the action (Haggard, 2005; Libet, 2004). The mind does not tell our brain to do something that is then actualized materially. Rather, the readiness potentials and the plans for our acts are evident in the brain before one consciously knows that they are: "The brain . . . decides to initiate . . . or prepares to initiate the act at a time before there is any reportable subjective awareness that such a decision has taken place" (Libet, Gleason, Wright, & Pearl, 1983, p. 640).

This observation has provoked some researchers to comment that free will is illusory (Wegner, 2002). But, whether one actually has free will or not, what is beyond dispute, and is deeply entrenched, is that people feel that they are in control of their own actions. The conditions under which people do and do not feel themselves to be agents and the extent to which they feel in control of their own actions, whether correctly or not, may have deep psychological consequences for the individual (Morsella, 2005). Bandura (2001) stressed the importance of agency, which depends not just on being exposed to stimuli but on active agentic exploration, on brain development, on health, on self-esteem, and on self-enabling behavior. Whether people are considered to be sane or schizophrenic and delusional and whether they accept personal responsibility for their actions is consequent to their feelings of agency. On a more mundane note, our metacognitions of agency may well determine our preferences for one tennis racket, carving knife, golf club, computer, car, or surgical instrument over another. People like feeling in control and seek out instruments that afford this feeling.

Among those studies that have looked at people's feelings of controlling their own actions, a number of intriguing and dramatic misattributions of agency have been documented. People sometimes perform actions themselves but attribute them to an external force. That people are so intrigued when others make such blatant misattributions speaks to how deeply engrained are people's feelings of controlling their own actions. Wegner (2002) and his colleagues cited such phenomena as alien hand syndrome, dissociative identity disorder, schizophrenic auditory hallucinations, hypnosis, Ouija boards, water dowsing, speaking in tongues, spirit possession, and channeling. In these cases, a person is the source of the action and yet fails to claim it as their own, instead attributing the act to an outside agent. Wegner (2003) proposed

Janet Metcalfe and Matthew Jason Greene, Department of Psychology, Columbia University.

Correspondence concerning this article should be addressed to Janet Metcalfe, Department of Psychology, Columbia University, New York, 10027. E-mail: jm348@columbia.edu

that the feeling of willing something comes about when a thought appears in consciousness just before an action, when the thought is consistent with the actions, and when there are no alternative causes for the action. The attribution of self-agency is post hoc by this view: It is assessed once the outcome is seen.

Supporting evidence comes from a study by Wegner and Wheatly (1999), in which they showed that when normal college students were primed with an intention before making a movement (that was, in fact, irrelevant), they were more likely to claim that they had executed an intentional act despite the fact that they were not the cause of the act. Similarly, Wegner, Sparrow, and Winerman (2004) showed that hearing instructions in advance resulted in participants expressing a greater feeling of control over the movement of someone else's hands that were cleverly contrived, via smocks and mirrors, to look like they might be the participant's own hands. By Wegner's (2003) post hoc attribution theory, one might expect that people would feel themselves in control when they have a conscious goal in advance of the action (such as to do well on a task) and when the outcomes of a series of movements that they perceive themselves to take are favorable as they intended (cf., Knoblich & Sebanz, 2005).

Other researchers, too, have investigated misattributions of agency (e.g., Blakemore, 2003; Blakemore, Wolpert, & Frith, 2002; Knoblich, Stottmeister, & Kircher, 2004) such as those that occur when a person with schizophrenia attributes his or her own actions to an external source or when hypnotized individuals claim that the hypnotist is in control of their actions, for example. In these studies, people's feeling of agency was assessed by means of their attribution of control to themselves as compared to another person, rather than a simple assessment of whether or not they were in control. Such syndromes have been used to study people's ability to correctly attribute actions to their veridical source: the self or the other. Although the question of whether people correctly determine whether the source of an action is themselves or someone else is undoubtedly related to their own feelings of being in control of their own actions, authorship and feelings of agency are not necessarily identical. One might recognize one's own movements and distinguish them from another's, either concurrently or retrospectively (Tsakiris, Haggard, Franck, Mainly, & Sirigu, 2005), for example, without necessarily feeling that one is controlling the act. Thus, a feeling of control or agency might differ from an ascription of authorship (Sato & Yasuda, 2005). Indeed, one could make a correct attribution of self-authorship but still not feel in control of the action. There have been many interesting demonstrations that people can retrospectively recognize their own performance (Flach, Knoblich, & Prinz, 2004; Knoblich & Prinz, 2001; Repp & Knoblich, 2004). Furthermore, the prediction of consequences is better when the actor was oneself: A video of oneself throwing a dart allows better predictions of where the dart will land than does a similar video of another throwing a dart (Knoblich & Flach, 2001). The involvement of the self appears to provide some privileged access, even retrospectively. But whether this superior recognition of self-generated patterns of action results because of a feeling of agency or occurs only because one is more familiar with the nuances of one's own patterns than with those of someone else is not yet clear.

Some patients exhibit great difficulty with this, however. Daprati et al. (1997) compared participants without schizophrenia and people with schizophrenia with and without hallucinations and/or delusional experiences. These people executed simple finger and wrist movements without direct visual control of their hand. The image of either their own hands or alien hands executing the same or a different movement was presented on a television screen in real time. People without schizophrenia made this discrimination well. People who were hallucinating and deluded schizophrenic patients were more impaired in discriminating their own hand from the alien hand than were the nonhallucinating patients and tended to misattribute the alien hand to themselves. Knoblich, Stottmeister, and Kircher (2004) have reported a similar failure to accurately monitor their own actions in people with schizophrenia.

Leube, Knoblich, Erb, and Kircher (2003) observed, in a functional magnetic resonance imaging study, that there was a specific activation in the right fronto-parietal cortex when there was an abrupt mismatch of one's own body movement and its visual consequences, as compared with a similar abrupt mismatch between one's own movement and somebody else's visually perceived hand movement. Other studies (Leube, Knoblich, Erb, Grodd, et al., 2003) have also implicated the parietal lobe as well as the dorsolateral prefrontal cortex and the cingulate (Jahanshahi & Frith, 1998), the cerebellum (Blakemore, Wolpert, & Frith, 1998), and Brodmann area 6—the mirror neuron area (Gallese & Goldman, 1998)—in monitoring one's own and others' actions. Two studies suggested that activation in the anterior insula is related to feelings of control, whereas less control, or attributions of agency to another rather than to the self, is related to activation in the right inferior parietal cortex (Farrer et al., 2003; Farrer & Frith, 2002). These studies suggested that there may be a deeply embedded neurocorrelate associated with feelings of agency. Unfortunately, people are rarely asked about their subjective impressions in these studies, so it is not known whether being subjectively aware of agency is involved. Nevertheless, the findings are provocative and suggest that there may be a deep and fundamental reason, and perhaps even a dedicated brain circuit, underlying our feelings of self-agency.

A number of researchers have proposed feed forward models of motor control (Jordan & Rumelhart, 1992; Kawato, Furukawa, & Suzuki, 1987; Kinsbourne, 1995; Miall & Wolpert, 1996), whereby our motor system monitors what is happening at a proprioceptive level, and this is matched and calibrated by what is called an efference copy of the motor plan. The idea that people run an internal simulation of their movements, as well as monitoring the feedback from the movements themselves, was originally formulated as an explanation of and model of fine-grained motor movements rather than as a model of subjective feeling states. However, such a model lends itself to an explanation of subjective feelings of agency. Many researchers (Blakemore, 2003; Jeannerod, 2001; van den Bos & Jeannerod, 2002; Wolpert, 1997; Wolpert, Ghahramani, & Jordan, 1995; and see, especially, Sebanz & Prinz, 2006) have proposed that concordance or discrepancy between predicted sensory-motor feedback and actual feedback is assessed by a comparator, and this comparison might allow for feelings of agency. If there is concordance, the person him- or herself enacted the action; if there is discrepancy, the action was external. A lesion or distortion in either system could readily result in misattributions, such as those that occur under hypnosis, with alien hand syndrome, or in people with schizophrenia (see, Blakemore & Frith, 2003; Frith, 1992; cf., Gallagher, 2004).

Despite these intriguing findings and models, metacognitions of agency in normal people have rarely been studied. One reason for the lack of research on the parameters that affect feelings of agency with normal people under conditions not contrived to provoke illusions may be the perception, among researchers, that the feeling that the "I" is causing something to happen (Haggard, 2005; Haggard, Clark, & Kalogeras, 2002; Lau, Rogers, & Haggard, 2004; Wegner, 2004, 2005) is tricky to pin down. Fourneret and Jeannerod (1998) have shown, for example, that the discrepancies between what a person does and what the outcome is in the world can be very large, without the person consciously noticing anything unusual. This observation led these researchers to suggest that there is limited spontaneous conscious monitoring of motor performance. Nisbett and Wilson (1977) noted, many years ago, that people have little introspective appreciation of any of the processes of cognition. As Metzinger (2003) has put it,

> The first difficulty in describing the content of volition is perhaps the biggest and most fundamental one, and it is very rarely made explicit: The phenomenology of volition is thin and evasive. Thin means that the conscious experience of willing something is not as crisp and vivid as, say, the conscious experience of seeing colors or of feeling pain. . . . "Evasive" means that attempts to deliberately focus one's introspective attention on the different steps and stages of conscious will-formation often lead to a state in which the target of introspection—will itself—seems to *recede* or to *dissolve*. (pp. 2–3)

Nevertheless, people do feel varying degrees of control, and their sense of self-agency may be important in many domains of life (Bandura, 2001). People may feel nonagentic when things seem to be happening by accident; when things are going wrong; when the world or chance circumstance intervenes; when their plans do not correspond to their perceptions of the outcome; when their own attributions, emotions, or feeling states (see Clore, 1992) are faulty; or when the outcome of their actions is not salient. We suggest that we may be able to investigate such metacognitions of agency by using methodologies that have been successfully applied to other kinds of metacognitions. In the study of other metacognitions (including people's judgments of learning, Dunlosky & Nelson, 1992; Hertzog & Hultsch, 2000; Koriat, 1997; Koriat & Bjork, 2005; Metcalfe & Kornell, 2005; Son, 2004; Thiede, 1999; Thiede & Dunlosky, 1999; their feelings of knowing, Metcalfe, 1993; Nelson & Narens, 1994; or their tip of the tongue states, Schwartz & Smith, 1997), multiple observations are obtained in varying circumstances, often in factorially designed experiments, allowing the researchers to investigate the factors that contribute and do not contribute to the metacognitive judgments. These metacognitive judgments have proven to be highly reliable and replicable and to vary in predictable and systematic ways. Such studies allow exploration of the heuristics (Koriat, 1997) that people use to make these judgments. We suggest that metacognitions of agency, too, may be investigated by similar means.

Despite the importance of feelings of self-agency or control, then, people's metacognitions concerning their own agency have been accorded little systematic investigation in relatively normal circumstances. Past experiments, in which researchers looked at pathological cases of things like alien hand syndrome, Ouija boards, hypnosis, and schizophrenia, although intriguing, failed to capture the nuances of personal metacognition of agency that people live out in their normal lives. In this article, we begin to study some of the many variables that may contribute to these metacognitions. Wegner (2004, p. 658) noted,

> Although the proper experiments have not yet been done to test this, it seems likely that people could discriminate the feeling of doing from other feelings, knowing by the sheer quality of the experience just what has happened. The experience of willing is more than a perception of something outside oneself, it is an experience of one's own mind and body in action.

We undertake to do just this.

Experiment 1

In this experiment, participants used a computer mouse to move a small box back and forth across a horizontal track on the screen as 16 stimuli, Xs and Os, randomly distributed across the play area, scrolled by. Participants were instructed to collect one type of stimulus by moving the mouse and "catching" them with their box and to avoid the other type. If a hit occurred, the target would disappear; if not, it would continue traveling downward on the screen. Participants played the game for 15 s on each trial, then made a judgment of agency (JoA). In this and subsequent experiments, we varied a number of parameters that allowed us to evaluate whether people were able to make JoAs that were reasonable and well-determined by their actual control and whether they responded in systematic ways to parameters that might be expected to alter their JoAs. The main parameters that were manipulated in the first experiment were turbulence of the mouse control, speed of the scroll, and density of the targets. We expected that each of these parameters would affect people's performance and sought to determine whether (and which) factors that affected performance would have an impact on people's JoAs.

Method

Participants. The participants in this experiment were 24 Columbia University or Barnard College students and faculty members and 1 New York University faculty member, 13 of whom were female participants and 11 of whom were male participants, who either volunteered without pay or received either course credit or pay for participating. Participants were treated in accordance with the American Psychological Association regulations.

Apparatus. All experiments were conducted on individual iMac computers, which were used with a mouse and mouse pad.

Procedure: The instructions given to participants were,

> In this experiment, we are interested in people's **metacognitions of control** or their feelings about when they are causing things to happen. You have probably heard stories, and maybe even have had the experience, of going to an arcade and you start playing a video game—thinking you are controlling what is happening. But when you let go of the levers, it just keeps on happening: you weren't actually controlling anything, but you *thought* you were. On the other hand, Ouija boards were popular, in the past, because although the people who were holding the board were actually moving it themselves, they didn't *think* they were. As a result they were willing to believe that external spooky forces were moving the board (and that they were communing with the spirits!). Our point is that they were actually moving it themselves, but their metacognition of control was very low: they didn't think they were doing it. Metacognition of control can differ in more mundane circumstances, too. Sometimes if you are

> driving, the steering maybe very loose and you don't feel like you are in control. At other times (like when you're at the wheel of a hot sports car) there's a great "road feel" and you simply know that you are controlling every move. Regardless, then, of whether you actually are in control or not (which is not our question here) you may sometimes feel like you are in control (and hence have a high metacognition of control) or feel like you are not (and have a low metacognition of control).
>
> In this experiment you are going to play a game of many trials in which you will use the computer mouse to move a box on a blue track. Your job is to touch all of the Xs as they come into range and to avoid touching any of the Os. After playing for a while, a screen will pop up to ask you for your metacognition of control during the immediately preceding trial. If you felt like you were very much in control click toward the "in control" end; if you felt like you were not in control, click toward the "not in control" end, to reflect your own feeling of control. There will be many trials, so don't labor over any one answer. And do your best to touch all the Xs and none of the Os.

Participants were given a practice trial at playing the game (in which the Xs or Os disappeared as soon as the person touched them but continued on screen below the mouse track, if the person did not touch them). The practice was at a slow rate. They then made their JoA, on a sliding scale for which they had to pull a slider to the value of felt control they wanted to indicate. The experimenter then asked whether there were any questions, and if there were, he or she answered them. If not, he or she left the room to allow the participant to proceed uninterrupted through the experiment. At the end of the experiment, the participant was questioned about what they had done, was debriefed, given credit or pay, and thanked.

Design. The experiment was a 2 × 2 × 2 within-participants design, in which the variables were turbulence (no turbulence or turbulence), speed (fast or slow), and density of targets and nontargets (equal or 7:1), with eight replications randomly ordered in each of the eight treatment combinations. There was an additional variable—whether the participant was told to touch Xs and avoid Os or the reverse—that was counterbalanced between participants. The first variable, turbulence, was manipulated as follows. In the no turbulence condition, movements of the mouse were directly reflected in movement of the box on the screen. In the turbulence condition, with each mouse movement made by participants, a noise component was added:

$$\Delta x' = \Delta x + \sigma \sin(2\pi t/3),$$

where Δx is the distance the participant actually moved the mouse, t is time in seconds, σ is the amplitude of the noise wave, and $\Delta x'$ is the resultant movement on the computer screen. This unpredictable noise algorithm was used in favor of a completely random noise distribution in order that the participants not feel as if their actions had no effect whatsoever on the outcome.

The second within-participant variable was the speed of the scroll, which either passed by the mouse track quickly or slowly. On average, about 18 objects appeared in the slow condition, whereas 54 objects appeared in the fast condition in each 15 s interval. The third variable was density of the targets relative to nontargets. In the equal condition, they were split 50/50; in the 7:1 condition, they were 87.5% Xs and 12.5% Os (or the reverse, i.e., 7 targets for each nontarget). We expected that this manipulation would affect performance and might influence JoAs.

There were two dependent variables of interest. The first was people's performance: How well did they do on touching the Xs and avoiding the Os? We computed performance using hit rate (reported as a proportion) and hits minus false alarms. We could not compute d′ consistently because there were a large number of trials on which people made no errors. The hits and the hits minus false alarms measures produced identical patterns of results in this and subsequent experiments, and so, we will report only the hit rate, though details of the other analyses are available through correspondence with Janet Metcalfe. The second dependent variable was JoA. This was the value that the participant entered on the slider after having completed the game trial. The JoA analogue scale produced a value between 0 and 1, continuously, depending on where it was placed. Figure 1 shows an example of the screen a participant would have seen after having just completed a trial on which he or she gave a moderate JoA.

Results

We used partial eta squared as our measure of effect size, throughout, for analysis of variance (ANOVA) results and r^2 for t-test results. A $p < .05$ was taken as the criterion for significance. Whether the targets were Xs or Os had no effect, therefore we collapsed over this variable for all of the analyses reported below.

Two of the three variables of interest affected hit rate. People had a higher hit rate in the slow than in the fast condition, $F(1, 23) = 367.28$, mean square error (MSE) = 0.01, $p < .001$, $\eta_p^2 = .94$. Hit rate was higher in the nonturbulent condition than the turbulent condition, $F(1, 23) = 61.64$, $MSE = 0.01$, $p < .001$, $\eta_p^2 = .73$. There was no difference in hit rate as a function of density ($F < 1$, $p > .1$). All means and standard deviations for both the hit rates and the JoAs for each treatment combination are shown in Appendix A. Our interest, of course, was in how these variables interacted with people's JoAs, that is, in the interactions between hit rate and JoAs, which were analyzed with a multivariate analysis of variance (MANOVA), reported below.

There were three interactions with measure (i.e., hit rate as compared with JoA) that were significant. The first was turbulence, $F(2, 22) = 44.83$, $p < .001$, $\eta_p^2 = .80$. As can be seen in Figure 2A, when there was no turbulence, people's hit rate and their JoAs tracked each other better than when there was turbulence. The difference between the hit rate and JoA in the no turbulence condition (.06) was smaller than was the difference in the turbulence condition (.19), $t(23) = 5.22$, $p < .001$, $r^2 = .54$. With high turbulence, even though people's performance was quite good, they judged themselves to be less in control than in the no turbulence condition, $t(23) = 7.34$, $p < .001$, $r^2 = .70$. This was our first indication that people could sensitively assess their metacognitions of control.

The second variable that interacted with measure was speed, $F(2, 22) = 176.55$, $p < .001$, $\eta_p^2 = .94$. At a fast rate, JoAs and hit rate did not differ ($t < 1$, $p > .1$), whereas at a slow speed, hit rate was higher than JoA, $t(23) = 6.74$, $p < .001$, $r^2 = .66$. People's JoAs did increase in the condition in which they had better performance. The slow speed JoAs were slightly higher than were the fast speed JoAs, $t(23) = 4.62$, $p < .001$, $r^2 = .48$. In summary, speed affected hit rate more than it affected JoAs, as is shown in Figure 2B.

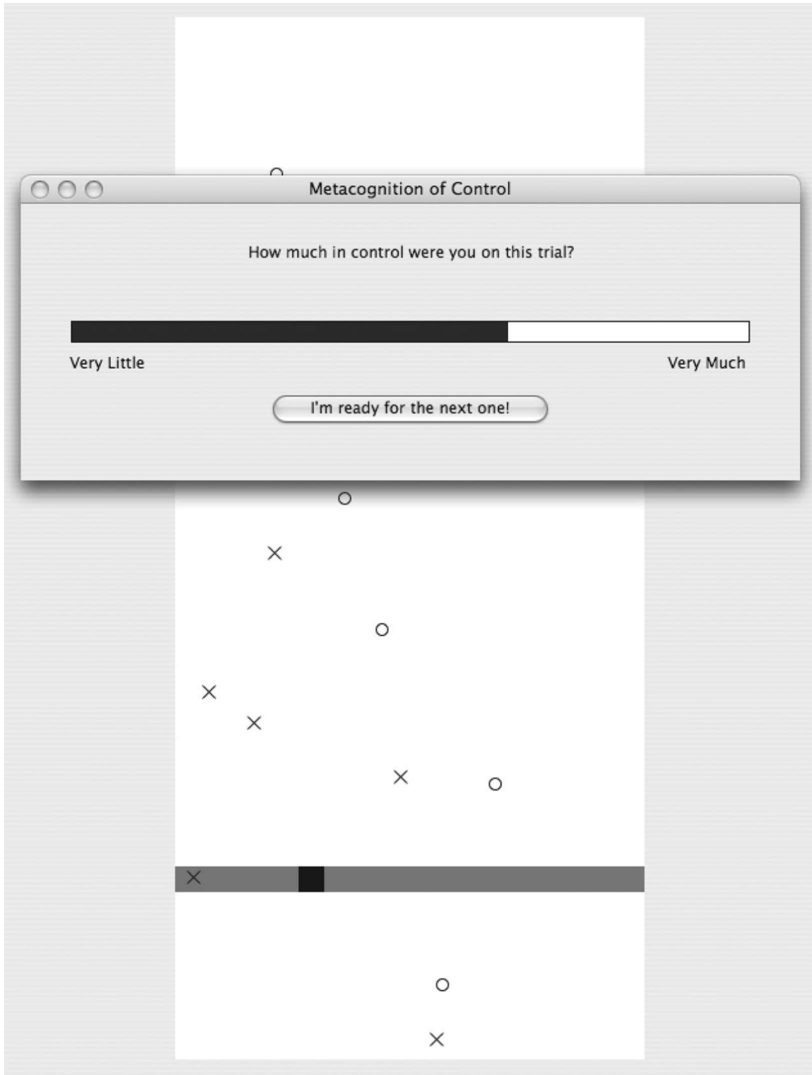

Figure 1. An example of the screen, including targets and nontargets, that a participant would see after having made a moderate judgment of agency.

Finally, there was a significant interaction between measure and density, $F(2, 22) = 5.78$, $p < .01$, $\eta_p^2 = .35$, though this effect was very small. As is shown in Figure 2C, although hit rate was the same in the equal density and the 7:1 density conditions ($t < 1$, $p > .1$), in the equal density condition, JoAs were lower (but only by 3%) than in the 7:1 condition, $t(23) = 3.26$, $p < .01$, $r^2 = .32$. The fact that there were more targets in the 7:1 condition may have led people to feel slightly more in control, even though their hit rate was unaffected.

To further investigate the relation between people's performance on the primary task and their JoAs, we collapsed across conditions and computed both gamma and Pearson's correlations between performance and judgments for each participant. The overall correlations were significantly different from zero. The Pearson's correlation was .37, $t(23) = 8.37$, $p = .001$, $r^2 = .75$; the gamma was .27, $t(23) = 7.91$, $p = .001$, $r^2 = .73$.

Discussion

The finding that there was an above zero correlation between JoAs and performance suggests a role for people's performance in their JoAs. When people performed well, regardless of treatment combination, they gave higher JoAs than when they performed poorly, suggesting that something about how well they did made its mark on their feelings of control. However, the finding that there were interactions between JoAs and hit rates suggested that a sensitivity to some additional factor or factors was important in people's assessments of their own control.

People's JoAs were sensitive, for example, to the turbulence that we had introduced into the mouse responsiveness. Performance differences between the turbulent and nonturbulent conditions could not fully account for this JoA difference. When people were not in good control of the mouse because we had modulated the

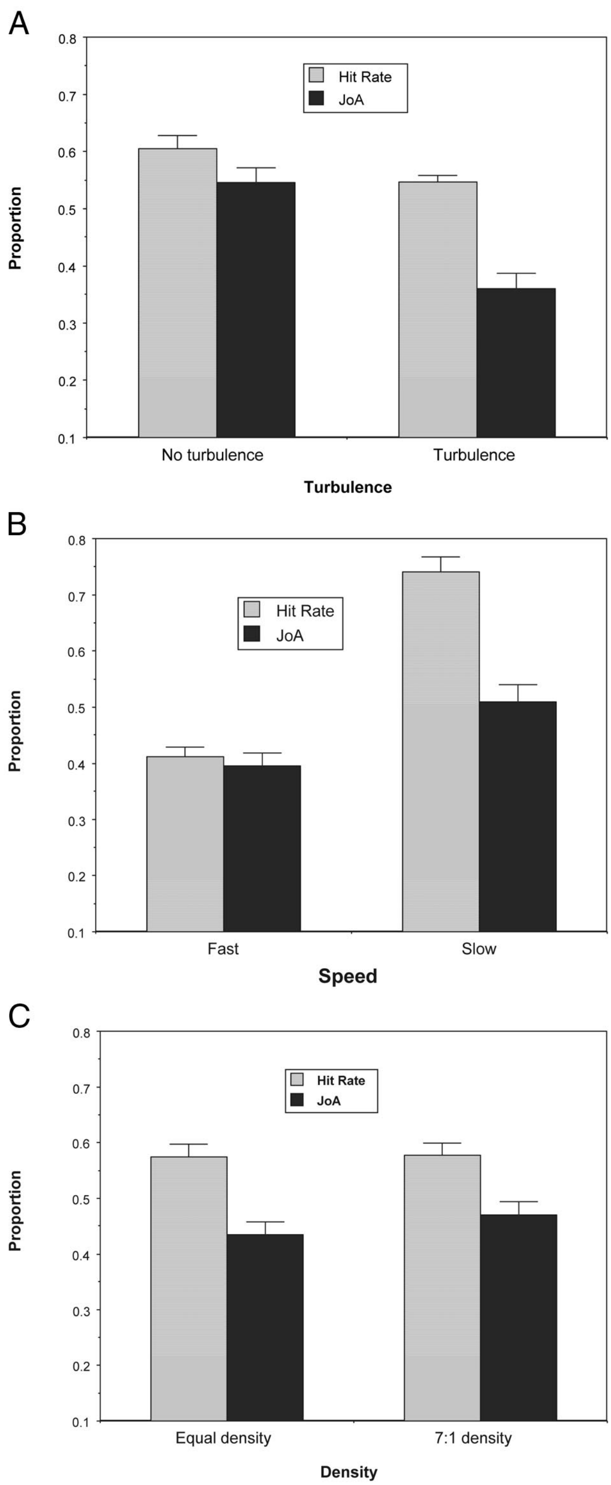

Figure 2. The significant interactions between hit rate and judgments of agency (JoAs) in Experiment 1. A: This panel shows the data for turbulence. B: This panel shows the data for speed. C: This panel shows the data for density. The error bars indicate standard errors of the mean.

function relating their movements to what happened to the cursor on the screen in a nonlinear way, this lack of control—although not showing up much in their performance—was reflected in large negative changes in their JoAs. A similar interaction occurred with differences in speed. Encouraged by these results, we decided to attempt a replication in Experiment 2. We also decided to introduce two new variables—"magic" and auditory feedback.

Experiment 2

Experiment 2 was similar to Experiment 1, though with some variations. We eliminated the density variable, which had produced no differences in performance and very small (though significant) differences in JoAs. Eliminating this variable allowed us the time to include other factors that could potentially be of great interest.

In Experiment 2, we introduced two new variables that we hypothesized would impact the JoAs in a specific way. The first new variable we called magic. In the no magic condition, the person had to place the small box on the screen (which was controlled by the mouse) directly on the target to get credit for a hit and to make the target disappear. (They also had to place the cursor directly on the nontarget to be considered to have made a false alarm.) In the magic condition, though, the criterion for a hit was more lenient: They only had to place the cursor within 10 pixels (about half a centimeter) of the target to be given credit for the hit and for the target to disappear. We left the nontargets' response circumference unchanged: People still had to place the cursor directly on the nontargets to be docked with a false alarm. Thus, their performance in the magic condition was artificially inflated because we had made the task easier by magic.

Note, though, that the targets were not disappearing randomly. The person still had to be in the vicinity, in a controlled way, to be scored with a hit. If JoAs were based on the outcome of people's performance, as the correlations between performance and JoAs in Experiment 1 suggested, then they should have increased with the magic manipulation. However, if people were, instead, closely monitoring the real-world consequences of their intended actions with the feedback from those actions, then this magic manipulation may have resulted in a discrepancy and lower JoAs than would be expected on the basis of performance alone.

We also varied whether people were given auditory feedback, in the form of a ping sound, when they got a hit, and a thud sound when they made a false alarm. We thought that feedback might make their performance more apparent to them and that this would better entrain their JoAs to their performance. In the case in which the result was less pronounced, as in the no auditory feedback condition, monitoring should have been more difficult.

The other variables were speed and turbulence, as in Experiment 1. We expected to replicate the results of Experiment 1 on these variables. In Experiment 2 and subsequent experiments, though, we doubled the amplitude of the turbulence so that it would be more salient to the participants. We expected to replicate (and perhaps increase) the differential effect of turbulence on performance and JoAs that we had observed in Experiment 1.

Method

Participants were 24 Columbia University students or Barnard College students, 14 females and 10 males, who were given either course credit or pay for participating. They were treated in accordance with the American Psychological Association standards of ethics. An additional participant was observed to be mistakenly treating the wrong stimulus type as the target type; that participant's data were eliminated.

The design was a 2 × 2 × 2 × 2 within-participants design, in which the factors were turbulence, speed, magic, and auditory feedback. The new variable, magic, was programmed such that in the no magic condition, the cursor had to be directly over the target to be scored as a hit, whereas in the magic condition a hit was credited (the target disappeared and auditory feedback might be given) if the cursor were within 10 pixels of the target. Auditory feedback (the pings and thuds) was given in the feedback conditions, whereas no feedback other than the visual disappearance of the X or O when it was a hit or false alarm, because of the positioning of the participant's cursor, was given in the no feedback condition. The only other change from Experiment 1 to Experiment 2 was that the turbulence or the magnitude of the noise function in the turbulent condition was doubled.

Results

The hit rates were higher in the nonturbulent condition than in the turbulent condition, $F(1, 23) = 19.02$, $MSE = 0.05$, $p < .001$, $\eta_p^2 = .45$; they were higher in the slow condition than in the fast condition, $F(1, 23) = 349.05$, $MSE = 0.02$, $p < .001$, $\eta_p^2 = .94$. They were higher in the magic condition compared with the no magic condition, $F(1, 23) = 239.89$, $MSE = 0.03$, $p < .001$, $\eta_p^2 = .91$. There was an interaction between turbulence and magic, $F(1, 23) = 14.41$, $MSE = 0.01$, $p < .001$, $\eta_p^2 = .39$. There was also an interaction between magic and speed, $F(1, 23) = 32.61$, $MSE = 0.01$, $p < .001$, $\eta_p^2 = .59$. There was also a triple interaction among turbulence, magic, and speed, $F(1, 23) = 6.64$, $MSE = 0.01$, $p = .02$, $\eta_p^2 = .22$. All cell means and standard deviations for the hit rate as well as for the JoAs are given in Appendix B.

To investigate the relation of JoAs to hit rate, we conducted a MANOVA. There was an interaction between hit rate and JoA on the turbulence variable, $F(2, 22) = 17.31$, $p < .001$, $\eta_p^2 = .61$. This interaction was like that observed in Experiment 1: In the no turbulence condition, the JoA rating was less different from the hit rate (.11), $t(23) = 3.62$, $p < .01$, $r^2 = .61$, than was the difference (.27) in the turbulent condition, $t(23) = 7.91$, $p < .001$, $r^2 = .73$. The difference between these differences was significant, $t(23) = 4.90$, $p < .001$, $r^2 = .51$, as is shown in Figure 3A.

There was an interaction with hit rate, compared with JoA, and speed, $F(2, 22) = 166.94$, $p < .001$, $\eta_p^2 = .94$. As in Experiment 1, people's performance increased greatly from the fast to the slow speed, $t(23) = 18.73$, $p = .001$, $r^2 = .94$, but their JoAs did not increase proportionately, $t(23) = 1.03$, $p > .1$. This interaction is shown in Figure 3B.

There was a significant interaction between the hit rate and JoAs with magic, $F(2, 22) = 125.00$, $p < .001$, $\eta_p^2 = .92$. Magic had a substantial effect on performance, as indicated above, increasing by 27% from the no magic to the magic condition. Participants' JoAs also increased, but only slightly (by 7%), with magic, $t(23) = 4.57$, $p < .001$, $r^2 = .48$. As is shown in Figure 3C, the difference between performance and hit rate in the no magic condition was smaller than this difference in the magic condition, $t(23) = 12.73$, $p < .001$, $r^2 = .88$.

There was a significant (though small) interaction between hit rate and JoAs, depending on whether people were provided with auditory feedback, $F(2, 22) = 9.87$, $p < .001$, $\eta_p^2 = .47$. Although the presence of feedback did not change performance, the means were .59 and .58, for no feedback and feedback, respectively ($t < 1$, $p > .1$); people's JoAs were higher (and closer to their hit rate) when they had auditory feedback (.42), as compared with when they were given no feedback (.37), $t(23) = 4.19$, $p < .001$, $r^2 = .43$. The mean difference between hit rate and JoA was smaller with auditory feedback than without it, $t(23) = 4.01$, $p < .001$, $r^2 = .41$. This is shown in Figure 3D.

There was also an interaction between magic and speed on hit rate compared with the JoAs, $F(2, 22) = 15.81$, $p < .001$, $\eta_p^2 = .59$. There was an interaction between turbulence and magic, $F(2, 22) = 7.21$, $p < .01$, $\eta_p^2 = .40$. There was an interaction between turbulence, magic, and speed, $F(2, 22) = 6.72$, $p < .01$, $\eta_p^2 = .38$. These were all very small, and the cell means are given in Appendix B. No other interactions were significant.

As in the previous experiment, we computed the correlations (both Pearson's r and gammas) between JoAs and performance. The overall correlations for each participant were significantly greater than zero: Pearson's $r = .35$, $t(23) = 7.94$, $p < .001$, $r^2 = .73$; $\gamma = .26$, $t(23) = 6.84$, $p < .001$, $r^2 = .67$.

Discussion

The results of Experiments 1 and 2 suggested that people might use their performance (or perhaps their perception of their performance) as the basis for their metacognitions of agency much of the time. In many of the conditions in these experiments, the JoAs were rather close to people's hit rates. Furthermore, the correlations between people's JoAs and their hit rates were also significantly greater than zero, in both experiments. Finally, when performance was made more salient by auditory feedback, people's JoAs were closer to their hit rates than when performance was less salient, in the no feedback condition.

In certain treatment combinations, though, the interactions between hit rate and JoA suggested that something more than performance underlies people's JoAs. When the mouse control was distorted, people's JoAs reflected this distortion, despite little impairment in performance. When the hits were accomplished by magic, once again, people did not take full credit for their enhanced performance but rather were conservative in their control ratings. These two cases are straightforward examples, so it would seem, in which something other than performance outcome might have impacted control ratings. They suggest that people's JoAs might have been picking up on the discrepancy between what they, themselves, did or did not do and what the outcome was.

The other case—speed—in which we found an interaction between JoA and performance level is less clearly explicable as the detection of a discrepancy between what the person did and what happened. Consideration of this case led us to entertain the possibility that people might be basing their JoAs not on their real performance but rather on their perception or metacognition of their performance. But we did not know what their perception of their performance was because we had not asked that question. For this hypothesis about the basis for people's JoAs to be a possibility, people's metacognitions about their performance would have

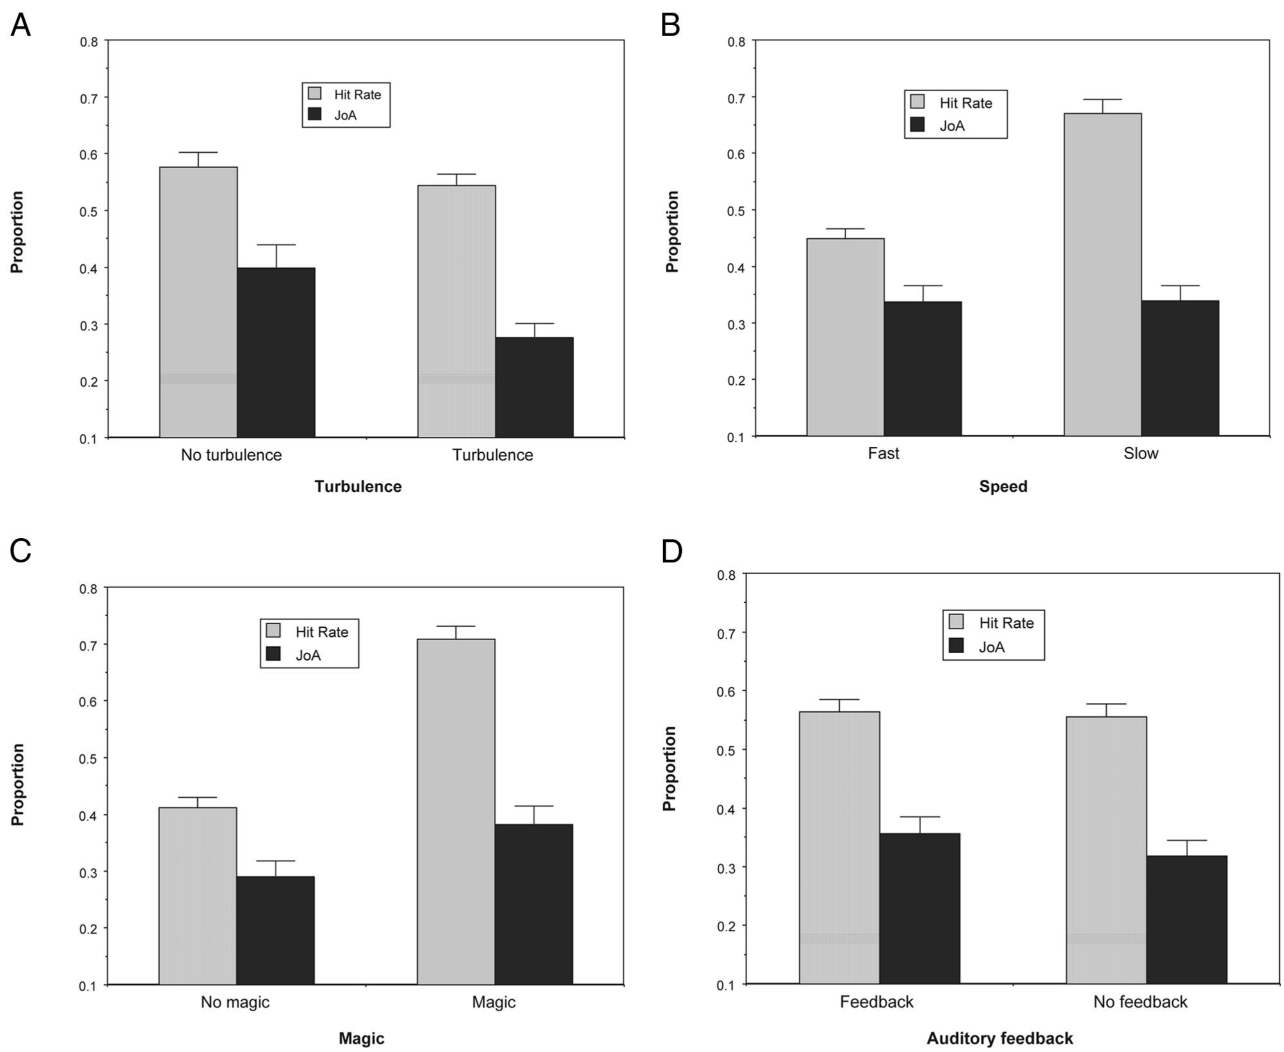

Figure 3. The significant interactions between hit rate and judgments of agency (JoAs) in Experiment 2. A: This panel shows the data for turbulence. B: This panel shows the data for speed. C: This panel shows the data for magic. D: This panel shows the data for feedback. The errors bars indicate standard errors of the mean.

to be biased and would have to be biased in the same way as were their JoAs. We had not requested judgments of performance (JoPs), however, so we had no information concerning whether people's perceptions of their performance might have been biased—and biased in just the right way to account for our JoA data.

There are many instances in the literature on metacognition in which systematic biases in judgments are in evidence, so we could certainly not rule out the biased metacognitions of performance hypothesis a priori. Perhaps, then, the reason for the interactions between the hit rates and the JoAs seen in the previous two experiments was that although they were basing their JoA on their metacognitions of performance and nothing more, those metacognitions of performance were themselves biased. Nothing in the first two experiments allowed us to rule out this possibility.

The judgment of how much in control participants are, by this view, reflects nothing more than the degree to which the outcome (their metacognition of performance) matches the intent. Translated into the terms of the present experimental situation, this means that given that their intent was to be perfect, their JoA might be just how closely they perceived themselves to have gotten to the goal of being perfect. If they assessed their performance level to be at 70% then their JoA would be 70%; if they assessed performance to be 50% then the JoA would be 50%. This hypothesis suggests that sensitivity to the nuances of parameters such as speed, which was important in the first two experiments, might have been a factor in altering JoAs only insofar as these parameters also altered people's metacognitions of their performance.

The fact that hit rate interacted with JoAs on the turbulence and magic variables—which were cases of distorted control—suggested that this may not be the case. But even these interactions did not definitively rule out the possibility that all people were doing was performance monitoring, given that we did not know that people's JoPs were not biased even in these cases, in which there were grounds to suppose that they might be picking up on discrepancies between what they, themselves, did as compared with what happened.

To test this hypothesis, we directly compared JoAs and JoPs in the experiment that follows. One hypothesis of Experiment 3, then, was that people were simply interchanging JoAs with JoPs. Such an explanation of the JoAs would require that people were not veridical in their assessments of performance, that is, that JoPs would systematically fail to track performance in just the way that JoAs had failed to track performance. Such a finding would certainly not be anomalous in the metacognitive literature, which has documented many cases of systematic biases in metacognitive judgments, including both underconfidence and overconfidence.

The alternate hypothesis was that people were doing something different from assessing their performance when they made JoAs. If people's JoPs interacted with their JoAs in this experiment, in which the two were directly compared, then the interactions found in the previous two experiments could more confidently be ascribed to metacognition of agency, per se.

Experiment 3

Method

The participants were 38 Columbia University or Barnard College students, 20 of whom were male and 18 of whom were female, who received course credit or pay for participating. They were treated in accordance with the American Psychological Association standards of ethics. The experiment was like Experiment 2, except that participants were randomly assigned to a judgment type, in which they made either JoAs or JoPs. Of the participants, 19 were randomly assigned to the JoA condition, and 19 were randomly assigned to the JoP condition. All other variables were the same.

The instructions given in the JoA condition were the same as had previously been used. The instructions in the JoP condition were as follows:

> In this experiment, we are interested in people's **metacognitions about their performance**, that is, their assessments of how well they did. You are going to play a game of many trials in which you will use the computer mouse to move a box on a blue track. Your job is to touch all of the Xs as they come into range and to avoid touching any of the Os. After playing for a while, a screen will pop up to ask you for your metacognition about your performance during the immediately preceding trial. If you felt like you did quite well, getting many of the Xs and avoiding most of the Os, click toward the "very high" end; if you felt like you did not do well, click toward the "very low" end, or click anywhere in between, to reflect your own assessment of performance on the immediately preceding trial. There will be many trials, so don't labor over any one answer. And do your best to touch all the Xs and none of the Os. If you have any questions, please ask your experimenter now.

Results and Discussion

The hit rates showed a main effect of judgment condition, such that people performed better in the JoA (.65) than the JoP (.59) condition, $F(1, 36) = 4.22$, $MSE = 0.16$, $p < .05$, $\eta_p^2 = .11$.

There was an effect of turbulence on hit rate, such that the nonturbulent condition was higher than the turbulent condition, $F(1, 36) = 61.76$, $MSE = 0.03$, $p < .001$, $\eta_p^2 = .63$. There was an effect of speed, such that slow was better than fast, $F(1, 36) = 1524.06$, $MSE = 0.01$, $p < .001$, $\eta_p^2 = .98$. There was an effect of magic, such that the magic condition was higher than the no magic condition, $F(1, 36) = 897.57$, $MSE = 0.02$, $p < .001$, $\eta_p^2 = .96$. There was an effect of feedback, $F(1, 36) = 5.52$, $MSE = 0.01$, $p = .02$, $\eta_p^2 = .13$. There was also an interaction between turbulence and magic, $F(1, 36) = 13.42$, $MSE = 0.01$, $p < .001$, $\eta_p^2 = .27$; an interaction between magic and speed, $F(1, 36) = 46.68$, $MSE = 0.01$, $p < .001$, $\eta_p^2 = .57$; and an interaction among turbulence, speed, feedback, and judgment type, $F(1, 36) = 8.33$, $MSE = 0.01$, $p < .01$, $\eta_p^2 = .19$. The cell means for these hit rates, as well as for both JoAs and JoPs, along with their standard deviations, are given in Appendix C.

We conducted an ANOVA to investigate the primary interest in this experiment, namely, differences between the two manipulated judgment types. We were especially interested in the interactions with magic, speed, turbulence, and feedback, insofar as these had interacted with hit rate in the previous experiments. In this experiment, however, only turbulence and magic showed an interaction with judgment type. Speed and feedback both were insignificant.

The interaction between turbulence and judgment type was significant, $F(1, 36) = 8.24$, $MSE = 0.09$, $p < .01$, $\eta_p^2 = .19$. Under conditions of turbulence, as compared with no turbulence, people's JoPs changed less than did their JoAs: They recognized that their performance was not much affected by the turbulence (difference = .14), but they nevertheless felt out of control (difference = .28). The difference of differences was significant, $t(36) = 2.87$, $p < .01$, $r^2 = .19$. This interaction is shown in Figure 4A.

The interaction between magic and judgment type was also significant, $F(1, 36) = 32.80$, $MSE = 0.02$, $p < .001$, $\eta_p^2 = .48$, such that people realized that their performance was considerably higher in the magic condition (JoP $M = .54$), as compared with the no magic condition (JoP $M = .35$), but their JoAs did not reflect this difference in perceived performance, (JoA magic $M = .47$; JoA no magic $M = .40$). This difference of differences was significant, $t(36) = 5.67$, $p < .001$, $r^2 = .47$. People recognized the increased performance but did not attribute it to their own actions. This is shown in Figure 4B.

The interaction between speed and judgment was not significant ($t < 1$, $p > .1$). The interaction between feedback and judgment was also not significant ($t < 1$, $p > .1$). Of the remaining 12 possible interactions with judgment type, one was significant: Turbulence × Magic × Speed × Judgment Type, $F(1, 36) = 4.22$, $MSE = 0.01$, $p = .047$, $\eta_p^2 = .11$. This small quadruple interaction is very difficult to parse. It may be due to a slight tendency for the two significant double interactions between judgment type with magic (see Figure 4A) and judgment type with turbulence (see Figure 4B) to be very slightly more prominent at a slow speed than at a fast speed. The cell means and their standard deviations are given in Appendix C.

Finally, we once again computed the gammas and the Pearson's correlations to examine the relation between the judgments and people's performance. We expected that these correlations would be higher when people were making JoPs than when they were making JoAs. The Pearson's correlation for the JoPs was .61; for the JoAs, it was .41. These were significantly different from one another, $t(36) = 4.66$, $p < .001$, $r^2 = .38$. The gammas showed a similar result. Gamma was .46 for the JoPs, whereas, for the JoAs, it was only .29, $t(36) = 4.62$, $p < .001$, $r^2 = .37$. These correlations indicated that people were doing something differently when

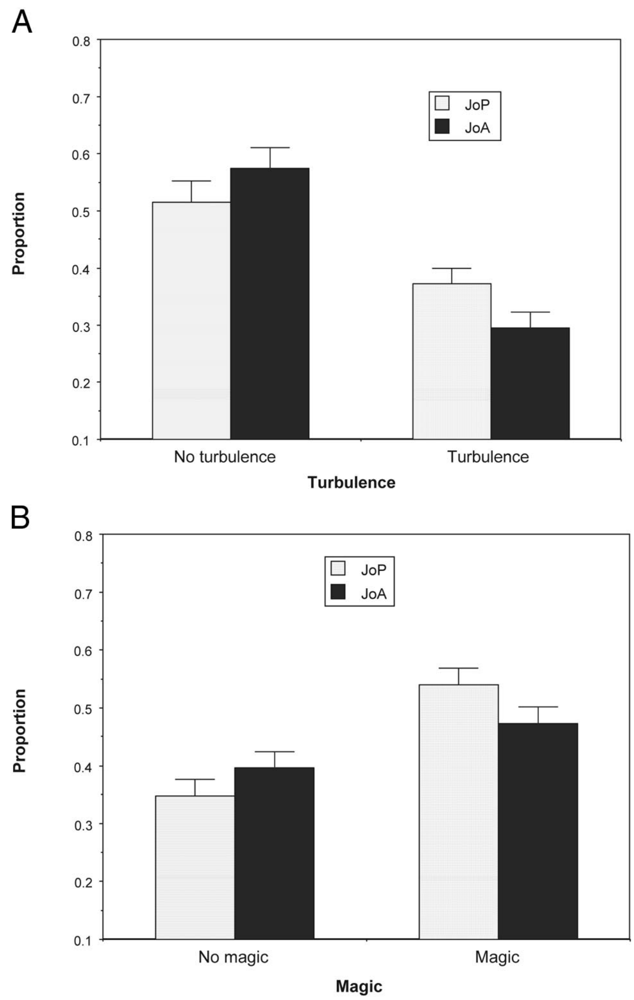

Figure 4. The relation between judgments of agency (JoAs) and judgments of performance (JoPs) in Experiment 3. A: This panel shows the interaction with turbulence. B: This panel shows the interaction with magic. The errors bars indicate standard errors of the mean.

they made JoAs compared with when they made JoPs. However, the fact that the correlations between JoAs and performance were significantly greater than zero, $t(18) = 13.85$, $p < .001$, $r^2 = .91$, for Pearson's *r,* and $t(18) = 13.10$, $p < .001$, $r^2 = .91$, for gammas, indicated that although performance was not the only factor contributing to JoAs and, indeed, that it did not contribute as much as it did to JoPs, it nevertheless could not be ruled out as having some impact on JoAs.

Experiment 4

Experiment 4 was similar to Experiment 3, except that it was conducted within participants, so that each person made JoAs for the trials in one half of the experiment and JoPs for the trials in the other half, with the order counterbalanced. The judgment type was blocked so that people would not become confused about which judgment they were supposed to be doing. We also gave a spoken reminder (given by the computer) before each judgment type, to emphasize the judgments that were requested during each half of the experiment. Because we wanted to complete the experiment within a single session, we eliminated the speed variable, which was not, in any event, distinguishing between JoAs and JoPs. For the same reason we eliminated the feedback variable, always giving auditory feedback. But we included an additional magic condition, a so-called "bad magic" condition, in which the nontargeted items (the Os) that the participants were told to avoid hitting would disappear (and softly thud) as if contacted, even though the person's cursor was still rather far from the item.

Method

Participants were 24 Columbia University and Barnard college students, 12 of whom were female and 12 of whom were male, who participated for course credit or pay and who were treated in accordance with the ethical guidelines of the American Psychological Association. The design of the experiment was 2 (type of judgment: JoA or JoP) × 2 (turbulence: turbulent or not turbulent) × 3 (magic: good magic, no magic or bad magic). The good magic condition was the same as described in the previous experiments: People were given credit for a hit as long as they were within 10 pixels of the Xs, and there was no change in the false alarms, that is, they had to be exactly on the Os to be docked for a false alarm. The bad magic condition was the reverse: People were only given credit for a hit if they were precisely on the Xs but they were docked with a false alarm if they were within 10 pixels of the Os.

Results and Discussion

The hit rates did not show a significant main effect of judgment condition in this experiment ($F < 1$, $p > .1$). Because this difference was significant in the previous between-participants experiment, we also analyzed just the first half of the experiment, at which point people had only been exposed to one or the other judgment condition—which amounts to a between-participants design. Analyzed this way, there was still no effect of judgment type on overall hit rate, $t(22) = 1.57$, $p > .1$. There was an effect of turbulence, such that the nonturbulent condition had a higher hit rate than did the turbulent condition, $F(1, 23) = 215.95$, $MSE = 0.01$, $p < .001$, $\eta_p^2 = .90$. There was an effect of magic on hit rate, such that the good magic condition was higher than the no magic condition and the bad magic and no magic conditions were the same, $F(2, 23) = 606.99$, $MSE = 0.01$, $p < .001$, $\eta_p^2 = .96$, as expected. Of course, when we investigated the false alarm rate—which was what the bad magic targeted—there was a difference among conditions, $F(2, 23) = 417.46$, $MSE = 0.01$, $p < .001$, $\eta_p^2 = .95$, such that the good magic and no magic conditions were the same, $t(23) = 1.65$, $p > .1$, and lower than the bad magic conditions, $t(23) = 21.65$, $p < .001$, $\eta_p^2 = .96$, indicating that the bad magic did, indeed, have the deleterious effect on performance that was expected. There was an interaction, with hit rate as the dependent variable, between turbulence and magic, $F(2, 23) = 67.91$, $MSE = 0.01$, $p < .001$, $\eta_p^2 = .75$. There were no other significant effects on hit rate. The cell means for hits, JoAs, and JoPs, as well as their standard deviations, are given in Appendix D.

We conducted an ANOVA to investigate the primary interest in this experiment, namely, differences between the two manipulated judgment types. The interaction between turbulence and judgment type was significant, $F(1, 23) = 24.98$, $MSE = 0.02$, $p < .001$, $\eta_p^2 = .52$. Under conditions of turbulence, as compared with no turbulence, people's JoPs changed less than did their JoAs: People did not distinguish between JoAs and JoPs in the no turbulence condition (difference = .05, which is not different from zero), $t(23) = 1.60$, $p > .1$, but did in the turbulence condition (difference = .13), $t(23) = 3.97$, $p < .001$, $r^2 = .41$. This interaction is shown in Figure 5A.

The interaction between magic and judgment type was significant, $F(2, 23) = 28.81$, $MSE = 0.01$, $p < .001$, $\eta_p^2 = .56$, such that people realized that their performance was considerably higher in the good magic condition (JoP $M = .68$), as compared with the no magic condition (JoP $M = .43$), and in their JoPs, they recognized that they were not performing well in the bad magic condition (.30), but their JoAs did not reflect this difference in perceived performance, (JoA: for good magic = .52, for no magic = .41, and for bad magic = .36). This interaction is shown in Figure 5B. No other effects or interactions with judgment type were significant.

As in the previous experiments, people recognized the increased performance with good magic but did not fully attribute it to their own agency. It is interesting that unlike in the turbulence condition, which lowered performance and which in this and the previous experiment resulted in JoAs that were even lower than the corresponding JoPs, the bad magic condition resulted in JoAs that, although lower than those in the no magic conditions, were slightly higher than were people's JoPs. One possible reason for this difference might be that people did not put much stock in false alarms, as far as their assessments of their own agency was concerned—they focused on hitting the targets rather than avoiding the nontargets. The turbulence condition affected performance both by decreasing the hit rate and by increasing the false alarm rate, whereas the bad magic had its effect only on the false alarm rate. To further investigate this possible underweighting of false alarms in people's assessments, we conducted two multiple regressions on the data from this experiment, extracting beta weights for each participant. In the first, which regressed hits and false alarms as predictors of JoPs, we found that the standardized beta weights for hits ($M = .58$), $t(23) = 17.84$, $p < .001$, $r^2 = .93$, were indeed considerably higher, $t(23) = 3.11$, $p < .01$, $r^2 = .30$, than were those for false alarms ($M = -.38$), $t(23) = 10.61$, $p < .001$, $r^2 = .83$. The overall R^2 was .70. The multiple regression directed at JoAs showed a much lower overall R^2 of .38. This difference, $t(23) = 9.36$, $p < .001$, $r^2 = .79$, itself bolsters the idea that people were indeed relying somewhat on their performance in making JoAs but much less so than when they made JoPs. The standardized beta coefficients directed at JoAs were .41 for hits, $t(23) = 11.65$, $p < .001$, $r^2 = .86$, and $-.27$ for false alarms, $t(23) = 7.92$, $p < .001$, $r^2 = .68$, again indicating less reliance on false alarms.

Finally, we once again computed the gamma and the Pearson's correlations to examine the relation between the judgments and people's performance (hit rate) and to allow comparison with previous experiments. As with the multiple regressions, we expected that these correlations should be higher when people were making JoPs than when they were making JoAs. The Pearson's correlation for the JoAs was .52; for the JoPs, it was .75. These were significantly different from one another, $t(23) = 6.79$, $p < .001$, $r^2 = .67$. The gammas showed a similar result. It was only .41 for the JoAs, whereas for the JoPs, it was .60, $t(23) = 6.00$, $p < .001$, $r^2 = .61$. These correlations indicate that people were doing something differently when they made JoAs as compared with when they made JoPs. However, the fact that the correlations between JoAs and performance were significantly greater than zero, $t(23) = 17.03$, $p < .001$, $r^2 = .93$, for Pearson's r against zero; $t(23) = 17.28$, $p < .001$, $r^2 = .93$, for gammas against zero, indicated that although performance was not the only factor contributing to JoAs and, indeed, that it did not contribute as much as it did to JoPs, it nevertheless could not be ruled out as having some impact on JoAs.

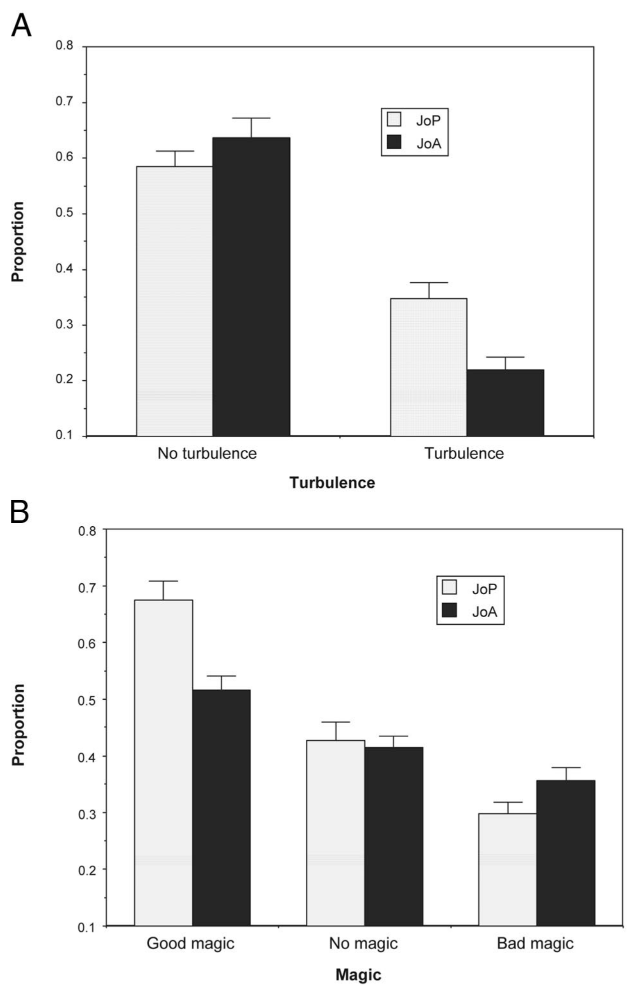

Figure 5. The relation between judgments of agency (JoAs) and judgments of performance (JoPs) in Experiment 4. A: This panel shows the interaction with turbulence. B: This panel shows the interaction with magic. The error bars indicate standard errors of the mean.

Conclusion

These experiments confirmed that people are able to metacognitively appreciate their own agency in a manner that is both well disciplined and that conforms closely to external parameters, tightly related to their real control. When in situations in which distortions of control were minimal, people's JoAs largely tracked their JoPs. If what happened corresponded to what the individual intended to happen, then the grounds for agency were fulfilled. On the assumption that people intended to do perfectly, their JoAs corresponded remarkably well to their metacognitions of how well they had actually performed. When they thought they had performed well, their judgments, by and large, reflected that performance. When they thought they had performed poorly, their judgments, again, tracked performance. Interestingly enough, when feedback made the outcome of the person's actions more salient—as this formulation would suppose—people's JoAs tracked their performance more closely. When different speeds resulted in distorted JoPs, JoAs indistinguishably tracked the biased JoPs. Furthermore, there was a correlation in all of the experiments between people's JoAs and their performance. Thus, these experiments revealed that one factor that contributes to people's JoAs is people's assessment of their performance and that this, in turn, is closely but not invariably related to their performance itself.

However, the notion that JoAs are based on perceptions of outcome alone is not the whole story nor is it the most interesting part of the story told by the present experiments. When, in these experiments, there really were distortions in the relation between the person's motor movements and what happened to either the cursor position (as when turbulence was introduced) or the cursor effect (as when hits or false alarms occurred by magic when the target had not really been contacted), people no longer relied exclusively on their perception of the outcome. Instead, they were able to monitor their own actions and modified their JoAs appropriately. They did not perceive a high level of control when the mouse was turbulent, even though their performance did not mirror the lack of control. When things went wrong by magic, they deflated their judgments of control. They also did not take credit spuriously for the inflated outcomes that occurred by good magic. They knew that it was not they who had been responsible for the favorable outcome.

Such appropriate denials of control seem consistent with the proprioceptive feed forward models that have been used to explain pathological disturbances of feelings of agency. However, in our cases, the disturbances were not pathological. Instead, the discrepancies were real between what the purported motor plan was and what the proprioceptive feedback indicated: There was something going on in the environment that decreased the actor's control, and those actors responded appropriately to it by modulating their JoAs. People were correct to feel less in control when turbulence was introduced into the relation between the cursor placement and its effect. They were also correct to deny themselves agentic credit, even for the excellent outcomes (which they were well able to recognize in their JoPs) when those outcomes occurred by magic rather than by their own precision.

The principled manner, revealed in this article, by which people made their JoAs indicates that to a first approximation, people are beautifully sensitive to the kinds of variables that they should be responsive to in agency monitoring. Many more empirical relations between people's JoAs and the factors that alter these judgments will need to be explored and elaborated in detail if we are to understand the underlying mechanisms by which people make these judgments, judgments that are so central to our sense of ourselves. The contribution of these first experiments lies in the small beginning that they make to chart out the fascinating empirical domain of how we know that we are, ourselves, intentionally making things happen by our own actions.

References

Bandura, A. (2001). Social cognitive theory: An agentic perspective. *Annual Review of Psychology, 52,* 1–26.

Blakemore, S.-J. (2003). Deluding the motor system. *Consciousness and Cognition, 12,* 647–655.

Blakemore, S.-J., & Frith, C. (2003). Self-awareness and action. *Current Opinion in Neurobiology, 13,* 219–224.

Blakemore, S.-J., Wolpert, D. M., & Frith, C. (1998). Central cancellation of self-produced tickle sensation. *Nature Neuroscience, 1,* 635–641.

Blakemore, S.-J., Wolpert, D. M., & Frith, C. (2002). Abnormalities in the awareness of action. *Trends in Cognitive Science, 6,* 237–242.

Clore, G. L. (1992). Cognitive phenomenology: Feelings and the construction of judgment. In L. L. Martin & A. Tesser (Eds.), *The construction of social judgments* (pp. 133–163). Hillsdale, NJ: Erlbaum.

Daprati, E., Franck, N., Georgieff, N., Proust, J., Pacherie, E., Dalery, J., & Jeannerod, M. (1997). Looking for the agent: An investigation into consciousness of action and self-consciousness in schizophrenic patients. *Cognition, 65,* 71–86.

Dunlosky, J., & Nelson, T. O. (1992). Importance of kind of cue for judgments of learning (JOL) and the delayed-JOL effect. *Memory & Cognition, 20,* 374–380.

Farrer, C., Franck, N., Gerogeiff, N., Frith, C. D., Decety, A., & Jeannerod, M. (2003). Modulating the experience of agency: A positron emission tomography study. *NeuroImage, 18,* 324–333.

Farrer, C., & Frith, C. D. (2002). Experiencing oneself versus another person as being the cause of an action: The neural correlates of the experience of agency. *NeuroImage, 15,* 596–603.

Flach, R., Knoblich, G., & Prinz, W. (2004). Recognizing one's own clapping: The role of temporal cues. *Psychological Research, 69,* 147–156.

Fourneret, P., & Jeannerod, M. (1998). Limited conscious monitoring of motor performance in normal subjects. *Neuropsychologia, 36,* 1133–1140.

Frith, C. D. (1992.) *The Cognitive Neuropsychology of Schizophrenia.* Hillsdale, NJ: Erlbaum.

Gallagher, S. (2004). Neurocognitive models of schizophrenia: A neurophenomenological critique. *Psychopathology, 37,* 8–19.

Gallese, V., & Goldman, A. (1998). Mirror neurons and the simulation theory of mind-reading. *Trends in Cognitive Science, 2,* 493–501.

Haggard, P. (2005). Conscious intention and motor cognition. *Trends in Cognitive Sciences, 9,* 290–295.

Haggard, P., Clark, S., & Kalogeras, J. (2002). Voluntary action and conscious awareness. *Nature Neuroscience, 5,* 382–385.

Hertzog, C., & Hultsch, D. F. (2000). Metacognition in adulthood and aging. In T. A. Salthouse & F. I. M. Craik (Eds.), *Handbook of aging and cognition* (pp. 417–466). Mahwah, NJ: Erlbaum.

Jahanshahi, M., & Frith, C. D. (1998). Willed action and its impairments. *Cognitive Neuropsychology, 15,* 483–533.

Jeannerod, M. (2001). Neural simulation of action: A unifying mechanism for motor cognition. *NeuroImage, 14,* S103–S109.

Jordan, M. L., & Rumelhart, D. E. (1992). Forward models: Supervised learning with a distal teacher. *Cognitive Science, 16,* 307–354.

Kawato, M., Furukawa, K., & Suzuki, R. (1987). A hierarchical neural network model for control and learning of voluntary movement. *Biological Cybernetics, 57,* 169–186.

Kinsbourne, M. (1995). Awareness of one's own body: An attentional theory of its nature, development, and brain basis. In J. Berm'dez, N. Eilan, & A. Marcel (Eds.), *The body and the self* (pp. 206–223). Cambridge, MA: MIT Press.

Knoblich, G., & Flach, R. (2001). Predicting the effects of actions: Interactions of perception and action. *Psychological Science, 12,* 467–472.

Knoblich, G., & Prinz, W. (2001). Recognition of self-generated actions from kinematic displays of drawing. *Journal of Experiment Psychology: Human Perception and Performance, 27,* 456–465.

Knoblich, G., & Sebanz, N. (2005). Agency in the face of error. *Trends in Cognitive Science, 9,* 259–261.

Knoblich, G., Stottmeister, F., & Kircher, T. (2004). Self-monitoring in patients with schizophrenia. *Psychological Medicine, 34,* 1561–1569.

Koriat, A. (1997). Monitoring one's own knowledge during study: A cue-utilization approach to judgment of learning. *Journal of Experimental Psychology: General, 126,* 349–370.

Koriat, A., & Bjork, R. A. (2005). Illusions of competence in monitoring one's knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31,* 187–194.

Lau, H. C., Rogers, R. D., & Haggard, P. (2004, February 20). Attention to intention. *Science, 303,* 1208–1210.

Leube, D. T., Knoblich, G., Erb, M., Grodd, W., Bartels, M., & Kircher, T. T. J. (2003). The neural correlates of perceiving one's own movements. *NeuroImage, 20,* 2084–2090.

Leube, D. T., Knoblich, G., Erb, M., & Kircher, T. T. J. (2003). Observing one's hand become anarchic: An fMRI study of action identification. *Consciousness and Cognition, 12,* 597–608.

Libet, B. (2004). *Mind time.* Cambridge, MA: Harvard University Press.

Libet, B., Gleason, C. A., Wright, E. W., & Pearl, D. K. (1983). Time of conscious intention to act in relation to onset of cerebral activities (readiness-potential): The unconscious initiation of a freely voluntary act. *Brain, 106,* 623–642.

Metcalfe, J. (1993). Novelty monitoring, metacognition, and control in a composite holographic associative recall model: Implications for Korsakoff amnesia. *Psychological Review, 100,* 3–22.

Metcalfe, J., & Kornell, N. (2005). A region of proximal learning model of study time allocation. *Journal of Memory and Language, 52,* 463–477.

Metzinger, T. (2003). Conscious volition and mental representation: Toward a more fine-grained analysis. *Consciousness and Cognition, 4,* 1–52.

Miall, R. C., & Wolpert, D. M. (1996). Forward models for physiological motor control. *Neural Networks, 9,* 1265–1279.

Morsella, E. (2005). The function of phenomenal states: Supramodal interaction theory. *Psychological Review, 112,* 1000–1021.

Nelson, T. O., & Narens, L. (1994). Why investigate metacognition? In J. Metcalfe & A. P. Shimamura (Eds.), *Metacognition: Knowing about knowing* (pp. 1–25). Cambridge, MA: MIT Press.

Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review, 84,* 231–259.

Repp, B. H., & Knoblich, G. (2004). Perceiving action identity: How pianists recognize their own performances. *Psychological Science, 15,* 604–609.

Sato, A., & Yasuda, A. (2005). Illusion of self-agency: Discrepancy between predicted and actual sensory consequences of actions modulates the sense of self-agency, but not the sense of self-ownership. *Cognition, 94,* 241–255.

Schwartz, B. L., & Smith, S. (1997). The retrieval of related information influences tip-of-the-tongue states. *Journal of Memory and Language, 36,* 68–86.

Sebanz, N., & Prinz, W. (2006). *Disorders of volition.* Cambridge, MA: MIT Press.

Son, L. (2004). Spacing one's study: Evidence for a metacognitive control strategy. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30,* 601–604.

Thiede, K. W. (1999). The importance of accurate monitoring and effective self-regulation during multi-trial learning. *Psychonomic Bulletin & Review, 6,* 662–667.

Thiede, K. W., & Dunlosky, J. (1999). Toward a general model of self-regulated study: An analysis of selection of items for study and a self-paced study time. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24,* 1024–1037.

Tsakiris, M., Haggard, P., Franck, N., Mainly, N., & Sirigu, A. (2005). A specific role for efferent information in self-recognition. *Cognition, 96,* 215–231.

van den Bos, E., & Jeannerod, M. (2002). Sense of body and sense of action both contribute to self-recognition. *Cognition, 85,* 177–187.

Wegner, D. M. (2002). *The illusion of conscious will.* Cambridge, MA: MIT Press.

Wegner, D. M. (2003). The mind's best trick: How we experience conscious will. *Trends in Cognitive Sciences, 7,* 65–69.

Wegner, D. M. (2004). Précis of the illusion of conscious will. *Behavioral and Brain Sciences, 27,* 649–692.

Wegner, D. M. (2005). Who is the controller of controlled processes? In R. R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The new unconscious.* New York: Oxford University Press.

Wegner, D. M., Sparrow, B., & Winerman, L. (2004). Vicarious agency: Experiencing control over the movements of others. *Journal of Personality and Social Psychology, 86,* 838–848.

Wegner, D. M., & Wheatly, T. (1999). Apparent mental causation: Sources of the experience of will. *American Psychologist,* 481–492.

Wolpert, D. M. (1997). Computational approaches to motor control. *Trends in Cognitive Sciences, 1,* 209–216.

Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995, September 29). An internal model for sensorimotor integration. *Science, 269,* 1880–1882.

Appendix A

Cell Means and Standard Deviations From Experiment 1

	No turbulence				Turbulence			
	Fast		Slow		Fast		Slow	
Measure	1:1 target density	7:1 target density	1:1 target density	7:1 target density	1:1 target density	7:1 target density	1:1 target density	7:1 target density
Hit rate								
M	.45	.43	.76	.79	.39	.38	.70	.71
SD	.12	.13	.20	.20	.13	.11	.20	.17
JoA								
M	.46	.48	.59	.64	.30	.33	.39	.42
SD	.22	.21	.22	.24	.20	.21	.24	.25

Note. JoA = judgment of agency.

Appendix B

Cell Means and Standard Deviations From Experiment 2

	No magic				Magic			
	Feedback		No feedback		Feedback		No feedback	
Measure	Slow	Fast	Slow	Fast	Slow	Fast	Slow	Fast
No turbulence								
Hit rate								
M	.67	.38	.68	.39	.80	.70	.84	.68
SD	.30	.20	.27	.19	.25	.19	.22	.20
JoA								
M	.52	.49	.49	.44	.58	.57	.55	.51
SD	.29	.28	.30	.28	.30	.30	.30	.29
Turbulence								
Hit Rate								
M	.54	.26	.55	.27	.81	.60	.79	.60
SD	.24	.18	.22	.18	.23	.16	.24	.16
JoA								
M	.25	.23	.22	.20	.31	.35	.28	.30
SD	.22	.23	.22	.21	.24	.26	.23	.26

Note. JoA = judgment of agency.

(*Appendixes continue*)

Appendix C

Cell Means and Standard Deviations From Experiment 3

Measure	No magic				Magic			
	Feedback		No feedback		Feedback		No feedback	
	Slow	Fast	Slow	Fast	Slow	Fast	Slow	Fast
No turbulence								
JoA condition								
Hit rate								
M	.76	.40	.70	.43	.93	.76	.91	.76
SD	.23	.19	.23	.17	.11	.16	.13	.15
JoA								
M	.61	.52	.57	.49	.60	.65	.57	.59
SD	.25	.24	.28	.23	.23	.23	.25	.24
JoP condition								
Hit rate								
M	.62	.36	.65	.35	.88	.69	.87	.69
SD	.25	.17	.26	.15	.16	.18	.18	.17
JoP								
M	.48	.39	.46	.38	.61	.63	.60	.56
SD	.26	.19	.26	.19	.22	.20	.26	.21
Turbulence								
JoA condition								
Hit rate								
M	.58	.29	.59	.27	.87	.68	.89	.64
SD	.22	.12	.24	.12	.18	.17	.14	.18
JoA								
M	.27	.26	.26	.20	.33	.42	.32	.29
SD	.19	.19	.21	.16	.22	.24	.22	.21
JoP condition								
Hit rate								
M	.52	.24	.47	.23	.80	.64	.78	.60
SD	.24	.12	.24	.14	.21	.19	.22	.19
JoP								
M	.30	.26	.29	.23	.48	.52	.42	.48
SD	.19	.18	.22	.17	.22	.22	.21	.22

Note. JoA = judgment of agency; JoP = judgment of performance.

Appendix D

Cell Means and Standard Deviations From Experiment 4

	No turbulence			Turbulence		
Measure	Bad magic	No magic	Good magic	Bad magic	No magic	Good magic
JoA condition						
Hit rate						
M	.56	.58	.85	.34	.35	.81
SD	.17	.17	.14	.16	.16	.17
False alarms						
M	.30	.05	.04	.53	.16	.13
SD	.17	.07	.07	.19	.11	.10
JoA						
M	.55	.63	.73	.16	.20	.30
SD	.25	.25	.21	.14	.16	.20
JoP condition						
Hit rate						
M	.60	.58	.87	.31	.35	.82
SD	.18	.15	.14	.16	.16	.14
False alarms						
M	.29	.04	.03	.48	.13	.13
SD	.18	.06	.05	.19	.11	.11
JoP						
M	.42	.56	.77	.17	.30	.58
SD	.21	.22	.21	.16	.23	.25

Note. JoA = judgment of agency; JoP = judgment of performance.

Received June 23, 2006
Revision received October 26, 2006
Accepted November 7, 2006 ■