

Ability, chance, and ambiguity aversion: Revisiting the competence hypothesis

William M. P. Klein*, Jennifer L. Cerully, Matthew M. Monin
Department of Psychology, University of Pittsburgh

Don A. Moore
Tepper School of Business, Carnegie Mellon University

Abstract

Individuals are often ambiguity-averse when choosing among purely chance-based prospects (Ellsberg, 1961). However, they often prefer apparently ambiguous *ability*-based prospects to unambiguous chance-based prospects. According to the competence hypothesis (Heath & Tversky, 1991), this pattern derives from favorable perceptions of one's competence. In most past tests of the competence hypothesis, ambiguity is confounded with personal controllability and the source of the ambiguity (e.g., chance vs. missing information). We unconfound these factors in three experiments and find strong evidence for independent effects of both ambiguity aversion and competence. In Experiment 1, participants preferred an unambiguous chance-based option to an ambiguous ability-based option when the ambiguity derived from chance rather than uncertainty about one's own ability. In Experiments 2 and 3, which used different operationalizations of ambiguity in choice contexts with actual consequences, participants attempted to avoid both ambiguity and chance insofar as they could. These findings support and extend the competence hypothesis by demonstrating ambiguity aversion independent of personal controllability and source of ambiguity.

Keywords: ambiguity, control, competence, choice.

1 Introduction

Folk wisdom advises, “Better the devil you know than the devil you don't.” Most people abide by this logic when choosing between options with known probabilities and others that have ambiguous probabilities (or second-order uncertainty). People often choose the former – manifesting what decision scientists call “ambiguity aversion.” In a classic demonstration by Ellsberg (1961), decision makers faced a choice between two urns; one contained 50% red and 50% black marbles whereas the other had an unknown distribution. Selection of a marble of a specified color produced a monetary prize. Participants showed no color preference, but they did show a clear preference for the 50/50 urn. This preference is hard to justify on rational grounds because the ambiguous urn cannot contain a lower proportion of both red *and* black marbles. Later work showed that people prefer the unambiguous

urn even when the expected value of the other (ambiguous) urn is higher (Keren & Gervitsen, 1999), and that they are willing to pay more (an “ambiguity premium”) to draw from the 50/50 urn (Becker & Brownson, 1964).

Several persuasive explanations have been offered to account for ambiguity aversion. Frisch and Baron (1988) proposed that decision makers are particularly averse to missing information, especially if that information can be known or manipulated by others (such as an opponent) or may become available at some later point in the decision process. These authors also suggest that people are concerned about the blame and regret that may follow a decision made in the context of missing information. As an example, parents are relatively less willing to vaccinate their children in a setting where the characteristics of children who experience side effects from the vaccination are unspecified (Ritov & Baron, 1990). Similarly, Heath and Tversky (1991) argue that participants in the Ellsberg paradigm mentally simulate the attributions that will be made following selection of the marble, and consequently worry about being blamed if the ambiguous urn turns out to present a smaller chance of winning. Indeed, Curley, Yates, and Abrams (1986) observed that ambiguity aversion is greater when the contents of the ambiguous urn will be revealed to others.

*This research was supported by a National Science Foundation grant (SES-0648044) to the first author. We thank Dana Aravich, Colin Brennan, Lee Ellis, Thomas Fuell, Victoria Prescott, and Katherine Smith for assistance with data collection. Correspondence may be sent to the first author who is now at Behavioral Research Program, Division of Cancer Control and Population Sciences, National Cancer Institute, 6130 Executive Boulevard, EPN Room 4060, Bethesda, MD 20892–7326. E-mail: kleinwm@mail.nih.gov.

An important feature of the Ellsberg paradigm is that the contents of the ambiguous urn are knowable — just unknown to the participants. The outcome of the decision also hinges on chance and is not in any way tethered to participants' own knowledge or abilities. Heath and Tversky (1991) suggest that if people feel competent in the decision domain — and if competence is related to the outcome — they should exhibit less ambiguity aversion. In an illustrative study, participants predicted outcomes of American football games over a five-week period and also assessed their own knowledge about each game. They then chose between betting on the team they predicted would win or a lottery in which the chances of winning were identical but determined by chance. As long as participants perceived the probability of winning to be .50 or higher, those with higher perceived knowledge chose the bet over the lottery, suggesting that perceptions of competence in the domain attenuated ambiguity aversion. Similar findings emerged in contexts where voters made predictions about outcomes in the November 1988 presidential election and where students made predictions about future news events (see also Chow & Sarin, 2001; Fox & Tversky, 1995; Fox & Weber, 2002; Grieco & Hogarth, 2004; Hogarth & Kunreuther, 1989; Yates & Zukowski, 1976; for review, see Camerer & Weber, 1992). Competence in the decision domain likely reduced the perception of missing information that Frisch and Baron posited to be key to ambiguity aversion. Thus, Heath and Tversky proposed that perceived competence affects perceived ambiguity and does not have an independent effect of its own.

However, other characteristics of this decision paradigm suggest that competence has an independent effect. Most importantly, people generally prefer controllable to uncontrollable outcomes (e.g., Goodie, 2003; Howell, 1971; Klein & Kunda, 1994), and are more optimistic about the outcomes of controllable outcomes (Harris, 1996). In most (though not all) tests of the competence hypothesis, the unambiguous option is chance-determined, meaning that the manipulation of competence may be confounded with outcome controllability. If the probabilities of the skill-based option were ambiguous, participants might still choose the skill option, suggesting they prefer skill over chance, enough to override any opposing effect of ambiguity. Controllability, rather than ambiguity, may also account for the choice of betting on familiar vs. unfamiliar domains of competence (e.g., weather in one's own geographical area vs. weather in a distal unknown geographical area), given that familiarity or proximity can be conflated with controllability due to magical thinking (e.g., Brun & Tieggen, 1990).

Another possibility in the current context is that people are responding not to competence or controllability per se but rather to the reason for the ambiguity. That

is, ambiguity could have different effects depending on its source. In tests of the competence hypothesis, the ability-based option is ambiguous because it is not associated with a clearly stated probability of winning (as there is for the chance option). The ambiguity is instead related to natural variation in competence and perceptions of competence, thereby confounding the source of ambiguity with competence. If the ambiguity associated with the ability-based option derived from factors unrelated to personal competence, such as chance, it is likely that the ability-based option would become less appealing. Above we suggested that chance-based options in tests of the competence hypothesis should possess ambiguity (i.e., a range of probabilities rather than a point estimate), and here we take this one step further by arguing that the competence-based options should possess the same kind of ambiguity (untethered to competence).

Finally, previous experiments offered competence options that were based on *a priori* beliefs, limiting the way competence was conceptualized. In no cases was competence manipulated — in particular, by giving participants feedback in the competence domain — to more directly test the role of competence in choices involving ambiguity.

We report three experiments that systematically separate ambiguity from controllability and the source of ambiguity in order to determine whether people continue to exhibit ambiguity aversion as predicted by the competence hypothesis. In Experiment 1, participants face a choice between an unambiguous chance-based prospect and an ambiguous ability-based prospect. Importantly, the latter prospect is ambiguous because its difficulty is determined by chance. We predicted that decision-makers would prefer the chance-based option — contrary to the typical pattern in tests of the competence hypothesis but nevertheless illustrative of ambiguity aversion.

Experiment 2 builds on this study by presenting participants with a chance option, a low-ambiguity ability-based option (with approximately a 60% chance of winning), and a high-ambiguity ability-based option (with win probabilities ranging from 20%-100%). Because the ambiguity inherent in the third option is due to chance (and because the first option is entirely chance-based), we expect participants to prefer the low-ambiguity ability-based option over the other two — demonstrating preferences for controllability and against ambiguity in the same context. In this experiment, we also give participants performance feedback upon which to base their choices rather than relying on *a priori* beliefs about competence. Experiment 2 builds on the standard competence hypothesis paradigm by pitting two ability-based options against each other while independently manipulating the level of ambiguity. Finally, we ask participants to make the same choice for someone else, thereby completely

separating ambiguity from controllability. If ambiguity aversion is independent of controllability, then choices for other people should mirror those for the self.

Experiment 3 tests a similar hypothesis using a different operationalization of ambiguity which again is unaffected by personal controllability. In this case, we give participants a choice among an unambiguous chance-based option, an ability-based option where they must outperform a *low ambiguity* opponent, and an ability-based option where they must instead outperform a *high ambiguity* opponent. As in Experiment 2, we expect people to prefer the ability-based options but to show ambiguity aversion by selecting the option in which they compete against the low ambiguity opponent.

An additional characteristic of previous studies in this literature is that the probability of winning the chance-based option was always unambiguous (i.e., a fixed probability). In Experiments 2 and 3, we also introduce ambiguity to the chance-based option in the form of a range of probabilities. Doing so provides an even more conservative test of whether preference for ability-based tasks in previous tests was due to ambiguity or simply a preference for ability-based over chance-based tasks. Moreover, by defining ambiguity in terms of the opponent in Experiment 3, we have yet again disentangled ambiguity from personal competence, and in a different manner than in Experiment 2 where we ask participants to make choices both for themselves and others.

In sum, we report three experiments designed to test and extend the competence hypothesis. Studies include both hypothetical and actual choices (and in the final two experiments more than one ability-based choice). We expect that people will exhibit ambiguity aversion even when controllability is held constant. They should prefer chance-based options only when the ability options are more ambiguous, but when choosing among similar ability-based options they should prefer those lower in ambiguity. Moreover, we expect the competence hypothesis to stand even when the chance-based option is ambiguous, when the source of the ambiguity is unrelated to personal competence, and when personal competence itself originates from performance feedback rather than from *a priori* beliefs.

2 Experiment 1

In Experiment 1, participants faced a choice between an unambiguous chance-based option and an ambiguous ability-based option — mirroring choices made in previous tests of the competence hypothesis. Importantly, as described below, the ambiguity in the ability-based option is due to chance rather than variation in perceived competence. Consequently, we predict that people will

prefer the chance-based option — reflective of ambiguity aversion but contrary to the typical preference for competence-based options in past work.

2.1 Method

Participants were 47 University of Pittsburgh undergraduates fulfilling a course requirement. They were told to imagine that they had to choose between two games affording them the opportunity of winning a \$100 prize:

Option A: A computer will randomly choose a number between 1 and 10, and if the number is less than 8, you will win the \$100 prize. Because there are 7 numbers below 8 and therefore 7 ways to win, your chances of winning the prize are 7 in 10 or 70%.

Option B: You will take a 10-item quiz on pop culture (an example question is “Who is the host of the Tonight Show?”). If you answer all the questions correctly, you will win the \$100 prize. The quiz will be taken from one of two possible sets of items — Set 1 or Set 2. It turns out that about 50% of Pitt students are able to answer all the Set 1 questions correctly, meaning that your chances of winning the prize if you get these questions right are about 50%. On the other hand, 90% of Pitt students are able to answer all the questions in Set 2, meaning that your chances of winning the prize in this case are about 90%. Thus, your chances of winning the prize are approximately 50% or 90%, but you won’t know in advance which it is because you won’t know whether the computer will take questions from Set 1 or Set 2.

As in tests of the competence hypothesis, Option A was chance-based (and unambiguous) and Option B was ability-based (and ambiguous). The average probability of winning either game across all participants was the same — 70%. However, in this problem, the ambiguity in Option B derived from chance, not ability. Importantly, pretests showed that students in this population considered their knowledge of pop culture to be above average.

2.2 Results and discussion

As predicted, a majority of participants chose the chance-based option (70.2%, $n = 33$), $\chi^2[1] = 6.90$, $p = .01$. These findings extend the competence hypothesis by showing that people are less sanguine about ability-based ambiguous options when the associated ambiguity is due to chance. Of course, the reversal from findings testing the competence hypothesis could be due to some aspect of the ability-based task. Ideally we would have included control groups in which the chance option was unambiguous and the ability option ambiguous. We were also relying on perceptions of high personal competence (which seemed reasonable based on pretest findings) rather than providing performance feedback, and it is also likely that

some of the ambiguity in the ability-based task derived from perceptions of the task (e.g., participants may have wondered what counted as pop culture). We addressed these limitations in the next two experiments. Both included two ability-based options following the receipt of actual performance feedback — thereby providing a more realistic scenario and one where there was less ambiguity about the task itself. These experiments also involved games with real monetary payoffs.

3 Experiment 2

In Experiment 2, participants choose among three games — an ambiguous chance-based game and two ability-based games. The latter two games differed in level of ambiguity, and the ambiguity itself was due to chance. We predicted that participants would generally prefer the ability-based game with the lower level of ambiguity to the other two games.

3.1 Method

In this experiment, 153 undergraduates fulfilling a course requirement first completed ten knowledge tests in a variety of domains (e.g., geography, music, mathematics, and identifying emotions from facial expressions). Importantly, pre-testing demonstrated that participants viewed their knowledge in these domains to be above-average. They were then given false feedback on the ten tests (in terms of percentage correct) ranging from 20% to 100% with a mean performance score of 60%. Note that this feedback was not linked to specific tests, so participants did not know on which tests they performed well. After receiving the performance feedback, participants faced a choice among three options, each of which offered a prize of \$5 (modeled after games used by Klein & Kunda, 1994).

Option A was entirely chance-determined; the computer would randomly choose a number between 1 and 100 during each of 10 trials, and if the computer ever chose a number in a certain range, the game would end immediately and they would not win the prize. Participants further learned that the range would be anywhere from 1–3 to 1–7 (with the specific range chosen randomly by the computer). Thus, the chance of winning this game ranged from $.93^{10} = 48\%$ (in the event that the losing range was 1–7) to $.97^{10} = 73\%$ (in the event that the losing range was 1–3). The experimenter explained these statistics. If the computer *never* chose a number in the specified range, the participant would win the \$5 prize. Note that this is a chance-based option with roughly a 60% chance of winning, on average, but with a high degree of ambiguity resulting from the salient missing infor-

mation about the losing range. We designed it this way in order to add ambiguity to the chance task, thereby reducing the confound between ambiguity and chance/ability.

Participants also had two ability-based options. In Option B, the computer would choose ten questions, one from each of the knowledge tests participants had just completed. If a participant answered 60% of the items correctly, he or she would win a prize. In this case, because the prospects of winning depended upon performance on a task that participants believed they had already completed successfully, we expected it to be attractive to them. The experimenter made clear that the chances of winning this game were high given that they had already achieved an average score of 60% on the pretests. This was defined as the low-ambiguity ability-based option.

Option C was similar; in this case, participants were told that the computer would select ten questions from the knowledge tests they had completed, but that the selection procedure was random and could lead to an uneven number of questions from the various tests (e.g., multiple questions from one domain and none from other domains). Thus, their chances of winning Game C were explained (veridically) to be somewhere between their lowest and highest test scores (i.e., 20% and 100%). Again, participants needed to answer 60% correctly in order to win the prize. If participants perceived high competence and were optimistic, this would seem to be an attractive option because of the potentially high probability of winning. We did not expect participants to be optimistic, however, despite perceiving competence in the domain — precisely because the ambiguity associated with this option was due to chance.

Participants then chose one game to play and were also asked to indicate how they would choose for another person whom they were to assume had performed just as they had on the original tests. The latter estimate allowed us to determine whether people were making judgments based on performance alone rather than on *a priori* beliefs.

3.2 Results and discussion

Of the 153 participants, 24 (15%) chose Option A, 89 (58%) chose Option B, and 40 (26%) chose Option C. Thus, participants tended to prefer each of the ability-based games (B and C) to the chance-based game (A) ($\chi^2(1) = 70.70, p < .000001$; and $p = .030$ for C vs. A by a one-sided exact binomial test). More importantly, and as predicted, they preferred the ability-based game associated with lower ambiguity (B) to the one with higher ambiguity (C), ($\chi^2(1) = 17.86, p = .00002$; and $p = .026$ by a one-sided binomial test for A vs. B and C combined). The inclusion of three different options therefore tests several features of the competence hypothesis. Replicating past

work, individuals preferred the ability-based options (B and C) to the chance-based option A in a study where the latter option was associated with ambiguous probabilities — a comparison that would not have been possible without the inclusion of Option A. Moreover, participants choosing an ability option were more inclined to select the one with less ambiguity, demonstrating ambiguity aversion in a case where neither of the options was purely based on chance.

The pattern of choices made for another person was similar with 23% choosing A, 56% choosing B, and 22% choosing C — again showing an overall preference for the ability options over the chance option as well as a preference for the less ambiguous ability option. Thus, the pattern above cannot be easily attributed to *a priori* beliefs about the participants' own ability in the tested domains — participants appeared to be basing their choices on pretest performance as well as the nature of the ambiguity inherent in the various game options.

In sum, participants who were given feedback on an ability-based task preferred a game linked to performance on a similar task over a game that was chance-determined with ambiguous probabilities of winning. This pattern is consistent with ambiguity aversion, but in a setting where competence was not confounded by ambiguity or the source of ambiguity. Moreover, participants preferred an ability option with win probabilities that were relatively less ambiguous, again demonstrating ambiguity aversion. To our knowledge, this is the first study testing the competence hypothesis in a context where participants have more than one ability-based option from which to choose, and also where the probability of winning the chance-based game was ambiguous. This design allowed an examination of whether people still prefer ability-based options to chance-based games when ambiguity is present for both, and also helped to show that when the source of ambiguity is due to an uncontrollable source, people choosing an ability-based option still demonstrate ambiguity aversion.

4 Experiment 3

Experiment 3 was a conceptual replication of Experiment 2 using a different operationalization of ambiguity. Participants again were faced with three possible games; the first was chance-based and the others were ability-based. In this case, success at the ability-based games hinged on outperforming an “opponent” who was relatively high or low in ambiguity (the previous participant or the average participant in the study, respectively). By defining ambiguity based on the opponent, we once again intended to disentangle the source of ambiguity from perceptions of personal competence. This manner of manipulating am-

biguity also tied performance to something other than the ostensibly random choices of a computer, allaying concern that participants in the previous study preferred the low-ambiguity game because they somehow mistrusted the computer algorithm.

4.1 Method

Participants were 250 undergraduates fulfilling a course requirement. They completed ten trials of a spelling error detection task (designed by Klein & Kunda, 1994) in which they were presented with various passages containing multiple spelling errors. For each passage, participants needed to circle as many spelling errors as possible in one minute. Finally, participants chose among three games, each of which offered a prize of \$10.

Option A was again chance-based. In this game, the computer would randomly select a number between 1 and 100 during each of ten trials, and if the computer ever chose a number in the “losing range” the game would end without any winnings. The losing range was selected randomly by the computer and could be anywhere from 1–2 to 1–5. Thus, the chance of winning this game was anywhere between $.95^{10} = .60$ (in the event that the losing range was 1–5) and $.98^{10} = .80$ (in the event that the losing range was 1–2). If the computer *never* chose a number in the losing range, participants would win the \$10. The experimenter clearly explained these probabilities. As in Experiment 2, then, the chance-based option was associated with ambiguous chances of winning.

The other two games were ability-based. In Option B, the computer would again choose a number between 1 and 100 for ten trials. If the computer selected a number between 1 and 10, the participant would have to complete another trial of the spelling error detection task. In order for the game to continue, the score on this task would need to be *higher than the average (mean) score of the previous participant* in the study across all trials on the spelling error task; otherwise the game ended without any winnings. If the computer selected a number between 11 and 100, the game went on to the next of the ten trials and the participant did not have to take the test. The experimenter told participants that the calculated chance of winning this game was about .70 if their ability on the task was about the same as that of the previous participant, and higher than .70 if they performed better.¹ Again, if they completed all ten trials, they would win the \$10. Because participants had no way of knowing who the previous participant was, this game represented the high-ambiguity ability option.

¹The chance of not getting an above average score is $1 - a$, where a is the chance of meeting the performance standard. The chance of winning the entire game, then, is $[1 - .10(1 - a)]^{10}$. When $a = .64$, the ostensible average score, the chance is therefore .70.

Option C was identical to Option B with one exception — whenever participants had to complete the spelling error task again, they had to outperform *the average (mean) score for all previous participants* rather than simply the unknown previous participant. Because the group mean includes so much less variance than that of an individual and because participants' own scores are relevant to determining the group mean, Option C was less ambiguous than was Option B (see Dawes & Mulford, 1994; Krueger, 2003; Krueger & Clement, 1996), which should make it more attractive.

In order to provide further evidence that people were basing their choices on perceived competence (i.e., their own performances) rather than controllability per se, we also asked participants to estimate their scores on the spelling task prior to making their game choice. These measures are similar to the knowledge estimates Heath and Tversky (1991) asked their participants to provide; the authors used these estimates to show that preferences for ability-based games were directly associated with higher perceptions of competence in the domain. In the current study, however, performance was fixed, providing more control over the influence of prior beliefs.

4.2 Results and discussion

Game choices conformed to predictions. First, 157 (63%) preferred one of the ability-based games (B or C) as compared with 93 (37%) who selected the chance-based game (A), $\chi^2(1) = 15.88, p = .00007$. Thus, under conditions where both chance-based and ability-based options possessed ambiguous probabilities of winning, individuals expressed greater interest in the ability-based options — consistent with the competence hypothesis. Of course, given that there were three options, this pattern could have occurred by chance. The more important finding is that, among those participants who chose an ability-based game, they clearly preferred C over B (109 vs. 48; $\chi^2[1] = 22.92, p = .000002$). In short, individuals preferred that their performance depend on comparisons with a less ambiguous opponent (the average participant) — once again demonstrating ambiguity aversion. This experiment once again provides support for the competence hypothesis in a context where ambiguity and controllability are unfounded.

As expected, participants' estimates of their scores on the task were significantly related to their game choices, $OR = .94$ (95% CI: .89-.99), $p = .05$: individuals estimating higher performance were more likely to choose one of the ability-based options (B and C) than the chance-based option (A). Thus, choosing an ability-based game seemed to result, at least in part, from beliefs that one could perform well in such a game. This is notable because the competence hypothesis suggests that decision

makers prefer ability-driven alternatives irrespective of judged probabilities (Heath & Tversky, 1991, p. 7). The current findings suggest that preferences for ability-based options may in fact be dependent on such probabilities.

5 General discussion

Three experiments tell a consistent story about how people behave when making choices among options that vary in both controllability and ambiguity. In short, ambiguity aversion and the documented preference for controllable over uncontrollable options appear to be independent tendencies. These results also provide strong support for the competence hypothesis in contexts that address limitations with past tests of this hypothesis. Experiment 1 showed that decision-makers prefer unambiguous chance-based options to ambiguous ability-based options (reflecting ambiguity aversion). This was the case even in a context where decision makers perceived personal competence. Experiments 2 and 3 demonstrated that individuals display ambiguity aversion even when choosing among ability-based options, using two different operationalizations of ambiguity — neither of which was related to personal controllability or ability.

Seminal work by Heath and Tversky (1991) suggested that people exhibit ambiguity aversion in the classic Ellsberg (1961) paradigm because they do not perceive personal competence in the domain — an example of what Frisch and Baron (1988) regarded as missing information. Heath and Tversky showed that people are willing to choose ambiguous options in domains where they feel competent. The current set of experiments extends these findings both conceptually and methodologically. In previous studies, many factors were potentially confounded including controllability and the source of ambiguity. Thus, the preference for ambiguous options may have resulted from preferences for options that were more controllable and where ambiguity was due to lack of precision in stated probabilities rather than due to something unrelated to personal control. Most studies involved hypothetical choices between one chance-based option and one ability-based option, and the chance-based option always included a fixed (unambiguous) set of probabilities. Perceived competence was typically based on *a priori* beliefs rather than performance feedback, and participants usually made choices for themselves and not others — obfuscating to some degree the role of perceived competence.

The current studies addressed these methodological issues. The second two experiments involved consequential choices with real outcomes and offered more than one ability option so that we could examine the role of ambiguity independent of the role of controllability. These

two studies also featured an ambiguous chance-based option so that all options incorporated some degree of ambiguity — again in an attempt to separate out ambiguity from controllability. In all experiments, ambiguity always existed independent of the individual's ability or performance (in the first two experiments due to chance and in the final experiment due to a characteristic of the opponent). In addition, the latter two experiments relied less on *a priori* beliefs about personal competence and more on actual performance feedback.

Inclusion of these methodological features allowed several extensions of the competence hypothesis. People do prefer ambiguous ability-based options to unambiguous chance-based options, as shown in conventional tests of the competence hypothesis (Grieco & Hogarth, 2004; Heath & Tversky, 1991). However, this preference may have more to do with controllability than ambiguity per se, and may depend greatly on the source of ambiguity. We showed here that people are ambiguity-averse in a context where the ambiguity associated with the ability-based option is unrelated to personal control — suggesting the very important conclusion that ambiguity aversion may override effects of competence or controllability. Given that people generally prefer controllable and ability-based courses of action to those deriving from chance (Goodie, 2003; Howell, 1971; Klein & Kunda, 1994), it was important to show that preferences followed predicted patterns of ambiguity-aversion and ambiguity-seeking regardless of controllability. By creating chance outcomes that themselves varied in ambiguity and by linking winnings to the performance of other people who themselves varied in ambiguity, we were able to accomplish this goal.

As predicted by the competence hypothesis, beliefs about competence do play an important role in driving preferences. We reinforced this point by showing that people use performance information to make decisions — and do so both for themselves and others. That they do so for others helps to separate out the potentially independent effects of competence and personal controllability.

One question that our findings do not address is the extent to which manipulating ambiguity outside of the control of the decision-maker is tantamount to increasing the perception of missing information that Frisch and Baron (1988) regarded as a major component of ambiguity aversion. Although we assumed that participants perceived choices we designed to be high in ambiguity as possessing missing information, we did not ask participants to report their perceptions directly. It is also likely that some types of missing information are viewed as more important than others; for example, had participants been told that the experimenter rather than the computer was choosing the numbers that would determine whether they won in a particular game, participants may have grown suspi-

cious of the experimenter and shown even more ambiguity aversion.

In addition to building on tests of the competence hypothesis, our findings have practical implications. People tend to be ambiguity-averse in response to risk communications, and the findings here may explain why. In particular, Han and colleagues (e.g., Han, Moser, & Klein, 2006) find that people respond with ambiguity aversion when they believe that recommendations concerning risk reduction in a given domain are ambiguous. For example, when people endorse the statement that “there are so many different recommendations about preventing cancer, it's hard to know which ones to follow,” they feel that cancer is less preventable and feel more worried and more at risk. The more worried they are, the greater the relationship between perceptions of ambiguity in cancer messages and the sense that cancer is not preventable. Our findings suggest why this may be the case — because the source of ambiguity is viewed as unpredictable and uncontrollable. If risk communications afforded greater perceptions of control, people might exhibit less ambiguity aversion in response to them.

Ambiguity is omnipresent in the feedback we receive, the information we possess, and the decisions we make on an everyday basis, and it can come in many forms. Important work on the competence hypothesis has helped us move beyond traditional demonstrations of ambiguity aversion. To the degree that we understand how people respond to ambiguity, we will better understand processes that underlie decision-making and behavior in multiple contexts of interest.

References

- Becker, S. W., & Brownson, F. O. (1964). What price ambiguity? Or the role of ambiguity in decision-making. *Journal of Political Economy*, *72*, 62–73.
- Brun, W., & Teigen, K. (1990). Prediction and postdiction preferences in guessing. *Journal of Behavioral Decision Making*, *3*, 17–28.
- Camerer, C., & Weber, M. (1992). Recent developments in modeling preferences: Uncertainty and ambiguity. *Journal of Risk and Uncertainty*, *5*, 325–370.
- Chow, C. C., & Sarin, R. K. (2001). Comparative ignorance and the Ellsberg paradox. *Journal of Risk and Uncertainty*, *22*, 129–139.
- Curley, S. P., Yates, F. J., & Abrams, R. A. (1986). Psychological sources of ambiguity avoidance. *Organizational Behavior and Human Decision Processes*, *38*, 230–256.
- Dawes, R. M., & Mulford, M. (1996). The false consensus effect and overconfidence: Flaws in judgment or flaws in how we study judgment? *Organizational*

- Behavior and Human Decision Processes*, 65(3), 201–211.
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, 75, 643–669.
- Fox, C. R., & Tversky, A. (1995). Ambiguity aversion and comparative ignorance. *Quarterly Journal of Economics*, 110, 585–603.
- Fox, C. R., & Weber, M. (2002). Ambiguity aversion, comparative ignorance, and decision context. *Organizational Behavior and Human Decision Processes*, 88, 476–498.
- Frisch, D., & Baron, J. (1988). Ambiguity and rationality. *Journal of Behavioral Decision Making*, 1, 149–157.
- Goodie, A. S. (2003). The effects of control on betting: Paradoxical betting on items of high confidence with low value. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 598–610.
- Grieco, D., & Hogarth, R. M. (2004). *Excess entry, ambiguity seeking, and competence: An experimental investigation*. UPF Economics and Business Working Paper 778.
- Han, P. K. J., Moser, R., & Klein, W. M. P. (2006). Perceived ambiguity about cancer prevention recommendations: Relationship to perceptions of cancer preventability, risk, and worry. *Journal of Health Communication*, 11, 51–69.
- Harris, P. (1996). Sufficient grounds for optimism? The relationship between perceived controllability and optimistic bias. *Journal of Social and Clinical Psychology*, 15, 9–52.
- Heath, C., & Tversky, A. (1991). Preference and belief: Ambiguity and competence in choice under uncertainty. *Journal of Risk and Uncertainty*, 4, 5–28.
- Hogarth, R. M., & Kunreuther, H. (1989). Risk, ambiguity, and insurance. *Journal of Risk and Uncertainty*, 2, 5–35.
- Howell, W. C. (1971). Uncertainty from internal and external sources: A clear case of overconfidence. *Journal of Experimental Psychology*, 89, 240–243.
- Keren, G. B., & Gerritsen, L. E. M. (1999). On the robustness and possible accounts for ambiguity aversion. *Acta Psychologica*, 103, 149–172.
- Klein, W. M., & Kunda, Z. (1994). Exaggerated self-assessments and the preference for controllable risks. *Organizational Behavior and Human Decision Processes*, 59, 410–427.
- Krueger, J. I. (2003). Return of the ego--Self-referent information as a filter for social prediction: Comment on Karniol (2003). *Psychological Review*, 110(3), 585–590.
- Krueger, J., & Clement, R. W. (1996). Inferring category characteristics from sample characteristics: Inductive reasoning and social projection. *Journal of Experimental Psychology: General*, 125, 52–68.
- Ritov, I., & Baron, J. (1990). Reluctance to vaccinate: Omission bias and ambiguity. *Journal of Behavioral Decision Making*, 3, 263–277.
- Yates, J. F., & Zukowski, L. G. (1976). Characterization of ambiguity in decision-making. *Behavioral Science*, 21, 19–21.