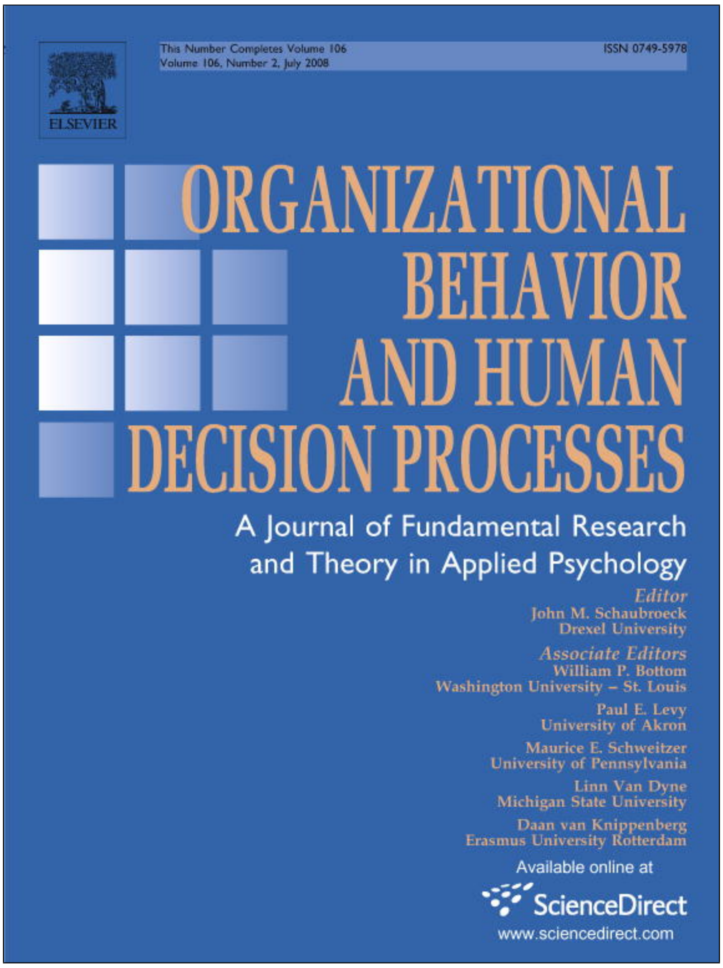
This Number Completes Volume 106
Volume 106, Number 2, July 2008
ISSN 0749-5978
ELSEVIER
ORGANIZATIONAL BEHAVIOR AND HUMAN DECISION PROCESSES
A Journal of Fundamental Research and Theory in Applied Psychology
Editor
John M. Schaubroeck
Drexel University
Associate Editors
William P. Bottom
Washington University – St. Louis
Paul E. Levy
University of Akron
Maurice E. Schweitzer
University of Pennsylvania
Linn Van Dyne
Michigan State University
Daan van Knippenberg
Erasmus University Rotterdam
Available online at
ScienceDirect
www.sciencedirect.com

Available online at www.sciencedirect.com
ScienceDirect
Organizational Behavior and Human Decision Processes 106 (2008) 125–142
ORGANIZATIONAL BEHAVIOR AND HUMAN DECISION PROCESSES
www.elsevier.com/locate/obhdp

The effect of safe experience on a warnings' impact: Sex, drugs, and rock-n-roll

Greg Barron [a,*], Stephen Leider [b], Jennifer Stack [c]

[a] *Harvard Business School, Baker Library, Boston, MA 02163, USA*
[b] *Department of Economics, Harvard Business School, Baker Library, Boston, MA 02163, USA*
[c] *Harvard Department of Economics, Boston, MA 02163, USA*

Received 20 April 2007
Available online 18 April 2008
Accepted by Robyn Dawes

Abstract

In many contexts we are warned against engaging in risky behavior only after having past safe experience. We examine the effect of safe experience on a warning's impact by comparing warnings received *after* having safe personal experience with those received *before* people start making choices. A series of five experiments studies this question with a paradigm that combines both descriptive information (i.e. the warning) and experiential information (safe outcomes). The results demonstrate two separate advantages to an early warning that go beyond the warning's mere informational content. When an early warning coincides with the beginning of a decision-making process, the warning is both weighted more heavily in future decisions (the Primacy Effect) and induces safer behavior that becomes the status quo for future choices (the Initial History Effect). While both effects operate indirectly through choice inertia, the primacy effect also operates directly on choices. This pattern of behavior is inconsistent with the "ideal" Bayesian for whom the order of information revelation does not influence subsequent behavior. The effect was robust across settings with and without forgone payoffs and when the consequences for risk taking are delayed until the end of the experiment. The results imply that, even after being adequately warned, some people may continue to take risks simply because they incurred good outcomes from the same choice in the past. Implications for policy and theory are discussed.

Keywords: Decision making; Experience; Description; Vioxx; Warnings; Cisapride; Sex education; RIAA

Introduction

Consider a recent example of a major safety warning: Vioxx, a non-steroidal anti-inflammatory drug developed by Merck & Co., was approved safe and effective by the Food and Drug Administration (FDA) on May 20, 1999. Subsequent research suggested that long-term daily use of Vioxx may increase the risk of heart attack[1] (Bombardier et al., 2000), and led the FDA to add black-box warnings, the strongest warning the FDA can impose on a pharmaceutical, to Vioxx labeling in April 2002. At the time the FDA warning was added, millions of people had already been taking Vioxx on a daily basis.[2] The FDA hoped that consumers would incorporate this new information about the risks of Vioxx into their decisions—which should lead many individuals to cease taking Vioxx. However, even after the subsequent withdrawal of Vioxx from the worldwide

* Corresponding author.
E-mail address: gbarron@hbs.edu (G. Barron).

[1] The increase in heart attacks was relative to that observed for Naproxen, another non-steroidal anti-inflammatory drug. While the VIGOR study did not include a control group, a later study with randomly controlled treatments confirmed that Vioxx increased the risk of heart attack. The drug was subsequently withdrawn worldwide (FDA, 2005).

[2] In 2002, at least 25.2 million people in the US received prescriptions for Vioxx and other cox-2 inhibitors (Dai, Stafford, & Alexander, 2005), the class of medications that shares Vioxx's risk profile.

0749-5978/$ - see front matter © 2008 Elsevier Inc. All rights reserved.
doi:10.1016/j.obhdp.2007.11.002

market, an estimated 2.25 million people continued taking the drug until their prescription ran out.[3]

In this paper we ask the question—Do warnings received *after* having safe personal experience have the same impact as those received *before* people start making choices? It is an important question as there is often a tradeoff between the costs of expediting a warning, by more extensive pre-testing for example, and those incurred when people partake in risky behavior before they are warned about the possible consequences. Policy makers implicitly assume that once a warning is provided, people can be expected to adjust their behavior regardless of when the warning is received relative to their own previous experience. In this paper we explicitly question this assumption. The question is particularly interesting since existing research leads to three competing hypotheses.

Warnings are an important information source designed to protect people from harm and have been shown to be effective in increasing safe behaviors (Cox, Wogalter, Stokes, & Murff, 1997). In particular, warnings are necessary when risky behavior may lead to extreme negative consequences but more frequently results in desirable outcomes. In these cases individual experience is likely a poor measure of the real consequences of risky behavior. Engaging in risky sexual behaviors, committing certain crimes, and taking medications with risky side effects are all examples where the typical personal experience belies the dangers addressed by the warning. For example, the salience of the risk imparted by the warning "unprotected sex carries the risk of HIV infection" may fade in light of the typically safe experience with unprotected sex.[4]

In these situations people have two types of information: *descriptive*, in the form of the warning, and *experiential*, from the outcomes of their past decisions. Existing decision-making research typically focuses on choices based on only one of these types of information (but not both). In description-based decisions, the information available to the decision maker consists of a symbolic representation of possible outcomes and probabilities (ex. choice between two gambles with known outcome distributions). Prospect Theory (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992) is one elegant summary of decisions of this type.

In experience-based decisions, on the other hand, there is typically no objective information about the alternatives and decisions are made based on the history of outcomes from making the same decision in the past. Various learning models have been proposed to describe how such decisions are made (Camerer & Ho, 1999; Denrell, 2005; Erev & Barron, 2005; Roth & Erev, 1995; Herriot, Levinthal, & March, 1985, to name but a few). In summary, the decision-making literature does not provide a clear sense of what to expect for decisions based on *both* descriptive (i.e. a warning) and experiential information.

The distinction between the two sources of information is important in the context of warnings about rare events as recent studies on description and experience-based decisions have shown distinctly different patterns of behavior when rare events are involved. Barron and Erev (2003) show that while people *overweight* the probability of rare events in description-based decisions, as predicted by Prospect Theory (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992), they appear to *underweight* the probability of rare events in experience-based decisions (see also Barron & Erev, 2003; Erev & Barron, 2005; Hertwig, Barron, Weber, & Erev, 2004).[5] Based on these findings, it is difficult to draw a prediction for situations where rare events are involved and the decision maker has both kinds of information, much less to hypothesize about how the timing or order of presentation will affect subsequent behavior.

Competing predictions

Three distinct hypotheses may be drawn regarding how people weight information obtained from different sequential presentations. Equal weighting, the first hypothesis, is implied by Bayes' theorem (holding fixed the set of observed outcomes[6]) and is a normative benchmark for how rational agents will integrate information revealed sequentially. Research has shown that experience plays a role in Bayesian reasoning by allowing the individual to directly experience base rates (Koehler, 1996; Hertwig & Ortmann, 2001). In applying Bayes' theorem to the current context of a warnings tim-

[3] A Harris Interactive (2004) poll of 2065 adults reported that 9% of people who taking Vioxx at the time of the withdrawal "Continued to take Vioxx until the prescription ran out". Independently, 10% reported as being "not at all concerned" that their health might be negatively effected.

[4] The odds of contracting HIV from a single act of heterosexual intercourse with a randomly selected person in the US are approximately 1:328,000 assuming a HIV/AIDS prevalence rate of 364 per 100,000 (CDC, 2004) and a mean transmission rate of 1 in 1000 (Royce, Seña, Cates, & Cohen, 1997).

[5] Fox and Hadar (2006) note that in general behavior reflecting the underweighting of rare events may be the result of biased samples (in small samples, one is more likely to encounter the rare event less frequently than expected than to encounter it more frequently than expected) or of errors in judgment. In our study no subject in the analysis sees the rare event but all subjects have a descriptive warning indicating its probability; moreover in experiments 2–4 all subjects have the same sample.

[6] A Bayesian agent could still have different posteriors if he samples differently from the two options (e.g. if the "hot stove" effect pushes the agent to cease sampling from the gamble that provided a large loss). We will rule this out in Experiments 2–4 by providing subjects with forgone payoffs. Thus subjects' choices will not affect their ability to update their posteriors for both gambles.

ing, it is simple to show that the theorem is insensitive to the order in which information is revealed. The theorem predicts that a warning followed by a series of safe outcomes will result in the same posterior risk estimate as a series of safe outcomes followed by a warning. Consequently, no matter where a warning is provided within a sequence of safe outcomes from experience, we expect subsequent risk taking (following the full sequence) to be the same.

A competing hypothesis can be drawn from the literature on experience-based decisions, which reports a tendency for decision makers to overweight *recent* information (De Bondt & Thaler, 1987; Kahneman & Tversky, 1982). Possibly caused by limitations on memory, this tendency is found both when the information is the monetary outcomes of repeated choices (Barron, Erev, & Yechiam, 2006) and when the information is a sample, repeatedly drawn from some distribution, with no immediate monetary consequences (Hertwig et al., 2004).

Survey-based research, on the role of experience in judging risks, is also consistent with the recency assumption. Halpern-Felsher et al. (2001) found that subjects reporting more experience with risky behaviors rated the chance of a negative outcome occurring (contingent on engaging in the risky behavior) as lower than did participants with no such experience. For example, people who had consumed alcohol in the past estimated the risk of getting into a car accident as lower than did people who reported never consuming alcohol. Presumably, those who had experienced the risk did so with a safe outcome in most of the time. In those cases, a safe outcome will often have been the most recent information available to those who subsequently estimated the risk of a bad outcome to be lower. Returning to our question of combined description and experience, it is possible that a description (i.e. a warning) provided within a sequence of safe experiences will also receive more weight when it is recent. Thus, a "Recency hypothesis" predicts that a warning presented after having safe experiences will have a larger impact on *subsequent* risk taking than a warning presented before the first decision, since it will have been experienced more recently.

A third hypothesis can be drawn from the classic psychological literature on memory. Studies of peoples' ability to recall sequentially presented lists of information demonstrated a serial position effect, with items presented at the *beginning* having a higher probability of recall, a phenomenon they called primacy (Madigan, 1971). Primacy is also thought to play a role in the "self fulfilling prophecy" of social interactions. One person, having formed initial expectations and beliefs about another person, acts in ways that cause the behavior of the target to appear to confirm these beliefs (Snyder & Stukas, 1999). Similarly, Denrell (2005) argues that when the probability of interacting with someone depends on current impressions, negative initial impressions will be more stable than positive impressions. New information that might disconfirm the initial negative impression is less likely to be collected, once the probability of future interactions has decreased. An additional dynamic, which suggests the early warning would have a larger impact by changing initial behavior, is seen in simple reinforcement-learning models that assume that choice probabilities depend only on the average reinforcement received from chosen alternatives (for example, Roth & Erev, 1995). If an alternative provides a very negative initial outcome, people will be less likely to choose it in the future and may potentially forgo desirable outcomes (a similar idea is explored in Denrell & March, 2001). In the current context of a warning's timing, a simple "Primacy hypothesis" predicts that the impact of a warning will be maximal when it is the first piece of information encountered, before acquiring information from personal experience. A more detailed prediction is that the initial warning will reduce long-term risk taking by reducing exposure to the desirable outcomes provided by the risky behavior.

Motivating examples: Sex, drugs, rock-n-roll

It is easy to find (and think of) additional real-world examples beyond Vioxx that are consistent with the Primacy hypothesis. For example, during a seven year period following the U.S. Surgeon General's 1986 warning about AIDS and HIV, condom use increased the most among women younger than 20. While there are undoubtedly many reasons for this, it is consistent with the idea that a warning has the most impact on those who encounter it before acquiring much personal experience. In a related vein, regular condom use was found to be highest when parent-adolescent sexual communication occurred at a younger age (Hutchinson, 2002). Again, it was those with less experience who were most affected by a warning against risky behavior that usually provides a desirable outcome.[7]

In 1993 Cisapride, a gastrointestinal promotility agent, was first marketed in the U.S. and, by 1995, had approximately 5 million users. In that year, the Food and Drug Administration (FDA) ordered a "black-box" warning regarding contraindications to using Cisapride. The warning, which must be added to the drug's label, was based on 61 reported incidents including 4 deaths. In a study that examined Cisapride usage before and after the black-box warning, the data show a minor *increase* in usage of 2% among repeat users but a decrease of 17% amongst first time users (Smalley et al., 2000). As with sexual risk communication, it is the inexperienced individual (in this case, the

[7] Unprotected intercourse is reported as being more pleasurable than using a condom (Thomsen, Stalker, & Toroitich-Ruto, 2004).

new user) who seems most affected by the warning. Similarly, a black-box warning describing the potential for an increased risk of suicidality with the use of antidepressants in a pediatric population had a higher impact on new prescriptions (Thomason, Riordan, Schaeffer, Cox, & Kratochvil, 2006).

Going beyond sex and drugs, in the second half of 2003 the Recording Industry Association of America (RIAA) sent out a clear warning to the estimated 35 million individuals a month who were downloading music through peer-to-peer networks in violation of copyright law. They identified and sued 261 individuals for copyright infringement and the settlements were typically for $3000 or more. By January 2004 the RIAA's legal campaign seemed to be working with downloading down 14% (Rainie & Madden, 2004). Since the RIAA was explicitly targeting "heavy" file sharers, one might expect the decrease in downloading to be largest among this group of more experienced users. However, the average number of music files acquired per user during the same period actually increased from 59 to 63 files, which suggests that the RIAA's legal tactics had more of an effect on the lighter and less experienced downloaders (NPD MusicWatch Digital, 2003).

Taken together, these examples support the Primacy account: warnings have a greater impact the earlier they are encountered, relative to personal experience, in cases where risky behavior provides good outcomes most of the time. In all three examples, a warning about a rare event had less of an effect on those who had already experienced positive outcomes from risky behavior—those who already had unprotected sex with negative consequences, those who were already taking Cisapride without an adverse reaction, and those were downloading large quantities of copyrighted music without being caught.

However, none of the examples above are controlled studies, nor are the different groups exogenously determined, and thus remain more suggestive then conclusive. People who consume alcohol may have different beliefs regarding the risk of car accidents, different risk preferences, or both. Heavy downloaders of copyrighted music may well be less risk averse by nature or may value music more highly.[8] Our paper's primary contribution is to clarify the effect of a warning's timing (before vs. after safe experience) through a series of controlled laboratory studies. A second contribution is identifying the underlying mechanism by testing the three hypotheses (equal weighting, recency and primacy). Identifying the mechanism allows us to derive non-trivial implications of our findings. More broadly, our paper serves as an initial step in understanding decisions based on both descriptive and experiential information.

[8] A similar bias may exist in the observation that swimmers experienced with diving into the shallow end of a pool were more likely to repeat the risky behavior despite a newly posted warning sign (Goldhaber & deTruck, 1988).

Study 1: The effect of a warning's timing on risk taking

Method

Design

Each participant performed a binary choice task under uncertainty one hundred times with immediate feedback. Participants repeatedly chose between two unmarked buttons presented on the screen (see Web Appendix A[9]). Each button was associated with one of two static distributions referred to here as S (for safe) and R (for risky). The S distribution provided a certain gain of $0.10, to be paid at the end of the experiment, while the R distribution provided a gain of $0.13 with probability .999 and a loss of 15 dollars with probability .001. Thus, the R distribution provides better outcomes than the S distribution 99.9% of the time but carries a significant risk. R was given a higher expected value (of $0.115 as compared to S's $0.10) so as to not trivialize the task. The position of the S and R buttons (right vs. left) was randomized for each participant.

Participants were randomly allocated to the two experimental conditions, "Early" and "Late". In the Early condition, subjects received a warning about the risky button before they began the binary choice task. The warning stated "Each time you hit the (Left/Right) button there is a 1 in 1000 chance (.001 probability) that you will lose $15. This is the only way you can loose money in the game". The second sentence was added to reduce potential heterogeneity stemming from concerns about hidden losses in the S distribution. Participants in the Late condition first began the binary choice task and then received the warning after completing half the task (50 trials). In addition to the warning above, these participants were told that "This has been true since the beginning of the game". In both conditions participants were told that outcomes were i.i.d (see instructions in Appendix A).

Participants

Sixty-two volunteers served as paid participants in the study. Participants in this and the other studies described in this paper were students (graduate or undergraduate) from several local universities. In addition to the performance contingent payoff, described above, participants received $10 for showing up. The final payoff was approximately $21.

[9] http://www.people.hbs.edu/gbarron/WebAppendix/WebAppendix1.pdf

Apparatus and procedure

Participants were aware of the expected length of the study (10–30 min), so they knew that it included many rounds. To avoid an "end of task" effect (e.g. a change in risk attitude), they were not informed that the study included exactly 100 trials.[10] Payoffs were contingent upon the button chosen. Two types of feedback immediately followed each choice: (1) the payoff for the choice, that appeared on the selected button for the duration of 1 second and (2) an update of an accumulating payoff counter, which was constantly displayed.

Results

Four participants incurred the loss of $15 and were removed from the remainder of the analysis. While their particular behavior is potentially interesting,[11] there were too few in order to perform any meaningful analysis. Following our motivating examples in the introduction, our main interest is in the vast majority of people who do not experience the rare event described in the warning.

Fig. 1 shows the mean proportion of R choices in the first and second halves of the experiment.

The difference between the two conditions in the proportion of R choices in the first 50 trials is significant ($t(56) = 5.45$, $p < .001$) but not surprising as only participants in the Early condition had received the warning about the risky option. More important is the comparison between choice proportions in the second half of the experiment, after participants in *both* conditions had both received the warning and had experienced feedback from 50 choices. In these 50 trials, participants in the Late condition were choosing the risky option 19% more often than participants in the Early condition ($t(56) = 1.67$, $p < .10$). This result is consistent with the Primacy hypothesis, that a warning presented before having good outcomes from risky behavior has a larger impact than a warning presented later on.

It is interesting to note the proportion of R choice in the first 50 trials in the Late condition (0.84). For these participants the task seems a simple one: the S button always provides $0.10 while the R button provides

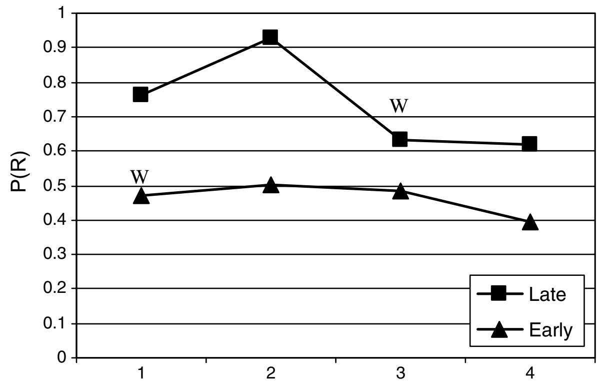

Fig. 1. Proportion of R choices in four blocks of 25 trials for Early and Late conditions (Experiment 1). "W" denotes the first block after the warning is received by subjects.

$0.13. The finding that people are slow to learn this is consistent with both previous research and descriptive models. In a very similar paradigm, where participants were repeatedly choosing between options providing 10 points with certainty or 11 points with certainty, Haruvy and Erev (2001) observed the modal subject choosing the EV maximizing option only 81% of the time in the first 100 trials. This figure increased to 93% with the addition of forgone payoff information. Descriptive reinforcement-learning models typically capture this slow adaptive process by assuming equally weighted initial propensities (for examples see Camerer & Ho, 1999; Erev & Barron, 2005; Roth & Erev, 1995) and a probabilistic choice rule (Luce, 1959). The result is a slow adaptive learning process even in an environment that provides certain outcomes. Indeed, it is precisely the "certainty" of the outcome that is being learnt.

Alternative explanations

Since the participants in the Late condition were choosing R more often, they were winning more money and their subsequent choices could be susceptible to "house money" effects (Thaler & Johnson, 1990). However, the repeated trial nature of the task should have diluted such effects quickly, with participants feeling increasingly endowed with their cumulative earnings and reluctant to risk them. Thus, the experimental task should effectively model repeated decision making after a few rounds.

While all the participants in the first half of the experiment both received the warning and feedback from 50 choices, the two conditions differed in one important aspect. Each participant was observing a unique series of feedback in the binary choice task. In particular, those in the Early condition were observing significantly more $0.10 outcomes as they were choosing the risky option less often (having received the warning from the very beginning). It is possible that these differences are responsible for the difference in risk taking in the

[10] Not knowing the length of the study also prevents participants from using probability-based reasoning (the focus on the likelihood of achieving a particular aspiration level) (Lopes, 1996). This type of reasoning bases choice on the probability of coming out ahead, which is a function of the number of choices to be made. A second reason for not telling participants the game's length is that this better approximates the real-world decisions that interest us. In such situations, the number of future choices to be made is often unknown and it is difficult to prescribe optimal behavior.

[11] Experiment 3 poses an alternative to the potentially objectionable elimination of participants. The interested reader may refer to Weinstein (1989) for a comparison between victims and non-victims of disasters. In a nutshell, the effect on victims (for example, seatbelt use after a traffic accident) is surprisingly short lived.

second half of the experiment. According this account, warnings may be weighted equally regardless of when they are received, contrary to the Primacy hypothesis. However, an early warning may bias future choices by reducing the likelihood that one observes good outcomes from risky choices. A similar argument was proposed by Denrell (2005) regarding impression formation with negative first impressions reducing the likelihood that disconfirming evidence be collected, known also as the "hot stove" effect (see also Denrell & March, 2001).

To evaluate this possibility, we replicated Experiment 1 with the addition of forgone payoff information. Thus, in the first half of the experiment every participant observed the exact same series of feedback in the binary choice task (i.e. 50 presentations of \$0.10 and \$0.13 from S and R, respectively). The only remaining informational difference between the two conditions is the timing of the warning. This will allow a stronger test of both the Bayesian hypothesis that predicts no difference in subsequent risk taking, and the Primacy hypothesis that predicts a larger impact for the Early warning even in the presence of forgone payoffs in the binary choice task.

Study 2: Replication of Study 1 with forgone payoff information

Method

Design

As in Experiment 1, each participant performed a binary choice task one hundred times with immediate feedback. The S distribution again provided a certain gain of \$0.10, to be paid at the end of the experiment, while the R distribution provided a gain of \$0.13 with probability .999 and a loss of 15 dollars with probability .001. After each choice, both the outcome of the chosen gamble, as well as the outcome of the other gamble, were displayed to the subject (in contrast with Experiment 1, where the subject only observed the outcome of his chosen gamble). Participants were randomly allocated to Early and Late conditions, where the subjects in the Early condition were informed about the loss outcome of the Risky gamble before the first choice while the subjects in the Late condition received the warning after period 50.

Participants

Sixty-seven volunteers served as paid participants in the study. The final payoff was approximately \$21.

Apparatus and procedure

Participants were aware of the expected length of the study (10–30 min), although they were not informed that the study included exactly 100 trials. Payoffs were contingent upon the button chosen. Two types of feedback immediately followed each choice: (1) the payoff for both gambles, which appeared on the respective buttons for the duration of 1 s and (2) an update of an accumulating payoff counter, which was constantly displayed.

Results

Fig. 2 shows the mean proportion of R choices aggregated across subjects in each condition and across the 50 trials in each half of the experiment. As in Experiment 1, participants who incurred the rare outcome of −\$15, or who observed it as a forgone payoff, were removed from the analysis.

Consistent with the visual impression from the figure, and with the results of Experiment 1, participants who had received the Early warning were taking significantly less risk in trials 51–100 (53% of choices as compared with 80% in the Late condition, $t(58) = 2.91$, $p < .01$). The finding of a larger impact for the early warning suggests that the effect is not driven by participants' specific series of observations that bias their future choices since as all participants saw the exact same series (50 outcomes of \$0.10 and \$0.13). This suggests that the observed effect on risk taking is not driven by a reduction in the likelihood that one observes good outcomes from risky choices (i.e. the "hot stove effect").

The addition of forgone payoff information also allows us to more cleanly refute the basic Bayesian hypothesis which predicted insensitivity to the timing of the warning. While the "ideal" Bayesian, with priors for all possible events, is indifferent to the timing, a more boundedly rational Bayesian in the Late condition may be surprised by the warning. This surprise, in turn, might cause our boundedly rational Bayesian to update

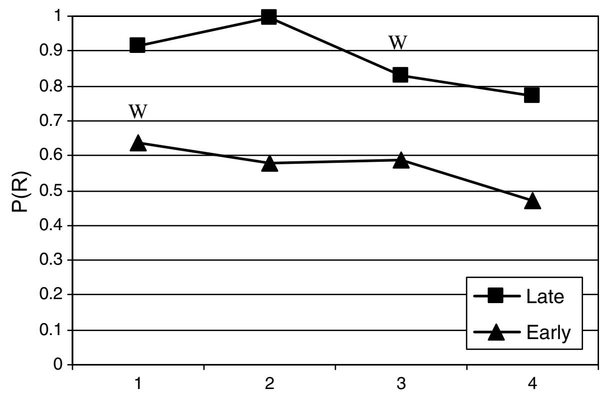

Fig. 2. Proportion of R choices in four blocks of 25 trials for Early and Late conditions (Experiment 2). "W" denotes the first block after the warning is received by subjects.

the probability of just such a surprise from choosing the risky button. For this reason, the warning included the information that only by pressing R could one could loose money in the experiment. Both Experiments' 1 and 2's results are consistent with the Primacy hypothesis.

We remained somewhat concerned that the results could be biased by the elimination of the few participants who observed the rare outcome. Participants were more likely to see the rare event the more often they chose the R option. The elimination of these, more risk taking, participants might differentially lower the mean observed proportion of R choices in the Late conditions (where R was being chosen more often). To control for this, we replicated Experiment 2 with the exception that the loss incurred by the rare event, if it occurred at all, was delayed until the end of the experiment. As a result, none of the subjects in either condition actually observe the rare event *during* the experiment when we are collecting data. Another advantage to this design is that it examines the robustness of the previous results to rare but delayed outcomes (a feature common to many health-related risks).

Study 3: Replication of Study 2 with delayed rare event

Method

Design

As in Experiments 1 and 2, each participant performed a binary choice task one hundred times with immediate feedback. The S distribution again provided a certain gain of $0.10, to be paid at the end of the experiment, while the R distribution provided a gain of $0.13 with probability .999 and a loss of 15 dollars with probability .001. Participants were randomly allocated to Early and Late conditions, where the subjects in the Early condition were informed about the loss outcome of the Risky gamble before the first choice while the subjects in the Late condition received the warning after period 50. However, in this experiment the warning was "Each time you hit the (Left/Right) button there is a 1 in 1000 chance (.001 probability) that you will lose $15 at the end of the experiment. This is the only way you may lose money in the experiment". As in Experiments 1 and 2, participants in the Late condition were told that this was true from the beginning of the experiment. After each choice outcomes were displayed for both gambles. However, since the loss would only be incurred at the end of the experiment, the Risky option always produced $0.13 as an immediate payoff.

Participants

Forty-six volunteers served as paid participants in the study. The final payoff was approximately $21.

Apparatus and procedure

Participants were aware of the expected length of the study (10–30 min), although they were not informed that the study included exactly 100 trials. Payoffs were contingent upon the button chosen. Two types of feedback immediately followed each choice: (1) the immediate payoff for both gambles, which appeared on the respective buttons for the duration of 1 s and (2) an update of an accumulating payoff counter, which was constantly displayed. At the end of the experiment subjects were informed of any losses they had incurred during the task.

Results

Fig. 3 shows the mean proportion of R choices aggregated across subjects in each condition and across the 50 trials in each half of the experiment. All 46 participants were used in the analysis.

The results replicate the findings of the previous two experiments with participants who had received the Early warning taking significantly less risk (46% of choices as compared with 71% in the Late condition, $t(44) = 2.27$, $p < .05$) in trials 51–100.

There are clear differences between Study 3's paradigm and that used in the previous two studies. Delaying the loss until the end of the experiment may introduce additional cognitive processes that were not present in the other studies. Most notably, one might expect discounting to affect the evaluation of the risk of a future loss (Fredrick, Loewenstein, & O'Donoghue, 2002; Liabson, 1997; Samuelson, 1937). While the standard interpretation of discounting would not suggest a difference between the treatments, since for a given trial the time span between the choice and the loss is the same in both conditions, one could imagine a discounting process where the time elapsed between an individual becoming aware of the loss and the actual occurrence

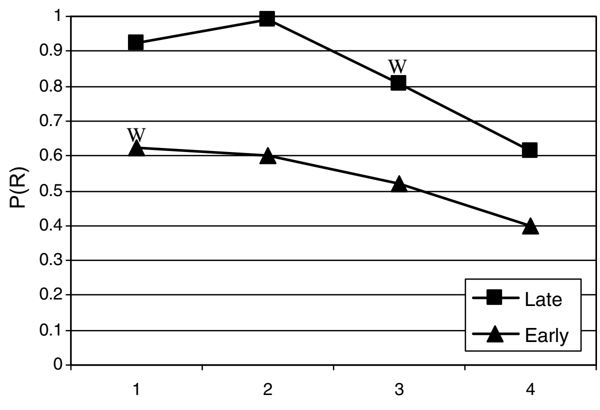

Fig. 3. Proportion of R choices in four blocks of 25 trials for Early and Late conditions (Experiment 3). "W" denotes the first block after the warning is received by subjects.

of the loss could play a role. However, to the extent that discounting occurred differently between the two conditions, the most natural interpretation would suggest that the Early condition would have a longer discount horizon and hence should have engaged in greater risk taking. In fact, since subjects in the Early condition engaged in less risk taking (as in Studies 1 and 2), either this effect of discounting is not present or it is overwhelmed by the effects of information primacy.

Choice inertia in Experiments 1–3

While the above results clearly indicate that delaying the introduction of the warning has a substantial impact on choices in the second half of the experiment, we must be careful in inferring the mechanism for this effect. There are in fact two differences between the subjects receiving the Early warning and the Late warning. First, the order of information (both descriptive and experiential) about the gambles is different. Second, however, the subjects in the Late warning treatment had a markedly different pattern of choices in the first half of the experiment (because they did not know about the negative outcome of the Risky lottery). If an individual's choice history affects his future choices, then this difference in choice histories could play a role in the differing choice behavior in the second half. To assess whether choice inertia (i.e. the propensity to choose more often the same gamble as the individual has chosen in the past) was playing a role in subjects individual choices throughout the experiment we display the proportion of risky choice as a function of the length of the previous "run" (i.e. a series of 1–5 consecutive risky or safe choices). If there is no inertia and current choices are independent of past choices this graph should be flat. Fig. 4 presents proportions of risky choices as a function of previous run length for Early and Late conditions, aggregated across experiments and subjects. The flat line in the figure is the expected proportion of risky choices assuming independence. While an econometric analysis follows, the visual impression clearly suggests significant inertia with subjects being more likely to repeat a choice the longer the run of identical past choices.[12]

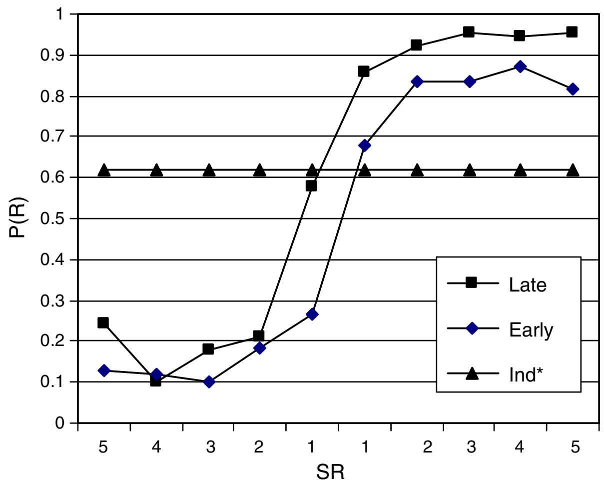

Fig. 4. Proportion of risky choices according to length of previous run. *Proportion of risky choices assuming independence.

Decomposing the treatment effect

Given the observed inertia in choices, there are two major differences induced by the treatment that could cause the difference in behavior we see in periods 51–100 (see Figs. 1–3). First, the timing of the warning may have changed the general preference of each group for choosing risky or safe, what we will call the Main Effect. In addition, in period 51 the two groups have different distributions of choice histories. Given that subjects exhibit decision inertia, their different histories mean that they will continue to make different choices for the rest of the experiment, what we will call the Initial History Effect. The following regression analysis decomposes these two effects to assess their relative magnitude. We regress subjects' risk decision on a set of individual fixed effects, as well as indicator variables for the subjects' recent decision history. We consider two specifications to capture inertia: the number of consecutive safe or risky decisions, and the average number of recent risky decisions (irrespective of order). Both yield similar results, and have similar explanatory power. While we focus on a horizon of five decisions, results are similar with a shorter horizon. We also pool the data from Treatments 1–3, results are similar for each Treatment analyzed separately.

Comparing the distribution of fixed effects between the two treatment groups provides a measure of the Main Effect of the treatment manipulation. The distributions are significantly different for both specifications (1: $\Delta = 0.073$ $t(164) = 3.62$ $p < .01$, 2: $\Delta = 0.075$ $t(164) = 3.39$, $p < .01$) Hence, even controlling for differences due to inertia there is a direct change in propensity to chose the risky option. The Late treatment group is significantly more likely to choose risky than the Early group, independent of the choice history.

The Main Effect will also have a secondary indirect effect on choices due to inertia. Even if the two treatments began the second half of the experiment with the same set of histories, as the Main Effect pushes choices towards or away from Risky their histories will become increasing different. The inertia of choices would then amplify the Main Effect. We can calculate

[12] The same pattern of inertia emerges when the analysis is constrained to trials 51–100 and to those participants whose overall proportion of risky choices was between 0.25 and 0.75. This suggests that the pattern is not driven only by those participants who chose the same option throughout the experiment.

this using our estimated magnitudes for inertia. For both the Early and Late treatment groups we take the total distribution of histories in period 51 and calculate forward what the expected distribution of histories should be for period 52, and again for period 53, and so forth. However, for each group we apply the average fixed effect for that group (the difference that generates the Main Effect). In each period we can then calculate the differing probability to choose Risky due to the increasingly different choice histories. This amplification of the Main Effect is evident in rows 1B and 2B of Table 2.

The magnitude of the Initial History Effect through the inertial channel can similarly be calculated using the estimated coefficients and the distribution of histories for each treatment group in period 51. For this calculation the two treatment groups begin with different histories, but have the same choice probabilities given a history (i.e. we exclude the Main Effect). The net inertial bias towards risky choices is displayed in rows 1C and 2C of Table 2.

For either specification of inertia, the inertial impact of the History Effect diminishes quite rapidly (due to the converging histories). We can compare, in Table 2 below, the relative magnitude of the Main and Initial History Effects in different periods. While the Initial History Effect's impact (via inertia) is much stronger immediately after the late warning, it diminishes rapidly. Hence the total treatment effect diminishes over time, but not to zero, since the Main Effect still operates. Our estimates indicate that the Main Effect accounts for between 70% and 90% of the overall treatment effect, while the indirect effect is between 10% and 30%.[13]

Experience as a shifting reference point

We may also consider whether our results can be explained with a narrower adaptation of the existing literature of description-based choice under uncertainty, i.e. Prospect Theory. According to this explanation, subjects in the Late condition have "gotten used to" receiving 13 cents every period and framed the prospect of receiving 10 cents for a safe choice as a loss. Typically the reference point around which gains and losses are determined is assumed to be the status quo level of wealth. However, when the individual is receiving a sequence of outcomes from the same choice problem the reference point may instead reasonably be determined by the (recent) sequence of outcomes received assuming they influence ones expectations (Kahneman & Tversky, 1979). If a subject has received 13 cents in every period, receiving only 10 cent may be viewed as a loss, even though the subject's wealth has increased. This would cause subjects who had chosen the Risky button frequently in the past (and thus have a high reference point) to view the choice as between a small certain loss (the Safe button) vs. a small probability of a large loss (the Risky button). This would increase their preference for the Risky option, relative to a subject with a lower reference point. Hence in our experiment, subjects in the Late treatment would have a higher reference point at the start of the second half (due to choosing Risky frequently before the warning) than subjects in the Early treatment, and hence would chose risky more often in the second half. Note that, since choices and outcomes were perfectly correlated in our experiments (on the subsample of subjects we consider), this model could be seen as the underlying mechanism for inertia (and thus the Initial History Effect), but not for the Main Effect identified previously.

We examine several standard formulations of Prospect Theory in the literature to assess how well this explanation fits out data (details are available in a Web Appendix B[14]). While these models can generate a difference in propensity to choose risky between subjects who have frequently chosen safe and those who have frequently chosen risky, for standard parameter values the largest difference was approximately 10%. Fig. 4 and Table 1 suggest that to explain the difference in choices in our data, the effect would have to be between 6 and 9 times larger. Hence we conclude that a moving reference point is not an effective modeling approach to explaining our data.

Brief summary and the role of decision-making processes

The econometric analysis above supports independent primacy and initial history effects. An early warning is associated with less subsequent risk taking because it is weighted more heavily than a late warning (i.e. primacy). Additionally, different behavior engaged in by the subjects in the first half of the experiment due to the timing of the warning is persistent in the second half (i.e. initial history). The role of a deliberate decision-making process is clearly important for the Initial History Effect to occur, since it operates through the inertia of decisions. An alternative inertial mechanism based on monetary outcomes, namely loss aversion, is shown to play either a lesser or non-existent role. However, the role of a deliberative decision-making process for the Main Effect (i.e. primacy) is less clear. Research on the importance of deliberation and reasoning suggests that they can have significant effect on attitudes,

[13] Calculating the indirect inertial component of the Main Effect by taking the difference between an estimate with different histories plus the average fixed effect difference and an estimate with different histories but without the fixed effect difference yields similar results.

[14] http://www.people.hbs.edu/gbarron/WebAppendix/WebAppendix1.pdf

Table 1
Individual choices controlling for history

Prob (choose risky)			
Indicator variable	(1)	Indicator variable	(2)
1 consecutive R	0.260*** (0.019)	1 R of last 5 choices	0.127*** (0.014)
2 consecutive R's	0.515*** (0.023)	2 R's of last 5 choices	0.271*** (0.017)
3 consecutive R's	0.489*** (0.024)	3 R's of last 5 choices	0.248*** (0.017)
4 consecutive R's	0.493*** (0.026)	4 R's of last 5 choices	0.438*** (0.016)
5 consecutive R's	0.506*** (0.016)	5 R's of last 5 choices	0.645*** (0.012)
2 consecutive S's	−0.147*** (0.021)		
3 consecutive S's	−0.140*** (0.022)		
4 consecutive S's	−0.180*** (0.023)		
5 consecutive S's	−0.160*** (0.015)		
Constant	0.359*** (0.014)	Constant	0.213*** (0.0082)
Observations	8300	Observations	8300
Number of subjects	166	Number of subjects	166
R-squared (within)	0.372	*R*-squared (within)	0.251
R-squared (between)	0.999	*R*-squared (between)	0.995
R-squared (overall)	0.755	*R*-squared (overall)	0.712

Note: The dependent variable is whether the subject chose the R button. Standard errors are reported in parenthesis. ***$p < .01$, **$p < .05$, *$p < .1$. Fixed-effects regression for treatments 1–3, period 51–100. Similar results obtain for estimating each treatment separately. Omitted variables: (1) 1 consecutive S, (2) 0 R's of last 5 choices.

judgments and choices (Shafir, Simonson, & Tversky, 1993; Wilson & Schooler, 1991; Wilson, Dunn, Kraft, & Lisle, 1989). Will an early warning reduce subsequent risk taking even in the absence of the deliberative process that accompanies choices made in the first half of the experiment?

The answer to the question has important practical implications - whether the correct interpretation of an "early" warning is "as soon as possible", or "early" relative to an individuals' decision opportunity (we leave the theoretical interpretations to the discussion section below). Consider two elaborations of the original primacy hypothesis. First, merely being exposed to an early warning may be sufficient to generate long lasting differences in the propensity for risk taking, even if the relevant decisions are made only after accruing experience. Alternately, individuals' *deliberation* over the two options while integrating both descriptive and experiential information may mediate the impact of an early warning on risk taking. Thus, under this account of the primacy hypothesis, an early warning is predicted to loose its impact when experiential information is accrued separate from, and prior to, decision making. Possibly, when both descriptive and experiential information are obtained *before* making a decision, individuals tend to rely more heavily on the experiential information regardless of the order.

To test these two versions of the primacy hypothesis, we ran a replication of Experiment 2, but removed the process of decision making from the first half of the experiment. Instead, participants merely observed outcomes sampled from *both* S and R options in trials 1–50. If the decision-making process is not a necessary mediator of the main results in Experiments 1–3 then the results of the new experiment should be similar since the information structure remains the same (subjects observe 50 pairs of draws from the two distributions). However, if deliberation is necessary for the primacy effect to occur, then we expect the Early and Late conditions to exhibit similar risk taking. Moreover, if experiential information is weighed more heavily in the absence of the primacy effect we should see greater total risk taking.

Study 4: Replication of Study 2 with sampled outcomes for both S and R

Method

Design

As in Experiments 1–3, each participant performed a binary choice task one hundred times with immediate feedback. However, for the first 50 trials, the S and R buttons were inactive and a third button was added to the middle of the screen. When pressed, this button produced outcomes for both the S and R buttons; these outcomes had no monetary consequences for the subjects. Participants were told that this constituted a sampling phase of the experiment and that no money would be won or lost. On trial 51 (and after the warning in the Late condition), the third button disappeared and the S and R buttons were activated. At that time, participants saw a message saying that the sampling phase had ended and that all choices were now for actual earn-

Table 2
Estimated effect magnitudes

Treatment effect	$t = 51$ Effect	$t = 51$ Total effect	$t = 51$ %	$t = 75$ Effect	$t = 75$ Total effect	$t = 75$ %	$t = 100$ Effect	$t = 100$ Total effect	$t = 100$ %	Avg. Effect	Avg. Total effect	Avg. %
(1) **Main Effect**		0.073	20.73		0.221	99.88		0.221	100.00		0.207	90.39
A Baseline propensity difference	0.073			0.073			0.073			0.073		
B Inertial effect of baseline difference	0.000			0.148			0.148			0.134		
Initial History Effect		0.279	79.27		2.6E−04	0.12		1.8E−07	0.00		0.022	9.61
C Inertial effect of Pd. 1–50 history difference	0.279			2.6E−04			1.8E−07			0.022		
Total		0.352			0.221			0.221			0.229	
(2) **Main Effect**		0.075	20.37		0.182	82.71		0.187	97.55		0.168	70.75
A Baseline propensity difference	0.075			0.075			0.075			0.075		
B Inertial effect of baseline difference	0.000			0.107			0.112			0.093		
Initial History Effect		0.293	79.63		0.038	17.29		0.005	2.45		0.070	29.25
C Inertial effect of Pd. 1–50 history difference	0.293			0.038			0.005			0.070		
Total		0.368			0.220			0.191			0.238	

ings. The S distribution again provided a certain gain of \$0.10, to be paid at the end of the experiment, while the R distribution provided a gain of \$0.13 with probability .999 and a loss of 15 dollars with probability .001. After each choice (after period 50), both the outcome of the chosen gamble, as well as the outcome of the other gamble, were displayed to the subject. Participants were randomly allocated to Early and Late conditions, where the subjects in the Early condition were informed about the loss outcome of the Risky gamble before the first choice while the subjects in the Late condition received the warning after period 50.

Participants

Sixty volunteers served as paid participants in the study. The final payoff was approximately \$16.

Apparatus and procedure

Participants were aware of the expected length of the study (10–30 min), although they were not informed that the study included exactly 100 trials. Payoffs were contingent upon the button chosen. Two types of feedback immediately followed each choice: (1) the payoff for both gambles, which appeared on the respective buttons for the duration of 1 s and (2) an update of an accumulating payoff counter, which was constantly displayed.

Results

Fig. 5 shows the mean proportion of R choices aggregated across subjects in each condition and across the 50 trials in the second half of the experiment. A single subject observed the rare event and was not included in the analysis.

There was no significant difference in the proportion of participants choosing the risky option ($t(58) = 1.00$, ns). This result supports the hypothesis

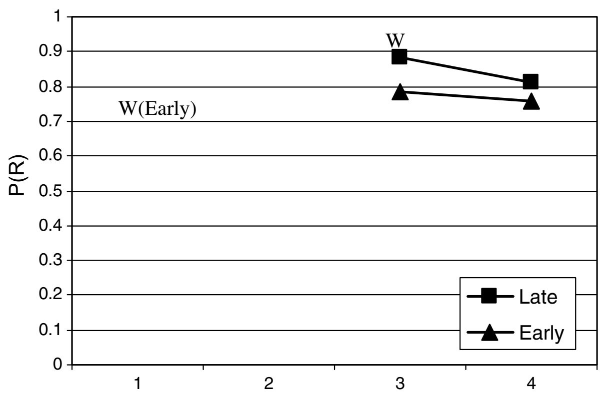

Fig. 5. Proportion of R choices in two blocks of 25 trials for Early and Late conditions (Experiment 4). "W" denotes the first block after the warning is received by subjects.

that the direct effect of the warning's timing is mediated by a decision-making process. In the current context, an early warning is associated with less risk taking only when the warning coincides with actual decision making.

The mean proportion of R choices in the Early condition (0.77) appears to be higher than that in Studies 1–3 (0.44, 0.53, 0.46, respectively). This supports the interpretation that absent the primacy effect, experiential information is relied on more heavily. Note that this increase in risk taking is distinct from the Initial History Effect (induced by inertia) since subjects made no decisions in the first 50 trials.

Study 5: Preferred behavior given full information

The large impact of an early warning in Studies 1–4 is particularly notable, given that the Risky lottery has a larger expected value than the Safe lottery. To substantially reduce the choice of the otherwise appealing Risky lottery to the extent observed, the effect of the warning's timing must have greatly increased the impact of a rare, large loss on the formation of the subjects' risk attitudes. However, to consider how this phenomenon should impact the development of public policy, one can take two perspectives. By providing an early warning, we are either helping subjects engage in the safe behavior they truly wish to pursue, or we are biasing their choices away from a risky but otherwise sufficiently rewarding activity. While our motivating examples suggest that society's interests are typically best served by inducing safe behavior, under the latter interpretation an early warning that reduces risk taking may be undesirable from the individual's point of view. To better understand whether a proactive policy of early intervention to prevent risky behavior could serve the underlying interests of the individual, we consider a paradigm in which individuals can make choices between the two options given full descriptive information about the distribution of outcomes for both lotteries.[15] Even under full descriptive information the theoretical and empirical literature suggests that it may be difficult to identify an individual's "true" preferences. However, we can find the limits of the level of risk the individual is willing to take on given the potential rewards.

We consider a paradigm of multiple prospective choices, both for comparability to our earlier studies, and to give scope for individuals to pursue expected value maximization. While Prospect Theory predicts that, given a full description of the safe and risky distributions, people will prefer the safe distribution in a one-shot choice since it avoids the (overweighted) small probability of a significant loss (that looms large, relative to gains), previous research has shown that when *multiple* choices are made under full descriptive information, behavior is more consistent with expected value maximization, the risky distribution in the current context (see Keren & Wagenaar, 1987; Keren, 1991; Lopes, 1981; Wedell & Bockenholt, 1990). Therefore, the level of risk taking we observe is an upper-bound on the amount of risk the individual is willing to take on. Hence if subjects prefer the safe behavior in a setting where we might expect to observe expected value maximization, we can be more confident that the choice behavior in the Early condition of our previous studies is more in line with the individual's underlying risk preferences than the Late condition. If individuals exhibit a strong preference for the Safe lottery absent potentially biasing manipulation of the presentation of information (relative to experience) then public policy makers can be more confident that early intervention is good for both society and the individual.

Method

Design

Subjects were presented with descriptions of the same two distributions used in the previous experiments ([\$0.10, 1] and [\$0.13,.999; −\$15]). Subjects were randomly allocated between two conditions: "Repeated" and "Allocation". In the Repeated condition, subjects were asked to make a single choice between either the safe or risky gamble (that were not labeled) and were told that their earnings, or losses, would be the sum of 100 independent plays of their chosen gamble (see Appendix A). The Allocation condition was identical except that subjects were asked to allocate 100 choices between the two gambles, i.e. could select a combination of both Risky and Safe choices. Subjects had to enter a separate Allocation (0–100) for each gamble and the sum of the two allocations had to equal 100.

Participants

Fifty volunteers served as paid participants in the study. The final payoff was approximately \$20.

Results

A clear preference for the safe gamble, with the lower expected value, was observed in both the Repeated ($P(\mathrm{R}) = 0.28$) and Allocation ($P(\mathrm{R}) = 0.24$) conditions ($z = 2.2$, $p < .05$ and $t(24) = 2.97$, $p < .01$, respectively). Interestingly, only one subject in the Allocation condition chose a mixed allocation. While the results have interesting links to similar studies on multiple prospec-

[15] This is how the vast majority of the literature thinks about eliciting risk preferences (see Hertwig et al. (2004) for a brief review).

tive choice,[16] for our purposes it is sufficient to note that early warnings appear to move behavior towards subjects' preferred level of risk taking under full information.

Discussion

In many real-world contexts, the *timing* of warning against risky behavior seems to affect its impact. Specifically, those who receive warnings *before* having (usually profitable or enjoyable) experience with risk taking exhibit less risk taking than those warned only *after* having such experience. Our paper demonstrates that this pattern is robust, even in controlled settings, which eliminate selection according to risk preferences as a factor. Subjects who received an early warning took less risk than those who received a late warning even though all subjects had the same warning and feedback, albeit in different order.

The second contribution is to identify the underlying mechanisms: primacy (mediated by decision making) and initial behavior (working through inertia[17]). While the present application is novel, these mechanisms have been well documented (e.g. Madigan, 1971 on primacy, and Samuelson & Zeckhauser, 1988 on status quo bias). One difference between the classic primacy studies (where primacy occurred with visually presented lists) and Experiment 4 (where we found no primacy from mere observation of outcomes) is that we used two different types of information, descriptive and experiential. Future research should examine the hypothesis that when the types of information are sufficiently different, an integrative process, such as that required when making choices, is necessary for primacy to manifest.

The current findings also extend the research on rare events in decision making. While description-based decisions have been shown to exhibit overweighting of small probabilities and experience-based decisions to exhibit underweighting of small probabilities, no previous research that we are aware of has examined behavior in the presence of both descriptive information and experience. The experiments presented here examine the extreme case where the rare event is only described but never experienced. By comparing risk taking in Experiments 1–4 to that observed in Experiment 5 (with a full description) we can draw the preliminary conclusion that experienced gains are more salient than symbolically described gains and can increase risk taking when they are associated with risky behaviors. Future research should examine the extent that experience may mitigate the effects of a described rare event.

Limitations and future research

While in all five experiments the risky option had a higher expected value than the safe option, in applied contexts we do not know to what extent risky behavior maximizes the individual's expected utility. For example, Viscusi (1992) argues that smokers' decision to continue smoking can be seen as a rational choice in the EUT framework. However, Study 5 demonstrates that even in the current context where the Risky option has a 15% higher expected value, individuals overwhelmingly prefer the safe choice (under their "natural" full information, repeated-choice risk attitudes). Moreover, as a society it is clear that in many situations, we would prefer people to choose more safe options even when risky behavior may be preferred by some individuals. For specific applications, however, policy makers may want to assess the attractiveness and risk level of the risky behavior, and weigh the merits of an early intervention for both society and the individual.

While the differences between the risky and safe behavior are quite stark in Study 5, the mix of descriptive and experiential information presented in Studies 1–4 likely reduced the discriminability of the two options (i.e. choosing safe vs. risky is a "tough call"). Previous work on the status quo bias suggests that the bias is stronger when people are less able to clearly discriminate between their options according to their true preferences. To the extent that our observed choice iner-

[16] This result is not inconsistent with past studies that observed more expected value maximization in a repeated task like the one used here. While the expected value maximization rates were low in both conditions it is reasonable to expect even less expected value maximization in a one-shot choice between these two distributions, as is predicted by Prospect Theory. Note that the one-shot decision is the choice between earning 10 cents with certainty and earning 13 cents while risking a loss of 15 dollars with probability .001. More interesting however is the comparison between expected value maximization rates in the current experiment and those observed in Experiments 1–4. The difference between the experiments is that the information about the positive outcomes, 10 cents for Safe and 13 cents for Risky, was descriptive in Experiment 5 and was experiential in Experiments 1–4. This change appears to make the risky gamble more attractive (the mean proportion of risky choices in the Early condition in Experiments 1–4 ranged from 0.46 to 0.77). This suggests that experiencing relatively good outcomes from risk taking can increase risk taking substantially as compared to merely knowing about these outcomes by way of a description. Future research should examine more closely the relative decisions weights given to described and experienced information in different contexts. It is interesting to note that the almost total absence of mixed allocations in the Allocation condition does not support Chen and Corter (2006) who find that people often prefer a mix of sure and risky options over pure bundles of either. One potential explanation for the contradicting findings lies in the specific gambles used in the studies. That paper used gambles that people were relatively indifferent between when they were forced to choose between a purely risky bundle or a purely safe bundle (i.e. approximately half the participants chose one while the other half chose the other). In our study however, in the repeated condition, people had a strong preference for the safe option to begin with. Thus, the ability to mix options in the Allocation condition provided little added benefit. This suggests that Chen and Corter's results may be most important when there is relative indifference between the options.

[17] i.e. the tendency to repeat past decisions (March, 1994).

tia is related, our paradigm is thus likely to be quite sensitive to the effects of information timing. However, most of the relevant applications and particularly our motivating examples, have additional rich contexts beyond our stark setting. For settings where the additional detail increases the disciminability of the possible behaviors, we may expect the timing of a warning to be less germane. In contrast, when the richness of particular application further clouds the issue of whether the safe or risky option is more attractive, the role of an early warning could be just as strong. Future research should examine the robustness of the effect to changes in the ability to discriminate between risky and safe options.

Of course, the role of a richer context can have other implications for the generalizability of our findings. Risky or safe behavior may be associated with different norms at different times or from different perspectives (e.g. sexual behavior and adolescence vs. adulthood, the criminality of "copying" vs. "stealing" music, etc.). The issue of the optimal timing for sex education to impact sexual behavior has long been studied. This question is complicated since adolescence, i.e. the age at which many individuals change from not being sexually active to being sexually active (Brady & Halpern-Felsher, 2007 is also the transition from childhood to adulthood. At this time the changing (and relaxing) norms about sexual behavior may reduce the perceived costs and increase the perceived benefits of engaging in risky behavior (Goldberg, Halpren-Feisher, & Millstein, 2002). Our study suggests that one factor that might play a role in this complicated issue is the relative timing of sex education with the onset of sexual decision making. In particular, as Study 4 emphasizes, it is important that the warning be contemporaneous with decision making. While this may make the optimal timing a difficult problem, in the following section we discuss the potential for role-playing and simulated decisions to enhance the efficacy of an early warning by increasing the alignment between the warning and decision making. Alternately, initiating sexual education at a young age, to be continued throughout adolescence, (see Kirby et al., 1994 for a review and Frost & Forrest, 1995 for evidence of the efficacy of targeting younger adolescents) may be sufficient, given possible heterogeneity in the onset of decision making. Early education will provide a contemporaneous warning to early decision makers, while continuing education still provides a contemporaneous warning to average and late decision makers. This suggestion is consistent with the observation that successful sexual education programs have been found to be longer in duration (Kim, Stanton, Li, Dickersin, & Galbraith, 1997). Admittedly, our study does not consider multiple or ongoing descriptive warnings, and hence this possibility should be studied explicitly in future research. Lastly, while much discussion has developed over the content of warnings about risky sexual behavior: e.g. whether sex education should be "abstinence only" (i.e. telling adolescents to postpone having sex but not suggesting condoms as a means of reducing the risks of sex) or "comprehensive" (providing adolescents with both messages), our findings have little to say about which kind of sex education is likely to be more effective (instead, see Bruine de Bruin, Downs, & Fischhoff, 2007; Kirby, 2002; McKay, 2000). However, our findings *do* suggest that for any form of sex education content, beginning the program at or shortly before the onset of decision making is likely to be important for its efficacy.

The broader literature on risk communication is concerned with many types of warnings and risks beyond those experienced by participants in our studies. We used a descriptive warning with numeric information (i.e. probabilities and outcomes). While this type of warning may be thought of as a quantitative summary of expert knowledge (Morgan, Fischhoff, Bostrom, & Atman, 2002, p. 6), other types of warnings are presented as simple statements of advice. For example, "Don't do it, it's dangerous!" The risk faced by participants in our studies (i.e. loosing $15) was individual, controllable, and was neither catastrophic nor dreadful. As such, it falls low on factor 1 of Slovic, Lichtenstein, & Fischoff (1980) three factor hazard space. The second factor in this space relates to characteristics such as observability and immediacy. The studies presented here cover more then one point on this factor. While the risk of loss was immediate in Studies 1 and 2 it was delayed and unobservable in Study 3. Additionally, in our studies the warning about the risk was always described as a probability of a loss given one trials worth of risky behavior, the same time span over which people were making decisions. Shaklee & Fischhoff (1990) show that when the time horizon used in risk communication is different then that used by people in their risk assessments, people tend to misestimate the risk implied by the original communication. Yet another limitation of the current context was that while our warnings indicated a risk, their absence could not be interpreted as implying safety. Meyer (2004) explores the difference between warnings systems that instruct us to perform an action when the warning is on (compliance), and those where we are to refrain from performing an action when the warning is off (reliance). More research is needed to validate the current findings on a broader set of warnings and risk types.

Lastly, in our setting all subjects faced the same distribution of outcomes. However, in many settings (particularly in the domain of health) individuals are heterogeneous in their probability of incurring the rare outcome. Hence, positive or negative outcomes are informative both about the general consequences of a given behavior and about an individual's "riskiness type". Several studies have suggested that individuals may over-attribute the non-occurrence of a rare event, e.g. pregnancy, as indicative of being a low probability type, i.e. infertile (see Downs, Bruine de Bruin, Murray, & Fischhoff, 2004; Zikmund-Fisher, 2004). These contexts may be particularly interesting avenues in which to extend our research, since an effective descriptive statement may be particularly useful in a setting where individuals appear to have great difficulty in interpreting experiential information.

Practical implications

In themselves, the five studies reported here can not address the larger question of the magnitude of the mere timing effect in warnings against HIV, adverse drug reactions, or being caught downloading music illegally. Still, the demonstration of any potential effect can have large implications in situations where we previously thought there to be no effect whatsoever. Consider the original example regarding the use of NSAID's such as Vioxx. The current results suggest that the addition of an FDA warning only after Vioxx was on the market for almost 3 years lead to two groups exhibiting excess risk taking: those who took (or prescribed) Vioxx before the warning due to incomplete information, and those who continued to use Vioxx after the warning due to their earlier positive experience. The latter group develops risky habits and incurs the costs, for themselves and for society, associated with those habits. Such costs can be avoided by an early warning, which is weighed more heavily if it coincides with the beginning of decision making, and influences initial choices that, in turn, tend to persist in the future. However, in some contexts it may be infeasible (or too costly) to acquire sufficient information to provide an accurate "early" warning for many individuals. For drugs like Vioxx, the usage of the medication by early adopters may have been necessary to bring to light the risks of rare negative side effects (absent vastly larger pre-approval studies). Hence, when only a late warning is possible (e.g. after a drug is widely available in the market), it is important to know how to make the late warning as effective as possible (e.g. it may need to be presented more thoroughly, saliently and/or extensively, and with greater cost, to overcome behavioral inertia and have its intended impact). Alternately, there may be some amount of propensity to take risks that cannot be overcome by a late warning. While such conclusions cannot be drawn directly from our findings, they do suggest that it may be important for policy makers that future research examine these possibilities explicitly.

On the other hand, early warnings are often costly themselves if large amounts of data are needed to identify a risk and if data collection comes at a cost. In the case of phase 3 drug trials designed to evaluate the drug with large populations of affected patients, larger trials that could identify rare adverse effects are more expensive. These costs must be balanced against the risk of a previously undetected adverse drug reaction manifesting itself only after a drug goes to market with possible economic repercussions for the drug company.

Our results are also potentially related to the known difficulty in recalling defective products by the Consumer Product Safety Commission. Past research shows that response rates are often as low as 5–10% (see for example Tobin, 1982). While difficulty in locating product owners and product longevity are often blamed for the low success rates, our results suggest that even when information about a recall does reach consumers who posses the product, safe past experience might lower its efficacy. This is consistent with studies showing that response rates decline the longer the consumer owns the product controlling for both the ability to notify consumers and the product's longevity (Murphy & Rubin, 1988). Recalls of defective automobiles by the National Highway and Traffic Safety Administration were also found to have lower completion rates for cars that were owned for more than two years (Hoffer, Pruitt, & Reilly, 1994). Future research should estimate the magnitude of the effect of safe experience in these applied contexts and examine possible solutions.

The current findings raise questions about an organization's responsibility for peoples post-warning behavior. If an organization is held responsible for the outcomes of risky behavior *before* a warning is issued, as is the case when companies put a product on the market that leads to injuries or death, should they be held (partially) responsible for outcomes *after* the warning is issued? For example, in liability cases regarding a drug's adverse side effects pharmaceutical companies frequently argue "Don't blame us; we got the labeling right" (Avorn, 2004). While a legal treatment of the question is beyond the scope of this paper, our findings suggest that warning *content* is only part of the story and that some amount of risky behavior following a warning may be attributable to the *timing* of the warning. In other words, some people who take a risk by using a product after a warning has been issued may not have taken the same risk had the warning been issued earlier, before the accrual of safe experience. This raises non-trivial questions about an organizations post-warning liability.

Returning to the example of RIAA's legal campaign against illegal file sharing, it seems possible that the campaign is having less of an effect on the target population (i.e. heavy users) since this group has had the most "good" file-sharing experience before the campaign was initiated. If providing a warning is costly, and it certainly is when litigation is involved, then understanding the warnings differential effect on populations with different levels of experience is vital for an accurate cost-benefit analysis.

Lastly, the current results underscore the importance of ''education'' in the sense of early risk communication. The policy implications of the current results are easiest to apply when large populations begin making choices at approximately the same time, and thus a proactive warning program can specifically target those individuals who are about to begin risky or safe behavior (otherwise, timing the warning to coincide with the start of decision making is less straight forward). Teenagers entering adolescence, and thus exposed to opportunities for risky behaviors (sexual and others) with frequently reinforcing outcomes, are perhaps the best example where the timing of the warning relative to the onset of decision making for the individual is feasible. The message of the current research is very clear: while it may be true that ''timing is everything'', it would be an oversimplification to prescribe ''the earlier the better''. As noted above, regarding sex education, preventive education that warns against the negative consequences of risky behavior should be timed to coincide with people's initial decisions. In Experiment 4, where warnings and outcomes were merely observed, without any decision-making process, there was no advantage to an early warning. However, in all three experiments where an early warning was coupled with a subsequent decision-making process, the early warning had a larger impact on reducing subsequent risk taking behavior. In other contexts, such as the availability of a new drug, where individuals are more heterogeneous in when the relevant decision making will occur (i.e. developing a medical issue), then it may be sufficient to focus on making a warning as early as possible.

Some existing research suggests that it may be sufficient to use simulated decisions, coupled with warnings, to achieve a reduction in risky behavior. Downs et al. (2004) find that an interactive DVD that educates and warns against the risk of sexually transmitted diseases (STD's) was more effective at reducing risky behavior and contraction of STD's in adolescent females than was a book containing the same information. An important feature of the DVD was the presentation of video taped scenarios where participants were required to answer the question ''What would you do?'' and choose between a number of options. These results, together with the current findings, emphasize the potential importance of incorporating some form of explicit decision making in programs designed prevent risky behavior.

Acknowledgments

We are grateful to the National Science Foundation for financial support (to S. Leider). We have benefited from discussions with and comments from Al Roth, Max Bazerman, George Baker, Ben Shenoy, Orit Avni-Barron and seminar participants at Harvard Business School, Carnegie Mellon University, Subjective Probability Utility and Decision Making, Society for Judgment and Decision Making, and several anonymous referees. All remaining errors are our own.

Appendix A. Instructions for Experiment 1

In this experiment you will be operating a ''money machine'' with two buttons. You can click on either button with the mouse and change as often as you wish. With each button press you will win or lose an amount of money. The amount of money you receive each time is randomly chosen from a predefined list of amounts and each button has its own list. Each amount on a list has the same chance of being chosen each time. Your goal is to finish the experiment with as much money as possible. Your total earnings will be constantly displayed below the money machine.

Your final earnings depend on the amount of money you have at the end of the experiment. You are also guaranteed an additional $10 show up fee.

Each subject's experiment is slightly different. Some of you may see a pop up screen with more information about the two buttons. If you do, read this information carefully.

Appendix B. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.obhdp.2007.11.002.

References

Avorn, J. (2004). *Powerful medicines: The benefits, risks, and costs of prescription drugs*. New York: Knopf.

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*(16), 215–233.

Barron, G., Erev, I., Yechiam, E. (2006). The judgment-decision paradox in experience-based decisions and the contingent recency effect. *Harvard Business School Working Paper, No. 07-003*.

Bombardier, C. et al. (2000). Comparison of upper gastrointestinal toxicity of rofecoxib and naproxen in patients with rheumatoid arthritis. *New England Journal of Medicine, 343*, 1520–1528.

Brady, S., & Halpern-Felsher, B. (2007). Adolescents' reported consequences of having oral versus vaginal sex. *Pediatrics, 119*, 229–236.

Bruine de Bruin, W., Downs, J. S., & Fischhoff, B. (2007). Adolescents' thinking about the risks and benefits of sexual behavior. In M. Lovett & P. Shah (Eds.), *Thinking with data*. Mahwah, NJ: Erlbaum.

Camerer, C., & Ho, T. H. (1999). Experience-weighted attraction learning in normal form games. *Econometrica, 67*(4), 827–874.

Center for Disease Control. (2004). *HIV/AIDS surveillance report: Cases of HIV infection and AIDS in the United States* (Vol. 16).

Chen, Y., & Corter, J. (2006). When mixed options are preferred in multiple-trial decisions. *Journal of Behavioral Decision Making, 19*, 17–42.

Cox, E., Wogalter, M. S., Stokes, L., & Murff, E. (1997). Do product warnings increase safe behavior?: A meta analysis. *Journal of Public Policy and Marketing, 25*, 195–204.

Dai, C., Stafford, R. S., & Alexander, G. C. (2005). National trends in cyclooxygenase-2 inhibitor use since market release: Nonselective diffusion of a selectively cost-effective innovation. *Archives of Internal Medicine, 165*(2), 171–177.

De Bondt, W. F. M., & Thaler, R. H. (1987). Further evidence on investor over-reaction and stock market seasonality. *The Journal of Finance, 42*(3), 557–581.

Denrell, J. (2005). Why most people disapprove of me: Experience sampling in impression formation. *Psychological Review, 112*(4), 951–978.

Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science, 12*(5), 523–538.

Downs, J. S., Bruine de Bruin, W., Murray, P. J., & Fischhoff, B. (2004). When "it only takes once" fails: Perceived infertility predicts condom use and STI acquisition. *Journal of Pediatric and Adolescent Gynecology, 17*(3), 224.

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review, 112*, 912–931.

FDA. (2005). Sequence of Events with VIOXX, since opening of IND. Retrieved April 6, 2007. Avaliable from FDA Web site: http://www.fda.gov/OHRMS/DOCKETS/AC/05/briefing/2005-4090B1_04_E-FDA-TAB-C.htm.

Fox, C. R., & Hadar, L. (2006). "Decisions from experience" = sampling error + prospecttheory: Reconsidering Hertwig, Barron, Weber & Erev (2004). *Judgment and Decision Making, 1*(2), 159–161.

Fredrick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature, 40*(2), 351–401.

Frost, J., & Forrest, J. (1995). Understanding the impact of effective teenage pregnancy prevention. *Family Planning Perspectives, 27*(5), 188–195.

Goldberg, J., Halpren-Feisher, B., & Millstein, S. (2002). Beyond invulnerability: The importance of benefits in adolescents' decision to drink alcohol. *Health Psychology, 21*(5), 477–484.

Goldhaber, G., & deTruck, M. (1988). Effectiveness of warning signs: Gender and familiarity effects. *Journal of Products Liability, 11*(3), 271–284.

Halpern-Felsher, B., Millstein, S., Ellen, J., Adler, N., Tschann, J., & Biehl, M. (2001). The role of behavioral experience in judging risks. *Health Psychology, 20*(2), 120–126.

Harris Interactive Inc. (2004, December 10). Nearly Half of U.S. Adults Closely Following News Coverage of Vioxx Withdrawal. Retrieved April 6, 2007. Avaliable from Harris Interactive Web site http://www.harrisinteractive.com/news/allnewsbydate.asp?NewsID=872.

Haruvy, E., & Erev, I. (2001). On the application and interpretation of learning models. In R. Zwick & A. Rapoport (Eds.), *Advances in experimental business research*. Boston: Kluwer Academic Publishers.

Herriot, S. R., Levinthal, D., & March, J. (1985). Learning from experience in organizations. *American Economic Review, 75*, 298–302.

Hertwig, R., & Ortmann, A. (2001). Experimental practices in economics: A methodological challenge for psychologists? *Behavioral and Brain Sciences, 24*, 383–451.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Research article decisions from experience and the effect of rare events in risky choice. *Psychological Science, 15*(8), 534.

Hoffer, G., Pruitt, S., & Reilly, R. (1994). When recalls matter: Factors affecting owner responses to automotive recalls. *The Journal of Consumer Affairs, 28*(1), 96–106.

Hutchinson, M. K. (2002). The influence of sexual risk communication between parents and daughters on sexual risk behaviors. *Family Relations, 51*(3), 238–247.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*, 263–291.

Kahneman, D., & Tversky, A. (1982). Intuitive prediction: Biases and corrective procedures. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgement under uncertainty: Heuristics and biases*. London: Cambridge University Press.

Keren, G. (1991). Calibration and probability judgments: Conceptual and methodological issues. *Acta Psychologica, 77*, 217–273.

Keren, G., & Wagenaar, W. A. (1987). Violation of utility theory in unique and repeated gambles. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*(3), 387–391.

Kim, N., Stanton, B., Li, X., Dickersin, K., & Galbraith, J. (1997). Effectiveness of the 40 adolescent AIDS-risk reduction interventions: A quantitative review. *Journal of Adolescent Health, 20*(3), 204–215.

Kirby, D. (2002). *Abstinence-only programs delay the initiation of sex among young people and reduce teen pregnancy?* Washington, DC: National Campaign to Prevent Teen Pregnancy.

Kirby, D., Short, L., Collins, J., Rugg, D., Kolbe, L., Howard, M., et al. (1994). School-based programs to reduce sexual risk behaviors: A review of effectiveness. *Public Health Reports, 109*(3), 339–360.

Koehler, J. (1996). The base rate fallacy reconsidered: Descriptive, normative, and methodological challenges. *Behavioral and Brain Sciences, 19*(1), 1–53.

Liabson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics, 112*(2), 443–477.

Lopes, L. L. (1981). Decision making in the short run. *Journal of Experimental Psychology: Human Learning and Memory, 7*(5), 377–385.

Lopes, L. L. (1996). When time is of the essence: Averaging, aspiration, and the short run. *Organizational Behavior and Human Decision Processes, 65*(3), 179–189.

Luce, R. D. (1959). *Individual choice behavior. A theoretical analysis*. New York: John Wiley & Sons, Inc..

Madigan, S. A. (1971). Modality and recall order interactions in short-term memory for serial order. *Journal of Experimental Psychology, 87*, 294–296.

March, J. (1994). *A primer on decision making: How decisions happen*. New York: The Free Press.

Meyer, J. (2004). Conceptual issues in the study of dynamic hazard warnings. *Human Factors, 46*(2), 196–204.

McKay, A. (2000). Prevention of sexually transmitted infections in different populations: A review of behaviourally effective and cost-effective interventions. *Canadian Journal of Human Sexuality, 9*(2), 95–120.

Morgan, M. G., Fischhoff, B., Bostrom, A., & Atman, C. J. (2002). *Risk communication: A mental models approach*. New York: Cambridge University Press.

Murphy, R., & Rubin, P. (1988). Determinants of recall success rates. *Journal of Products Liability, 11*, 17–28.

NPD Group press release. (2003). RIAA lawsuits appear to reduce music file sharing. Retrieved April 6, 2007. Avaliable from NPD Group Web site: http://www.npd.com/press/releases/press_030825.htm.

Rainie, L., & Madden, M. (2004). Pew internet project and comscore media metrix data memo. Retrieved April 6, 2007. Avaliable from the Pew Internet and American Life Project Web site: http://www.pewinternet.org/pdfs/PIP_File_Swapping_Memo_0104.pdf.

Roth, A. E., & Erev, I. (1995). Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior, 8*, 164–212.

Royce, R. A., Seña, A., Cates, W., Jr., & Cohen, M. S. (1997). Sexual transmission of HIV. *The New England Journal of Medicine, 336*(15), 1072–1079.

Samuelson, P. (1937). A note on measuring utility. *Review of Economic Studies, 4*, 155–161.

Samuelson, W., & Zeckhauser, R. (1988). Status quo bias in decision making. *Journal of Risk and Uncertainty, 1*(1), 7–59.

Shafir, E., Simonson, I., & Tversky, A. (1993). Reason-based choice. *Cognition, 49*(1–2), 11–36.

Shaklee, H., & Fischhoff, B. (1990). The psychology of contraceptive surprises: Cumulative risk and contraceptive effectiveness. *Journal of Applied Psychology, 20*(5), 385–403.

Slovic, P., Lichtenstein, S., & Fischoff, B. (1980). Facts and fears: Understanding perceived risk. In R. C. Schwing & W. A. Albers (Eds.), *Societal risk assessment: How Safe is safe enough?* New York: Plenum.

Smalley, W., Shatin, D., Wysowski, D. K., Gurwitz, J., Andrade, S. E., Goodman, M., et al. (2000). Contraindicated use of cisapride. *JAMA, 284*(23), 3036–3039.

Snyder, M., & Stukas, A. A. Jr., (1999). Interpersonal processes: The interplay of cognitive, motivational, and behavioral activities in social interaction. *Annual Review of Psychology, 50*, 273–303.

Thaler, R. H., & Johnson, E. J. (1990). Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. *Management Science, 36*(6), 643–660.

Thomason, C., Riordan, H., Schaeffer, J., Cox, K., & Kratochvil, C. (2006). Antidepressant prescribing practices prior to and after black-box warnings. *AACAP, 33*, B6.

Thomsen, S., Stalker, M., & Toroitich-Ruto, C. (2004). Fifty ways to leave your rubber: How men in Mombasa rationalise unsafe sex. *Men's Sexual Health Special Edition, 80*(6), 430–434.

Tobin, R. J. (1982). Recalls and the remediation of hazardous or defective consumer products: The experiences of the consumer product safety commission and the national highway traffic safety administration. *Journal of Consumer Affairs, 16*(2), 278–306.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty, 9*, 195–230.

Viscusi, W. K. (1992). *Smoking: Making the risky decision*. New York: Oxford University Press.

Wedell, D. H., & Bockenholt, U. (1990). Contemplating single versus multiple encounters of a risky prospect. *The American Journal of Psychology, 107*(4), 499–518.

Weinstein, N. (1989). Effects of personal experience on self-protective behavior. *Psychological Bulletin, 105*(1), 31–50.

Wilson, T. D., Dunn, D. S., Kraft, D., & Lisle, D. J. (1989). Introspection, attitude change, and attitude–behavior consistency: The disruptive effects of explaining why we feel the way we do. In L. Berkowitz (Ed.). *Advances in experimental social psychology* (Vol. 19, pp. 123–205). San Diego, CA: Academic Press.

Wilson, T. D., & Schooler, J. W. (1991). Thinking too much: Introspection can reduce the quality of preferences and decisions [attitudes and social cognition]. *Journal of Personality and Social Psychology, 60*(2), 181–192.

Zikmund-Fisher, B. (2004). De-escalation after repeated negative feedback: Emergent expectations of failure. *Journal of Behavioral Decision Making, 17*(5), 365–379.