Failing to Learn from Experience about Catastrophes: The Case of Hurricane Preparedness

Robert J. Meyer
The Wharton School
University of Pennsylvania

March 2012
Working Paper # 2012-05

Forthcoming in *Journal of Risk and Uncertainty*

Risk Management and Decision Processes Center
The Wharton School, University of Pennsylvania
3730 Walnut Street, Jon Huntsman Hall, Suite 500
Philadelphia, PA, 19104
USA
Phone: 215-898-4589
Fax: 215-573-2130
http://opim.wharton.upenn.edu/risk/

THE WHARTON RISK MANAGEMENT AND DECISION PROCESSES CENTER

Established in 1984, the Wharton Risk Management and Decision Processes Center develops and promotes effective corporate and public policies for low-probability events with potentially catastrophic consequences through the integration of risk assessment, and risk perception with risk management strategies. Natural disasters, technological hazards, and national and international security issues (e.g., terrorism risk insurance markets, protection of critical infrastructure, global security) are among the extreme events that are the focus of the Center's research.

The Risk Center's neutrality allows it to undertake large-scale projects in conjunction with other researchers and organizations in the public and private sectors. Building on the disciplines of economics, decision sciences, finance, insurance, marketing and psychology, the Center supports and undertakes field and experimental studies of risk and uncertainty to better understand how individuals and organizations make choices under conditions of risk and uncertainty. Risk Center research also investigates the effectiveness of strategies such as risk communication, information sharing, incentive systems, insurance, regulation and public-private collaborations at a national and international scale. From these findings, the Wharton Risk Center's research team – over 50 faculty, fellows and doctoral students – is able to design new approaches to enable individuals and organizations to make better decisions regarding risk under various regulatory and market conditions.

The Center is also concerned with training leading decision makers. It actively engages multiple viewpoints, including top-level representatives from industry, government, international organizations, interest groups and academics through its research and policy publications, and through sponsored seminars, roundtables and forums.

More information is available at http://opim.wharton.upenn.edu/risk.

Failing to Learn from Experience about Catastrophes: The Case of Hurricane Preparedness

Robert J. Meyer*

*Robert Meyer is the Gayfryd Steinberg Professor of Marketing and co-director of the Center for Risk Management and Decision Processes at the Wharton School, University of Pennsylvania, 3730 Walnut Street, Philadelphia, PA 19104. Email: meyerer@wharton.upenn.edu. The author thanks Howard Kunreuther for comments on an earlier draft of this manuscript, and to Carol Heller for Editorial Assistance.

Failing to Learn from Experience about Catastrophes: The Case of Hurricane Preparedness

Abstract

The ability of individuals to learn from experience to invest in protection against catastrophic events is explored. Data are reported from two controlled experiments in which participants have a monetary incentive to learn from experience to make investments to protect against hurricane risks. A central finding is that investments display a short-term forgetting effect consistent with the use of reinforcement learning rules, where a significant driver of investments in a given period is whether storm losses were incurred in the precious period. Given the relative rarity of such losses, this reinforcement process produces a mean investment level that is below that which would be optimal for most storm threats. Investments are also found to be insensitive to the censoring effect of protection itself, implying that it is the size of experienced losses—rather than losses that are avoided—that is the primary driver of investment decisions.

"High dwellings are the peace and harmony of our descendants. Remember the calamity of the great tsunamis. Do not build any homes below this point."

--Inscription on a 600-year old tablet in the hills above Miyako, Japan, a city that was largely destroyed by the tsunami following the 2011 Great East Japan earthquake

Introduction

It often seems that individuals and communities have short memories for catastrophic losses. After Hurricane Katrina induced massive flooding in New Orleans and the Gulf coast in 2005, for example, the National Flood Insurance Program witnessed a 53% nationwide increase in new flood policies issued the following year—only to watch new purchases drop back to pre-2005 levels--with a 33% cancelation rate of existing policies--by 2008 (Michel-Kerjan, Lemoyne de Forges, and Kunreuther, 2012). Likewise, in the years since the Los Lomas earthquake of 1994, the state of California has watched the percentage of residents carrying earthquake insurance drop from 31% immediately after the quake to approximately 12% now, even in the face of increasing odds of a major event (Wilkinson 2008). Such examples of apparent forgetting are, of course, by no means limited to natural hazards; investors also seem to learn little from repeated experiences with stock market bubbles, as illustrated by the 2008 real-estate and equities crash (e.g., Hussam, Porter, and Smith 2008).

Yet, whether such anecdotes provide real evidence of learning failures is far from clear. Declines in the renewal of flood and earthquake insurance in quiet years after hazards, for example, could just easily be cast as an example of recovery from errors of short-term *over*-buying induced by exaggerated beliefs about risk that arise in the immediate aftermath of a disaster (e.g., Lerner, Gonzales, Small, and Fischhoff, 2003). Likewise, the opposite cases where disasters appear to induce *no* short-term adaptive response also need not be viewed as a failure to learn: a rational decision maker who uses Bayes' rule to update beliefs about the future likelihood of a future disaster will not necessarily increase protective investments after observing a single loss outcome, no matter how severe (e.g., Shafran 2011; Viscusi 1979). Because one can only gauge the optimality of a risky decision in terms of the information that was available to the decision maker at the time the decision was made, the question of whether failures to invest in protection provide evidence of rational learning versus irrational forgetting is difficult to infer from case studies alone.

The purpose of this paper is to attempt to resolve the question of whether there are, in fact, inherent limits to our ability to learn from experience about the value of protection against low-probability, high-consequence, events. We report the findings of a program of theoretical and experimental research in which participants are given a monetary incentive to learn from experience about how best to invest in both short and long-term protection against losses from an illustrative class of natural hazards, tropical cyclones. We build on prior work on reinforcement learning (e.g., Erev and Baron 2005; March 1996; Weber, Shafir, and Blais 2004; Shafran 2011) by first showing that oscillating cycles of increasing-then-decreasing investments in protection anecdotally observed in real-world settings are consistent with the temporal patterns one would see if individuals made decisions using simple trial-and-error rules that increase investments in protection after real or imagined losses but not after non-events or false alarms. Investments in protection are shown to be sustained only under the strong condition where the reinforcing strength of counter-factual losses place is as strong as that of equivalent *real* losses.

We then show that decisions over time made in two realistic laboratory tasks replicate this oscillating learning pattern. In an initial study of short-term hurricane mitigation decisions we find that, after controlling for objective storm risk and subjective worry, investments in protection are significantly driven by short-term lags in experienced losses. More importantly, the tendency to reduce investments in protection given the *absence* of past losses is observed regardless of whether the reason for this absence was the lack of a storm event or the presence of past mitigation. Hence, the actual experience of losses, not the counter-factual imagining of them, seemed a necessary condition for increases in investments.

We then replicate these findings in a second study where participants make long-term decisions to invest in adaptive improvements in hypothetical homes against seasonal hurricane threats. We again find evidence that investment decisions are driven by short-term reactions to experienced losses. In this case, however, because the effects of improvements are cumulative, asymptotic protection levels are much higher than those observed in the first study. Yet, even here investments remain below those which would minimize expected losses and are achieved only after experiencing a series of otherwise avoidable losses.

The discussion is partitioned in three phases. We first establish a theoretical background for the work by reviewing the findings of prior research on learning in probabilistic environments, and then explore the implications of applying trial-and-error learning rules to decisions to invest in protection against low-probability losses. We then describe the findings of a program of empirical research that examines learning abilities revealed in controlled laboratory settings. We conclude with a discussion of the implications of the work both for basic research on limits to learning in risky decision contexts and applied research on how to foster investments in mitigation against hazards in real-world settings.

2. Learning to Protect by Trial and Error

Decisions about how much to invest in protection against extreme events are not easy to make. Unlike decisions involving frequently-encountered risks, the probabilities associated with hazards such as earthquakes and hurricanes are often ambiguous, and the conditional effectiveness of investments in protection are rarely known with any certainty. Hence, while a new resident moving to coastal Florida may see wisdom in retro-fitting their home with window protection, there may no objective guidance to help assessments of which window technology would be most cost-effective to install, or indeed whether the investment should be made at all.

Yet, the fact that protective decisions may be difficult to make does not necessarily imply that that they will be made poorly by residents, particularly in the long run. One of the major findings of both theoretical and empirical work in decision making in complex tasks is that knowledge of the normative basis of a decision -- such as knowing the actuarial odds of a hazard—is by no means a necessary condition for behavior to *appear* optimal, at least in the aggregate (e.g., Fudenberg and Levine, 2000). Specifically, as long as individuals are placed in environments where optimal actions more reliably rewarded than sub-optimal actions, decisions made using simple trial-and-error rules will often converge to optimality, even among naïve decision makers who have few insights into normative behavior (e.g., Kalai and Lehrer, 1993; Meyer and Hutchinson, 2001).

This result, however, comes with a critical catch: if the decision environment is *not* ideal for learning—such as if feedback is rare and noisy—one can no longer be assured that trial-and-error policies will discover optimal equilibria (e.g., Camerer, Ho and Chong, 2001). Of particular relevance to the current investigation is recent work by Weber, Shafir, and Blais (2004) and Shafran (2011) who provide examples of how reinforcement rules can distort asymptotic choice frequencies when applied to repeated decisions involving rare events (see also March 1996). In an effort that has parallels to the present, Shafran (2011) posed experimental participants with a series of repeated gambles in which they could pay a fixed fee to reduce the probability of a loss by a known amount. One of the key findings was that even though participants were fully informed about the odds of each gamble, they nevertheless acted as if they updated their beliefs about whether they would win or lose based on the most recently-observed outcomes. As a consequence, participants were more likely to buy protection when faced with higher-probability, lower-consequence, gambles--where losses were observed more frequently--compared to lower-probability, higher-consequence, ones.

The tendency to make decisions by looking at the success or failure of recent outcomes provides a potential mechanism for explaining why investments in protection against rare events may be difficult to sustain in real-world settings. By their nature investments in protection are *counter-reinforcing*; because hazards are rarely encountered, risk-taking will rarely be punished by nature while prudence rarely rewarded. A resident who cancels his/her flood insurance for a year in an effort to save money, for example, is unlikely to be punished by experiencing a flood that year (or the next), while the community that invests large amounts in the construction of levees may fail to see tangible benefits for years (e.g., Browne and Hoyt 2000; Kunreuther, Sanderson, and Vetschera, 1985).

This lack of reinforcement may, in turn, be exacerbated by the censoring effect of the investments themselves; the more successful a protective investment is in precluding harm, the more difficult it becomes to directly observe its benefits. In such cases, reinforcement relies on the ability of the decision maker to imagine the counter-factual of *what would have happened* had the investment *not* been made. If individuals lack effective counter-factual reasoning skills, investments in protection may be, paradoxically, self-extinguishing; the more effective an investment in preventing harm, the more difficult it becomes for decision makers to recall the original need for the investments (e.g., Meyer, 2006).

2.1 The naïve behavior of naïve learning rules

To illustrate these points, consider what a simple reinforcement learning process would predict about how investments in protection would evolve over time in a task that we call the "storm shutter game." It works like this: a decision maker lives in a community that is prone to periodic impacts from major windstorms. At the start of each storm season, the decision maker has the opportunity to install shutters at a cost c that he or she knows will eliminate the risk of losses from potential storms. Specifically, if a storm occurs and shutters are in place the decision maker suffers no damage, but if it occurs and no shutters are in place he or she will experience a loss (with certainty) $l>c$. In each year there is a constant probability $p<1$ of the community being impacted by a storm, and p is sufficiently large that annual investments in mitigation are long-term optimal; i.e., in each year $E(l-c|p)>0$.

To explore whether individuals would learn this optimal long-term policy by trial-and-error alone, consider the case where the probability that the decision maker will invest in mitigation at a given time t, P_t, evolves as

$$P_t=\alpha P_{t-1}+(1-\alpha)x_{t-1}, \quad 1>\alpha>0. \qquad (1)$$

where α is inertia parameter that controls the rate of learning and x_{t-1} is a (0,1)-bounded measure of the degree of positive reinforcement received in year $t-1$ about the benefit of shutters. Expression (1) should be recognized as a variant of the family of linear operator models that have been widely used to characterize associative learning in psychology as well as repeated play in games (e.g., Bush and Mosteller, 1953; Erev and Baron, 2005; Weber, Shafir, and Blais 2004). To extend these formulations to the kind of learning that might take place in natural hazard contexts, we assume that reinforcement (x_{t-1}) can be received through one of two routes: by *direct experience* with real or precluded losses if a storm occurred in year $t-1$ or by *fictitious* *experience* if a storm did not occur. Fictitious experience reflects the strength of the decision maker's beliefs about the losses that *could* have occurred had a storm occurred the previous year without mitigation in place. A simple assumption about this feedback is that in each year x_{t-1} takes the form:

$$x_{t-1} = MAX(s_{t-1}, d_{t-1}) \qquad (2)$$

where s_{t-1}=1 if a storm was experienced in year $t-1$ and 0 otherwise, and d_{t-1} is a (0,1)-bounded measure of the imagined (or counter-factual) loss in a year when there was no storm.

It should be transparent that if counterfactual beliefs are constant (i.e., $d_{t-1}=d$ for all t) and less than one—that is, imagined losses are less effective in encouraging investment than real ones--a learning process described by (1) and (2) will yield a temporally-oscillating pattern of investments where the likelihood of shutters increases to $\alpha P_{t-1}+(1-\alpha)$ in years after a storm but decreases to $\alpha P_{t-1}+(1-\alpha)d$ in years after the absence of a storm, with a long-term mean of $pd+(1-p)$. These dynamics are illustrated in Figure 1, which plots the learning path yielded by equation (1) for the case of a simulated 50-year time horizon in which $p=.2$, the fictitious reinforcement d equals the baseline odds of a storm (.2), and shutter probabilities are moderately inertial (α=.6). The figure displays a pattern of forgetfulness similar to the anecdotal examples provided at the outset, where a storm experience in year t acts to boost mitigation in year t+1, but this diligence decays as the years since the last storm increase.

It is important to emphasize, of course, that this lack of learning is not a necessary property of expressions (1) and (2): they *could* allow a resident to learn a policy of persistent shutter installation simply if one is willing to make more favorable assumptions about the strength of fictitious reinforcement in periods without hazards. For example, Figure 2 plots the learning path that equation (1) yields for the same simulated 50-year horizon in the case when counter-factual beliefs about damage (d_t) evolve as a *cumulative function* of past experienced losses, where $d_t = (d_{t-1}+s_{t-1})/t$. Here, trial-and-error learning continues to evolve as an oscillating, "step forward, step back" process, but is now marked by gradual convergence toward a policy of persistent investment in mitigation.

Which of these characterizations is likely to offer the better portrait of reality? The answer is far from clear. On one hand, because natural environments afford decision makers with a far richer set of cues that could potentially aid learning than is assumed in these examples, one might predict that the second, more optimistic, view will be the better reflection of reality. In natural settings people *are* capable of having long memories, and have frequent opportunities to indirectly learn by observing the experiences of others. One does not have to experience an earthquake in one's own community, for example, to understand their dangers, and news reports of destructive quakes and hurricanes in other areas can serve as compelling reminders of the value of strong building codes and diligent mitigation practices.

On the other hand, one could just as easily argue the opposite, that in natural settings such favorable factors would be offset by psychological influences that act to inhibit learning to an even greater degree than modeled above. First, while individuals may have long memories, the long real-time intervals that typically arise between catastrophic events implies that ongoing-feedback will be dominated by observations of how they are *not* needed, and/or prompt beliefs that the base rate risk of the hazard has abated (the gambler's fallacy; see, e,g., Kahneman and Tversky, 1973). Likewise, while opportunities for indirect learning may well abound, it is not clear that remote observation is an effective replacement for direct experience. For example, in experimental work on "near miss" effects, Dillon and Tinsely (2008) found that close encounters with negative events that did not impose damage (but could have) served to *deflate* rather than reinforce propensities to invest in protection by encouraging misplaced optimism that future threats will also be near-misses. Likewise, similar findings showing that events affecting others do little to influence beliefs about personal risk have been offered by Raghubir and Menon (1998), Weinstein (1980), and others.

Finally, even if forgetting effects are small in natural settings, investments in protection may be asymptotically limited by the inability of experience to cure myopic reasoning (or hyperbolic discounting) biases—the tendency to focus more on up-front costs than longer-run benefits when making inter-temporal tradeoffs (see, e.g., Loewenstein and Prelec, 1992; Thaler 1980). Because, by definition, *all* protective decisions involve a decision to incur up-front costs to achieve a delayed benefit, even decision makers who accurately anticipate the long-run material consequences of failing to invest in protection may still under-invest in the short run due to a tendency to excessively focus on the immediate disutility of cost outlays (e.g., Kunreuther, et al. 2012).

In the next section we describe the findings of a program of empirical work designed to shed some light on this empirical uncertainty. The goal of the work is to examine the ability of individuals to learn optimal strategies for protective investment in a laboratory setting where optimal behavior can be objectively defined, experimental control can be exerted over the magnitude and incidence rates of hazards, and available cues in the learning environment can be precisely measured. In the work we examine both the degree to which empirical learning patterns display short-term feedback effects consistent with trial-and-error learning, and the degree to which these decisions show evidence of counter-factual reasoning, where decisions about how much to invest in protection are supported by beliefs about losses that *could have* occurred. We report the findings of two studies: one focusing on recurrent decisions to make short-term investments in protection against immediate threats (Study 1), and one focusing on longer-term mitigation against temporally distant threats (Study 2).

3. Study 1: Learning to Make Hurricane Preparedness Decisions

3.1 Overview

The study focused on the ability of individuals to learn how to make decisions about investing in preparedness for hurricane threats. This context emerged from a desire for a learning task that would be perceived as involving by participants, and where recurrent decisions to take protective action realistically arise in natural settings. The task had a simple structure: participants were told that they would be living for three years in a fictional country that was prone to periodic hurricane threats, and their goal was to make a series of decisions about how much to invest in protection against approaching storms that would maximize their total net wealth at the end of each year, defined as the initial value of their home minus losses due to hurricane damage and investments in protection. Participants were awarded a $10 show-up fee, and were told that at the end of the study the game records for three participants would be selected at random and paid up to $50 in additional compensation based on their final wealth scores.

3.2 Participants and procedure

Participants were 203 students and staff at a northeastern university who volunteered to participate in exchange for a monetary incentive. Play in the simulation proceeded through five steps, illustrated in the screen shots shown in the appendix:

1. *Basic instructions and orientation*. At the start of the simulation participants were seated at computer terminals displaying the main decision interface *(Figure A-1)*. This interface contained a map showing their home's location in the fictional country

of "Pentonia," as well as text boxes displaying their initial home value, cumulative expenditures on protection and storm losses (0 at the start), and their net wealth. In addition, the interface contained buttons that, when clicked, opened up new screens that provided information about both the relative risk of hurricanes at that location (the historical frequency) plus the cost and effectiveness of mitigation investments. Participants were told that their home had an initial value of 50,000 "Pentonian dollars" and their goal was to make decisions about how much to invest in protection against approaching storms in such a way as to maximize their year-end net wealth, or difference between their home's starting value and cumulative losses due to hurricane damage and expenditures on protection. Participants were told that they faced no budget constraint when buying protection, such that even if their net wealth score was negative at a particular point in time they could still buy whatever levels of protection they felt would be needed to prevent damage from an approaching storm. In addition, participants were told that at the start of each new year the slate would be wiped clean of all previous losses and expenditures on protection, such that each new year could be viewed a fresh run of the simulation[1].

2. *Instructions about potential losses and the role of mitigation.* Participants were told that the amount of damage caused by a storm would depend both on its strength and proximity. A minor tropical storm (category 0 strength), for example, could cause, at most, a 5% loss in home value, while an extreme hurricane (category 4 strength) could potentially destroy their home. These potential losses, however, could be mitigated by purchasing "protection points" each time a storm approached. These

[1] The slate was wiped clean after each year to allow participants to play three independent replicates of the basic decision task, with the only carry-over being expertise.

points were defined on a 0-to-100 scale, and corresponded to the approximate percentage of their home value that would be protected from hurricane damage. Purchased levels of protection remained in place only for the duration of a given storm; with each new storm a new decision would need to be made about how much to invest in protection.

3. *Hurricane development and motion (Figure A-2).* Time progressed in the simulation by clicking a button labeled, “let’s see if there is a hurricane this week.” If a hurricane developed, it appeared as an animated storm icon on the screen, and begin moving toward the Pentonian coast. As the storm intensified, the color of the icon would change to reflect increasing intensity on a 0-to-4 scale,—from green (category 0) to black (category 4). Storm motion, frequency, and intensity was programmed to mimic climatological norms in the North Atlantic basin, becoming more frequent and intense on average as the summer progressed, peaking in September. Likewise, the typical pattern of motion was west-northwest followed by recurvature to the northeast. The points of oceanic origin, direction, and time of recurvature was randomized. Frequencies were set such that participants would typically see 10-14 hurricanes develop over the course of a 20-week (four-month) hurricane season, 6 being major (category 3 or 4). On average 8 of all storms would make landfall some point along the coast.
4. *Protection decisions (Figure A-2).* The moment a storm developed participants were given the opportunity to purchase protection against its potential impact. To simplify this decision, investment levels were presented as a choice between two ordinal categories: limited protection at a low cost or complete protection at a high cost.

Limited protection was implemented by clicking on a button labeled, "board up" on the screen and complete protection was implemented by clicking on a button termed, "evacuate." Repeated clicking of "board up" allowed them to purchase up to 50 points of protection (of a possible 100) in 10-point blocks at a cost of $100 per block. Clicking "evacuate" immediately increased their protection to 100 points at a cost of $2,500. The nonlinear price structure was implemented to simulate the choice residents typically have to prepare for storms either by undertaking relatively low-cost, limited-effectiveness actions (such as buying extra water, filling a car with gas) versus high-cost, high-effectiveness actions (such as installing storm shutters, evacuating to a hotel). To insure that participants had enough time to make these decisions thoughtfully, they were given the ability to pause a hurricane's movement as it approached the coast. There was no time limit on the duration of these pauses, during which they could buy protection and/or read about the merits of protection. To simulate the difficulty of taking protective actions in the midst of a storm, however, the ability to invest in protection was disabled when the storm was within 2 seconds of uninterrupted game time before hitting their home location.

5. *Damage assessment and debriefing scales (Figure A-3).* After each storm (whether it hit land or not) participants were taken to a "damage assessment" screen that revealed how much damage (in dollars) they suffered from the storm, plus the theoretical amount of damage they *would* have suffered had they not invested in protection. After viewing this information, participants were then asked to respond to two short debriefing questions: one asking them to rate the degree to which they felt they had

over- or under-prepared for the storm, and another asking them to rate how worried they were that they might suffer damage during the course of the storm.

After experiencing three seasons of the simulation, participants responded to a short debriefing survey that elicited basic demographic data, their personal experience with hurricanes, and how closely they follow news coverage of hurricane impacts in other areas. This post-study debriefing indicated that while few (8%) had experience being in a hurricane, almost all (92%) indicated that follow new coverage of such storms.

3.3 Results

We partition our discussion of results into two phases. We first offer a general characterization of how investment decisions evolved over time in the task, focusing on the degree to which participants' protective investments departed from optimality. We then provide a more detailed analysis of how decisions were made by estimating a disaggregate model of trial-by-trial investment choices.

3.31 Aggregate Learning Performance.

Depending on the run of the simulation, of the 10 to 14 storms witnessed by participants, slightly more than half (52%) came close enough to threaten damage sufficient to make investments in mitigation worthwhile. Because each home had a value of $50,000 and the cost of maximum (100) mitigation was $2500, participants should have found it worthwhile to purchase maximum protection for any storm that threatened a 5% loss in home value or more---which participants were instructed would be true for all storms of category 1 (hurricane) or higher coming near their residences. Because most storms were of this strength, maximum protection was optimal for 86% of storms that had the potential of causing damage, while the optimal protection level was 50 for the remaining 14%. Intermediate levels less than 50 were never optimal[2]. Hence, for the most part the optimal mitigation policy was a simple binary one: if the storm was sufficiently strong and close to potentially cause damage, it was probably worthwhile to invest in 100% protection.

In Figure 3 we plot actual minus normative investment levels by normative protection and year in the task. The data provide a clear view of achievement: for the most part, participants displayed reasonably good skills at knowing when *not* to invest in protection, but were comparatively poor at knowing *how much* to invest. Specifically, the data show a slight (8%) over-investment bias when encountering storms that did not call for protection, but large and persistent *under*-investment bias when protecting against storms that posed a real threat, ranging from 23% under-investment when limited protection (50) was called for, and 44% under-investment when maximum (100) protection was called for.

Insights into the pattern of individual differences in achievement are provided in Figure 4, which plots a histogram of the average size of the under-investment bias for participants over all storms for which some mitigation was optimal. The figure suggests that the mean bias exhibited a single mode between 35 and 40, suggesting that an analysis of central tendencies provides a reasonable account of participants' behavior in the task.

Another metric for gauging performance is whether participants correctly invested in maximum protection when it was normatively called for by a storm versus erroneously investing in *no* protection at all when maximum protection was optimal. The data suggest that participants developed improved skills in achieving the former and avoiding the latter errors with experience in the task, but skill levels in the former were modest. Specifically, when maximum investment

[2] This optimality pattern was a consequence of the storm damage function programmed in the simulation, where the potential home loss was a discrete step function of a storm's strength and proximity.

was optimal participants did so only 33% of the time in year one, with this percentage increasing to 50% in year three. Conversely, the percentage mistakenly investing when none was called for decreased from 14% in year one to 10% in year three.

A more detailed longitudinal view of learning dynamics is provided in Figure 5, which plots mean investments in protection by each storm experienced by participants each of the three seasons of play, separating the learning paths for cases where maximum (100) mitigation was worthwhile for a given storm (solid lines in the figure) and when no mitigation was called for (dashed lines). Because of the comparative sparseness of cases of storms where limited mitigation was optimal, we plot only these two extreme cases. In addition, because the number of storms varied across runs of the simulation, the figure plots the averages for the first ten storms to insure comparable sample sizes (189) for each storm event.

Echoing the key insights from Figure 3, Figure 5 shows a gradual attenuation of the under-investment bias for maximum mitigation with experience in the task. The more interesting feature, however, is the shape of this learning path. When encountering storms that called for maximum investments in protection (the solid line in the figure), learning displays a discontinuous, “two steps forward, one step back” pattern: investments in protection increased over time during the course of the first simulated season but then regressed at the start of the next season.[3] By the third season we see no such regression at the start of the year, however mean investment (38.2%) was well below optimal.

Why did we observe this annual regression? Some initial insights into this are provided in Figure 6, where we plot the mean size of the avoidable loss incurred by participants as a

[3] The sharp decrease in mitigation observed at the very end of the simulation shown in Figure 4 might be attributed to an end-game effect, where the impending end of tenure in the home triggered a reluctance to make any fixed investments in protection—even though such a reduction would have no rational basis given the payoff structure of the game.

percentage of home value by storm for each year. The figure provides a possible explanation for the investment evolution: because storms tended to be stronger as each season progressed (following climatological norms), failing to invest in maximum protection would have had milder absolute consequences earlier in the season compared to later. Hence, participants appeared to correctly anticipate that storms earlier in the season would pose less of a risk, but failed to grasp that even milder storms were worth fully protecting against. The average amount invested thus seemed to proportionally track the magnitude of the average experienced loss.

Finally, it is worth noting that the costs to participants of the tendency to under-invest in protection were substantial for participants as measured by loss of total wealth, which was the scoring criterion in the task. In the first year of the simulation, participants' mitigation and loss costs averaged 274% higher per storm than those that would have been yielded had they made all decisions following an optimal mitigation policy, and remained at 186% above optimality in the third year.

3.32 Individual Mitigation Decisions: Modeling the Effect of Trial-by-Trial Feedback.

To provide a more rigorous account of the processes underlying individual decision making in the task, we estimated a family of disaggregate models describing the drivers of trial-by-trial decisions about whether and how much to mitigate. These analyses examine the degree to which the aggregate investment patterns displayed in Figures 3,4, and 5 could be explained by participants' engaging in a trial-and-error learning process that increased investments after experienced losses, but deflated investments after non-events. Such a learning process would be evidenced by a significant lag effect of experienced storm losses after controlling for individual mean propensities to invest in protection, the objective magnitude of the immediate threat (the storm's strength and proximity), and mean time (e.g., seasonal) trends. In addition, we investigated the degree to which investment decisions may have also been reinforced by two potential sources of fictitious or counter-factual feedback:

1. *Mitigated losses*, that is, losses that were avoided due to prior investments in protection; and
2. *Near-misses*, that is, losses that were avoided because of the storm's making landfall too far from the participant's home to risk damage.

If participants engaged in strong counter-factual reasoning when learning from past storm events, measures describing the *potential* losses imposed by a given storm—not the experienced losses themselves—should be the primary statistical driver of subsequent investment decisions, mediating the direct effect of experienced losses.

A challenge to this modeling task was that the data exhibited three complicating features: a left-censoring of investment levels in cases where a storm was too weak or too distant to cause damage (see Figure 3), a non-uniform distribution of protection choices available to participants when a storm *did* pose such a risk (10-50 or 100), and the likely presence of heterogeneity in attitudes toward investing in protection. We addressed these requirements by modeling trial-by-trial investment decisions as two differential-effect ordered logit models, one capturing participants' binary decisions whether or not to invest in protection and the other the conditional ordered quantity decision.[4] Hence, the analysis assumed that the probability that a decision maker i would be undertake the specific mitigation of level $j=k$ in response to storm t could be represented by the two-stage structure

[4] We also estimated an ordered logit model that assumed decisions of both whether and how much to invest were driven by a single, homogeneous process. This single stage model, however, provided a poorer account of the data than one that partitioned the data into two stages.

$$Pr(j{=}k|t)_i{=}Pr(j{>}0|t)_i \times Pr(j{=}k|j{>}0,t)_i \qquad (3)$$

where $j=0,1,..N$ indexes increasingly high ordinal levels of mitigation. Formally, the probability that individual i would invest in mitigation level k in response to storm t, $Pr(j=k|t)_i$ was represented by the binary and ordered logits

$$Pr(j{>}0|t)_i{=}\frac{1}{1+\exp(\beta'_{i0}X_t)} \qquad (4)$$

$$Pr(j{=}k|j{>}0,\, t)_i{=}\Big[\frac{1}{1+exp(\beta'_{im}X_t)} - \Big(\textstyle\sum_{j=1}^{k-1} \frac{1}{1+exp(\beta'_{im}X_t-\delta_j}\Big)\Big] \qquad (5)$$

Where β'_{io} and β'_{im} are vectors of the coefficients describing the effect of a common set of measured predictors X on decisions whether and how much to mitigate respectively, and δ_j is the threshold intercept for migration category j. Because investment levels were restricted to 6 discrete levels (10, 20, 30, 40, 50, or 100), the ordered choice model estimated five category intercepts.

We first estimated base models that represented decisions whether and how much to mitigate for each storm within each season as a function of five groups of explanatory variables:

1. 202 individual effect-codes designed to capture mean individual propensities to invest in protection across all storms in the simulation (the number of participants minus one);
2. Yearly and seasonal trends in mean investment levels, represented by year, storm number within year, and their interaction;

3. The objective threat posed by the storm, measured by its intensity (magnitude), distance of closest approach to the participants' residence (minimum distance) and their interaction;
4. Participants' stated degree of worry in advance of the current storm (a 3-point scale; modeled as two dummy-coded effect variables); and
5. First- and second-order lags in both real losses from previous storms.[5]

It was this last group of variables that formed the primary focus of the analysis; we wished to measure the degree to which decisions about protection were influenced by previous experienced losses after controlling for possible contemporaneous individual and situational drivers of these decisions.

In addition to these base models, we estimated two additional forms that designed to estimate the effect of three foregone loss measures:

1. *Foregone loss*, or the size of the loss that would have occurred given no mitigation from the last storm;
2. *Mitigation Benefit*, or the loss that was avoided from the last storm due to mitigation; and
3. *Near miss event*, a binary indicator of whether the last storm was strong enough to (category 4) to risk catastrophic damage, but passed far enough away to risk limited damage (defined as less than half the mean potential loss).

In Table 1 we report maximum likelihood estimates of the base specification of expressions (3) and (4). One initial insight that emerges from the table is that different

[5] In exploratory analyses we also examined models recognizing higher-order lag effects (up to fifth order), however, due to the loss of efficiency in estimates of higher-order lags, we focus here on the first two.

statistically processes appeared to drive the decisions of *whether* versus *how much* to mitigate. Decisions *whether* to mitigate were driven largely by one predictor: the proximity of the storm's approach to the participants' residence, regardless of strength and past experience. In contrast, conditional decisions of *how much* to mitigate were driven by a richer array of predictors: after controlling for individual mean differences in investment levels, protective investments for each storm were an increasing function of the proximity of the storm and its strength, the year in the simulation, and, most critically for this research, the size of the loss experienced from the last two encountered storms. Specifically, consistent with a reinforcement learning process, the data reveal a highly significant, positive, first-order lag effect of experienced losses (χ^2=23.63; $p<.0001$) and a weaker second-order lag (χ^2=7.55; $p=.006$). In other words, the greater (or weaker) the previous experienced storm losses, the greater (or weaker) the current investment in protection.

Evidence on the degree to which participants acted as if they engaged in counter-factual reasoning when making investment decisions is provided in Table 2, which reports maximum likelihood estimates of expression (5) when the predictors of the base model are expanded to recognize the effects of absolute potential losses (model 2a), precluded losses (model 2b), and near-miss events (model 2c). The data provide two, possibly surprising, insights. First, the data provide no evidence that participants' decisions were influenced by beliefs about the amount of damage that *could* have occurred from the last storm; neither lag absolute damage potential (2a) or precluded damage (2b) emerge as significant predictors of investment levels, while lag *experienced* losses remain robustly significant in both specifications. Hence, participants did not appear to consider take foregone losses into account when learning about the wisdom of protective investments; a failure to observe losses from one storm had the same suppressing effect on investments in protection against the next storm regardless of whether the absence of a loss was due to the absence of a threat or the presence of protection.

In contrast, the data *do* reveal evidence of a small “near-miss” effect when the previous storm was of category-4 strength but did not pass close enough to the participant’s location to risk high levels of potential damage. Consistent with a previous laboratory finding on near-miss effects reported above by Dillon and Tinsley (2008), we find that the effect of a near miss on preparation was negative; witnessing an extreme storm that *could have* had a major impact caused participants to marginally reduce their average investments in mitigation in advance of the next storm (χ^2=4.68; p=.03). We should note, however, that this effect appeared limited to the case of extreme storm events; defining a near-miss in terms of storms of strength 3 or higher did yield a similar significant effect.

4. Study 2: The Case of Long-Term Mitigation

4.1 Motivation and Method

The first study focused on biases that arose in the specific case of recurrent short-term mitigation decisions, where protective investments made for one storm had no carry-over benefit to the next storm. In many (or most) natural settings, of course, protective investments have at least some longer-term benefit; a decision to invest in storm shutters or a generator in advance of one hurricane, for example, will provide protective benefits long after a single storm season has passed, even short-term food and water supplies can have benefits for multiple storms. As such, the permanent nature of some types of mitigation investments provide a natural source of insurance against forgetting; as long as individuals are sufficiently motivated *at some point* to invest in protection, the benefit will remain even after the original perceptions of risk that motivated the investment have vanished.

On the other hand, longer-term investments in protection typically involve higher up-front costs and involve greater degrees of uncertainty than the short-term decisions studied above—factors that might serve to inflate rather than suppress the under-investment bias discovered above. If participants in Study 1 found that they under-or-over protected for a single storm, the slate was quickly wiped clean, and they had ample opportunities to make better decisions for future storms. In addition, with each storm, they faced comparatively little uncertainty about the *need* for protective investments, as they could delay purchasing protection until the size of the storm threat was known with virtual certainty. As noted above, participants were quite skilled at making this distinction, with the primary bias arising in knowing *how much* to invest given that some protection was needed.

To investigate individuals' abilities to learn to invest in longer-term mitigation,130 new participants completed an alternative version of the hurricane simulation in which investments in protection, once made, did not vanish after a storm had passed, but rather remained in place for future storms. The basic structure of the simulation was identical to that described above except for a major change in how protective investment decisions were made. Specifically, as in Study 1, participants were endowed with a home that had no storm protection, but in this case were told that they could purchase permanent or semi-permanent improvements that would increase its protection rating.

In this new design there were two types of improvements that participants could purchase: *permanent improvements* (such as shutters and roof bracings) that would stay with the house as long as they owned it, and *seasonal improvements* (such as supplies and landscape trimming,) that would stay with the house for the duration of a hurricane season, but needed to be re-purchased at the start of each new season. Permanent improvements were worth up to 50 protection points, and seasonal improvements another 50. Hence, a participant who purchased all possible improvements by the end of a given season would have a protection rating of 100 at the end of that season, but this would drop to 50 at the start of the next. Participants could purchase either type of improvement at any point in the simulation; they could buy all possible improvements at the start, or wait to see how the storm season evolved. Unlike Study 1, however, once a storm formed and began moving toward the coast, no mitigation improvements could be purchased; all decisions in a given simulation week had to be made prior to observing whether a storm would form that week.

In order to provide participants with some flexibility in revisiting their decisions about how much to invest in permanent protection, participants were told that they would be living in three different homes in "Pentonia" for each of two years. Each new home came endowed with no protection rating, hence each time they moved to a new home they had to make decisions about *both* seasonal and permanent improvements.

As in Study 1, participants were paid a $10 show-up fee. Participants were told that at the end of the study the records of two participants would be chosen at random and awarded up to $50 depending on final wealth scores.

4.2 Results

The price of both long-term and seasonal mitigation was set such that it was optimal for all participants to purchase maximum (100) protection for each home. Since there was no discounting (participants could not invest funds used to purchase protection for other purposes) the task thus had a simple optimal dynamic investment policy: purchase maximum levels of permanent and seasonal improvements as soon as possible given each new home and each new year.

In Figure 7 we plot the average protection rating of homes owned by participants by home number, year, and month, in this new version of the simulation. The figure offers two important insights about how the learning of investment behavior in enduring protection differed from those of short-term protection in Study 1. First, on the positive side, mean levels of protection are asymptotically higher than those observed in Study 1 when the same (100%) level of protection was optimal; while the mean investment level in such cases in Study 1 never consistently exceeded 65% with experience, by the third home and sixth year of ownership, protection levels provided through permanent and seasonal mitigation were asymptotically over 80%. On the negative side, however, this mean asymptotic investment level remained significantly below that which would have been optimal for the task, and the learning path displayed the same, relatively slow, "two steps forward, one step back" pattern observed in short-term investments (Figure 5). In essence, the learning path is one consistent with participants increasing their investments in structural improvements as each season progressed, but when seasonal improvements were withdrawn at the end of each year, they were not immediately replaced at the start of the next, escalating only as a result of storm experience in the course of each year.

To provide a more detailed analysis of mitigation decisions, we modeled the week-by-week changes in overall protection levels in each week of the simulation as a function of:

1. 129 individual-specific effects capturing mean individual differences in protection propensities;

2. 5 home-location effects capturing differences in objective risk exposure (storm impact frequencies) on protection levels across home sites;
3. Lag levels of protection;
4. Time point in the simulation (home number, year, month, and week); and
5. Lag value of loss in the previous week.

OLS estimates of model parameters[6], reported in Table 3, resemble those of Study 1: after controlling for mean differences in objective risk across locations (the frequency of storm encounters), individual differences in protection rates, and lag levels of protection, we find a significant positive effect of the size of the loss in the previous week. Note that unlike Study 1, protection, once it was invested in, was not withdrawn for the duration of a given year, so the lag effect here has a more limited, asymmetric, meaning compared to Study 1: it implies that a loss in one week tended to encourage investments that boosted protection levels the next—but here the *absence* of a loss, by constraint, does not lead to a *reduction* in protection.

5. Discussion

The goal of this research was to investigate the empirical basis of a criticism of decision making that routinely emerges in post-mortem accounts of catastrophic events: that individuals and communities are prone to under-invest in protection, a flaw that needlessly inflates losses (e.g., Daniels, Kettl, and Kunreuther, 2006). One need not go far for examples: the floods of Hurricane Katrina in 2005, for example, were widely attributed to failures to invest in the upkeep of levees built after Hurricane Betsy in 1965 (Brinkley, 2006; Michel-Kerjan 2010), the catastrophic losses from the 2009 earthquake in Haiti were widely seen as having been avoidable had greater investments been made in the enforcement of building codes (e.g., Bajak, 2010), and

[6] Mixed-model maximum-likelihood estimates derived under the assumption of more general (e.g., AR-1) error structures yield similar results.

the Deepwater Horizon disaster in the Gulf of Mexico in 2011 was widely blamed on BP's failure to invest in precautions that could have averted or at least limited the resulting oil spill. But hindsight is 20/20; all decisions under risk come with the chance that they will sometimes not pan out, hence it is difficult to say with certainty whether examples such as these serve to illustrate fundamental limitations in our ability to learn to protect against hazards or the mere fact that risky choices sometimes yield losses.

The goal of this paper was to take a step toward trying to resolve this question by exploring both theoretically and experimentally whether there are limits to our ability to learn to take protective action against low-probability, high-consequence events. We first showed how simple trial-and-error learning rules that work well when used for learning to perform frequently-repeated tasks with clear feedback tend to be far less effective—and can produce dysfunctional behaviors – when applied to settings where feedback is rare, noisy, and censored. Specifically, if memories of past losses are prone to erosion, such rules can produce the kind of "two steps forward, one step back" patterns of investment in protection over time that anecdotally seem to arise in real-world settings, such as recurrent market bubbles and crashes (see, e.g., Hussam, Porter, and smith 2008).

We then examined actual patterns of learning from experience observed in two controlled laboratory studies where participants made repeated decisions about their level of preparation against hurricane threats. Across both studies mean levels of an investment in protection were asymptotically below optimal levels, with this bias persisting regardless of whether decisions were being made for short-term protection (Study1) or longer-term (Study 2). Likewise, consistent with the predictions of a trial-and-error learning process, the amount invested in protection in one time period was an increasing function the size of the loss incurred during the last. Perhaps most disturbingly, the data also offered evidence that in order for losses to boost investment in protection they had to be real, not imagined; when this issues was investigated in Study 1 average investments in protection in advance of one storm declined whenever the previous storm failed to induce losses, *regardless* of whether the absence of loss was due to the success of previous investments in protection or the absence of a physical threat. In short, the direct experience of a loss emerged as a necessary condition for increases in investments in protection.

What is noteworthy about the persistent under-investment biases is that they were observed in controlled settings where decision makers had a clear financial motivation to make optimal protective decisions, and they received far more feedback about the effectiveness of their decisions than would have been the case in a real-world setting. After each storm encounter, participants were given information about what their losses would have been had they not invested in mitigation, they faced no budget constraints that might otherwise limit protective investments, and the simulation provided participants with precise information about the long-term value of different levels of protection.

While the tendency to decrease investments after failing to experience losses provides an explanation for the persistent pattern of under-investment, it should be emphasized that other mechanism may have worked to inflate the size of the effect. For example, one possibility is that investments in protective decisions were also suppressed by a tendency for out-of-pocket costs to loom larger than delayed opportunity costs in decision making (e.g., Loewenstein and Prelec 1992; Thaler 1980). To illustrate, while participants in Study 1 were likely aware that an approaching storm would induce damage, at the time the decision was made the magnitude of this consequence would have been both uncertain and temporally distant, whereas the cost of protection was both certain and immediate. While this explanation for why people may under-invest in protection is not a new one (e.g., Kunreuther et al., 2012), what is intriguing here is the evidence that it both can arise when the time difference between the expenditure (on protection) and the benefit (the preclusion of losses) is quite short (mere seconds), and that the under-investment bias did not vanish with decision making experience.

An obvious question for future work is to explore the degree to which the findings here are predictive of decision-making biases that arise in real-world contexts. Although the simulations were designed to re-create many of the real-world features of hurricane threats and mitigation, protective investments in natural settings are complicated by a range of factors not reproduced in here—factors that could either deflate or inflate the strength of the biases uncovered here. For example, perhaps the most salient difference is that our lab simulations did not reproduce the social context in which many preparedness decisions are made. Whether group-influenced decisions will be safer than individual decisions, however, is far from clear; while social settings provide greater opportunities for learning from the actions of others, it could also have the detrimental effect of making norms of under-investment harder to change. Likewise, another factor that could explain failures to learn by communities is the ironic effect of adaptation by out-migration. If the residents who suffer the greatest losses from disasters learn that the best response is to move away to low-risk areas, the remaining population is left without the critical experiential knowledge base essential for seeing the value in protection (see, e.g., Smith, et al. 2006).

This gives rise to a final question: assuming that the biases reported here *are*, at least to some degree, characteristic of real-world decision making, how might they be overcome? This is may be a difficult since the psychological mechanisms that lead to the biases may be hard-wired;

as long as we remain present-focused, prone to chasing short-term rewards and avoiding short-term punishment, it is unlikely that individuals and institutions will learn to undertake optimal levels of protective investment by experience alone. The key, therefore, is introducing decision architectures that allow individuals to overcome these biases through, for example, creative use of defaults (see, e.g., Goldstein, et al. 2008). As an illustration, protection levels in Study 1 would like have been higher if the choice task had been restructured so that participants were deciding whether to *remove* protection (for a cash rebate) in advance of storms that they felt posed little risk. Likewise, if participants in Study 2 were initially endowed with a fully-protected home, and the task was to decide whether to accept compensation for an alternative residence that was less well protected against storm threats.

Finally, while this work was set in the context of natural hazards, it should be clear that it potentially holds implications for understanding protective decision making in a broader range of risky decision-making contexts, such as financial investments and health care. Like emergency managers faced with the problem of convincing homeowners to invest in protection against natural hazards, doctors face similar challenges trying to convince patients to stay on medication regimens; persuading patients to take medication when they are experiencing pain is easy; convincing them to stay on that medication when the pain has been temporarily erased is much harder. Another important area for future research will thus be to explore the communalities and differences of under-investment biases as they arise across this range of settings.

References

Bajak, Frank (2010), "Chile was Ready for Quake, Haiti Wasn't", http://www.salon.com/news/2010/02/27/cb_tale_of_two_earthquakes)

Brinkley, Douglas (2006). *The Great Deluge: Hurricane Katrina, New Orleans, and the Mississippi Gulf Coast,* New York: Harper Collins.

Browne, Mark J., and Robert E. Hoyt (2000), "The Demand for Flood Insurance: Empirical Evidence", *Journal of Risk and Uncertainty*, 20, 291-306.

Bush, Robert R., and Frederick Mosteller (1953). "A Stochastic Model with Applications to Learning" *Annals of Mathematical Statistics*, 24(4), 559-585.

Camerer, Colin, Teck-Hua Ho, and , J. K. Chong (2001). "Behavioral Game Theory: Thinking, Learning and Teaching," Paper Presented at the Nobel Prize Symposium (Dec 2001).

Daniels, Ronald J., Donald F. Kettl, and Howard Kunreuther, eds. (2006). *On Risk and Disaster: Lessons from Hurricane Katrina.* Philadelphia: University of Pennsylvania Press.

Denrell, Jerker, "Adaptive Learning and Risk Taking", *Psychological Review*, 114(1), 177-187.

Dillion, Robin L. and Catherine H. Tinsley (2008). "How near-misses influence decision making under risk: A missed opportunity for learning." *Management Science* 54 (8), 1425-1440.

Erev, Ido, and Al Barron (2005). "On adaptation, maximization, and reinforcement learning among cognitive strategies " *Psychological Review*, 112, 912-931.

Goldstein, Daniel G., Eric J. Johnson, Andreas Herrmann, and Mark Heitmann (2008). " Nudge your customers toward better choices". *Harvard Business Review***,** 86(12), 99-105.

Fudenberg, Drew and David Levine (2000). *The Theory of Learning in Games*, Cambridge, MA: MIT Press.

Hussam, Reshmaan N., David Porter and Vernon L. Smith (2008), "Thar She Blows: Can Stock Bubbles be Rekindled with Experienced Subjects? *American Economic Review,* 98(3), 924-37.

Kahneman, Daniel and Amos Tversky (1973). "On the Psychology of Prediction", *Psychological Review*, 80, 237-257.

Kalai, Ehud and Ehud Lehrer (1993). "Rational Learning Leads to Nash Equilibrium", Econometrica, 61(5), 1019-1045.

Klayman, Joshua and Young-won Ha (1987). "Confirmation, Disconfirmation, and Information in Hypothesis Testing", *Psychological Review,* 94(2), 211-228.

Kunreuther, H., Robert J. Meyer and. Erwann Michel-Kerjan (2012). "Overcoming Decision Biases to Reduce Losses from Natural Disasters", in E. Shafir (ed.), *Behavioral Foundations of Policy*, Princeton University Press, in press

Kunreuther, Howard, Warren Sanderson and Rudolf. Vetschera (1985). "A Behavioral Model of the Adoption of Protective Activities", *Journal of Economic Behavior and Organization*, 6, 1-15.

Lerner, Jennifer S., Roxana M. Gonzales, Deborah A. Small and Baruch Fischhoff (2003). "Emotion and perceived risks of terrorism: A national field experiment." *Psychological Science* 14(2), 144-150.

Loewenstein, George, and Drazen D. Prelec (1992). "Anomalies in Intertemporal Choice: Evidence and an Interpretation", *The Quarterly Journal of Economics* 107, 573–597.

March, James G. (1996), "Learning to be Risk Averse", *Psychological Review,* 103, 309-313.

Meyer, Robert (2006), “Why we Under-Prepare for Hazards”, in R.J. Daniels, D.F. Kettl, and H. Kunreuther (eds.), *On Risk and Disaster: Lessons from Hurricane Katrina.* Philadelphia: University of Pennsylvania Press. Pp. 153-173.

Meyer, Robert J. and J. Wesley Hutchinson (2001), “Bumbling Geniuses: The Power of Everyday Reasoning on Multi-State Decision Making”, in Steve Hoch and Howard Kunreuther (eds.), *Wharton on Making Decisions*. New York: John Wiley and Sons, 37-61.

Michel-Kerjan, Erwann (2010), “Catastrophe Economics: The National Flood Insurance Program”, *Journal of Economic Perspectives*, 24(4), 165-186.

Michel-Kerjan, Erwann, Sabine Lemoyne de Forges and Howard Kunreuther (2012). “Policy Tenure under the Federal Flood Insurance Program”, *Risk Analysis*, in press.

Raghubir, Priya and Geeta Menon (1998). “AIDS and Me, Never the Twain Shall Meet: The Effects of Information Accessibility on Judgments of Risk and Advertising Effectiveness,” *Journal of Consumer Research,* 25 (June), 52-63.

Shafran, Aric P. (2011). “Self Protection against repeated low probability risks”, *Journal of Risk and Uncertainty*” 42, 263-285.

Smith, V. Kerry, Jared C. Carbone, Jaren C. Pope, Daniel G. Hallstron, and Michael E. Darden (2006). “Adjusting to Natural Hazards”, *Journal of Risk and Uncertainty*” 33, 37-54.

Thaler, Richard (1980), “Toward a Positive Theory of Consumer Choice”, *Journal of Economic Behavior and Organization,* 1, 39-60.

Viscusi, W. Kip (1979). “Insurance and Individual Incentives in Adaptive Contexts”, *Econometrica*, 47, 1195-1208.

Weber, Elke, Sharoni Shafir, and Anne Rene Blais (2004). Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review, 111,* 430-445.

Weinstein, Neil D. (1980). "Unrealistic optimism about future life events," *Journal of Personality and Social Psychology,* 39, 806-820.

Wilkinson, Claire (2008). "The California Earthquake Authority", Briefing, Insurance Information Institute, http://server.iii.org/yy_obj_data/binary/794172_1_0/CEA.pdf.

Table 1: Estimation Results for a baseline logit models of mitigation choice (top) and conditional mitigation quantity. Note that all parameters are reverse-signed.

Unconditional Mitigation Choice Model (Binary Logit)				
Term	**DF**	**Estimate**	**Chi Square**	**Pr>ChiSquare**
Subjects	202		398.5367	<.0001
Year 2 (dummy)	1	0.1068	0.4795	0.4887
Year 3 (dummy)	1	0.00267	0.0003	0.9874
Storm number	1	0.01240	0.1295	0.7190
Magnitude	1	0.2299	1.7335	0.1880
minimum distance	1	0.00359	78.4652	<.0001
mag.*min. distance	1	-0.00026	3.7777	0.0519
Worry-med(dummy)	1	-3.5895	301.2110	<.0001
Worry-high(dummy)	1	-5.0293	253.6135	<.0001
lag1loss	1	0.000009	1.0445	0.3068
lag2loss	1	0.000007	0.0054	0.9414
Model LR: 3316.685 Pr>ChiSquare(211) <.0001				

Conditional Mitigation Quantity Model (Ordered Logit)				
Term	**DF**	**Estimate**	**Chi Square**	**Pr>ChiSquare**
Subjects	201	-	591.2913	<.0001
Intercept 10	1	-.8715	0.0020	0.9646
Intercept 20	1	.0305	0.0000	0.9988
Intercept 30	1	.5786	0.0009	0.9765
Intercept 40	1	.9846	0.0025	0.9600
Intercept 50	1	2.5780	0.0173	0.8955
Year 2 (dummy)	1	-0.5341	22.4008	<.0001
Year 3 (dummy)	1	-0.8468	47.7501	<.0001
Storm number	1	0.0249	1.0107	0.3147
Magnitude	1	-0.5670	36.1629	<.0001
min. distance	1	0.00112	14.3324	0.0002
Mag x min. distance	1	0.00018	3.6422	0.0563
Worry-med(dummy)	1	-.5377	8.2784	.0004
Worry-high(dummy)	1	-2.6620	142.1904	<.0001
lag1loss	1	-0.0004	22.2738	<.0001
lag2loss	1	-0.0002	8.8517	0.0060
Model LR: 1430.8194 Pr>ChiSquare(210) <.0001				

Table 2: Estimation Results for conditional ordered logit models of mitigation quantity including total precluded losses (2A) and near-miss event indicator (2B; observation of a category-4 storm on the last occasion that did not risk damage). Note that all parameters are reverse-signed.

Conditional Mitigation Quantity Model Recognizing Precluded Losses (Ordered Logit)					
		Model 2A		**Model 2B**	
Term	**DF**	**Estimate**	**Pr>ChiSquare**	**Estimate**	**Pr>ChiSquare**
Subjects	201		<.0001		<.0001
Intercept 10	1	-0.8672	0.9648	-0.7833	0.9681
Intercept 20	1	0.0364	0.9985	0.1238	0.9950
Intercept 30	1	.5866	0.9762	0.6733	0.9726
Intercept 40	1	.9943	0.9596	1.0808	0.9560
Intercept 50	1	2.5905	0.8950	2.6843	0.8910
Year 2 (dummy)	1	-0.5335	<.0001	-0.2904	<.0001
Year 3 (dummy)	1	-0.8441	<.0001	-0.8391	<.0001
storm number	1	-.0199	0.4323	-.0335	0.1934
Magnitude	1	-0.5943	<.0001	-0.6074	<.0001
minimum distance	1	0.00110	0.0002	0.00109	0.0002
Mag x min distance	1	0.00019	0.0489	0.00020	0.0039
Worry-med (dummy)	1	-.5361	0.0041	-.5402	<.0037
Worry-high (dummy)		-2.6662	<.0001	-2.6681	<.0001
lag1loss	1	-0.00004	<.0001	-0.00003	<.0001
lag2loss	1	-0.00002	0.0115	-0.00002	0.0098
lag1Benefit	1	3.94E-6	0.3189	4.977E-6	0.3177
lag2Benefit	1	-6.22E-6	0.1386	-5.87E-6	0.2093
Near Miss	1	-		0.9037	0.0305
Model LR		1434.1561		1448.3891	
Pr>ChiSquare			<.0001		<.0001

Table 3: OLS estimates of a model of trial-by-trial changes in protection levels, Study 2.

General Linear Model of Long-Term Protection Levels, Study 2				
Term	**DF**	**Estimate**	**F/t**	**Pr>F/t**
Subjects	129		461.39	<.0001
Location 1	1	-5.96709	-17.17	<.0001
Location 2	1	-4.31430	-12.03	<.0001
Location 3	1	-4.65523	-12.48	<.0001
Location 4	1	-2.91331	-8.05	<.0001
Location 5	1	-2.42959	-7.44	<.0001
Lag Loss	1	0.00003	4.00	<.0001
Year	1	1.05840	5.98	<.0001
Month	1	0.53989	8.52	<.0001
Week	1	0.11171	1.42	0.1568
House number	1	1.71173	13.57	<.0001
Lag protection	1	0.81767	205.36	<.0001
Model	140		868.31	<.0001

R^2=.927

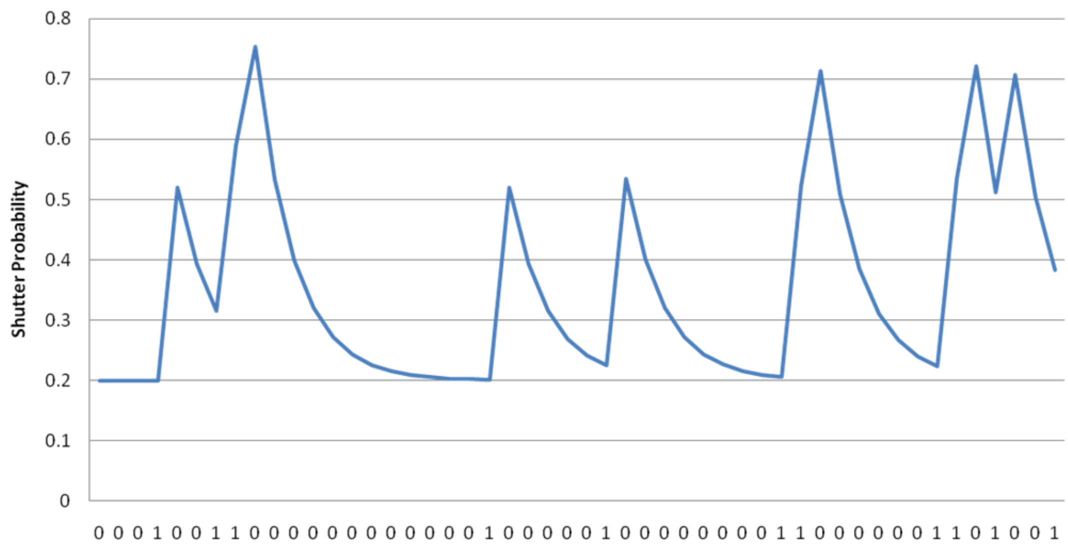

Figure 1: Simulated evolution of shutter probabilities given a simple reinforcement learning process. “1” on the ordinate axis denotes a storm incident in a given year.

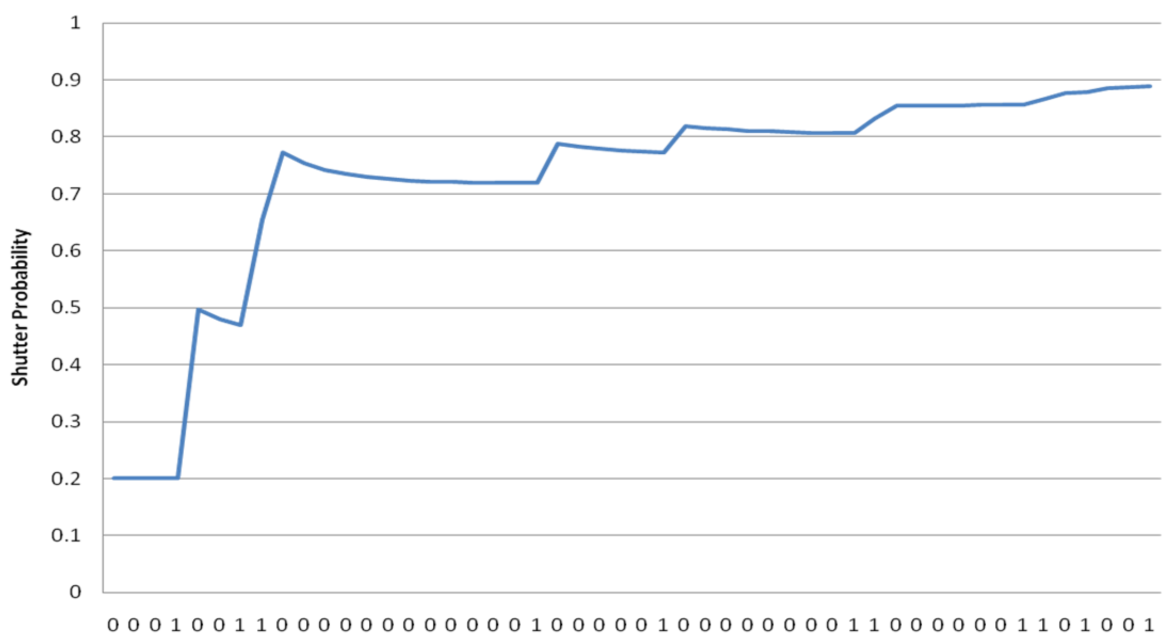

Figure 2: Modification of the learning path in Figure 1 when fictitious reinforcement is monotonically increasing in cumulative storm experience

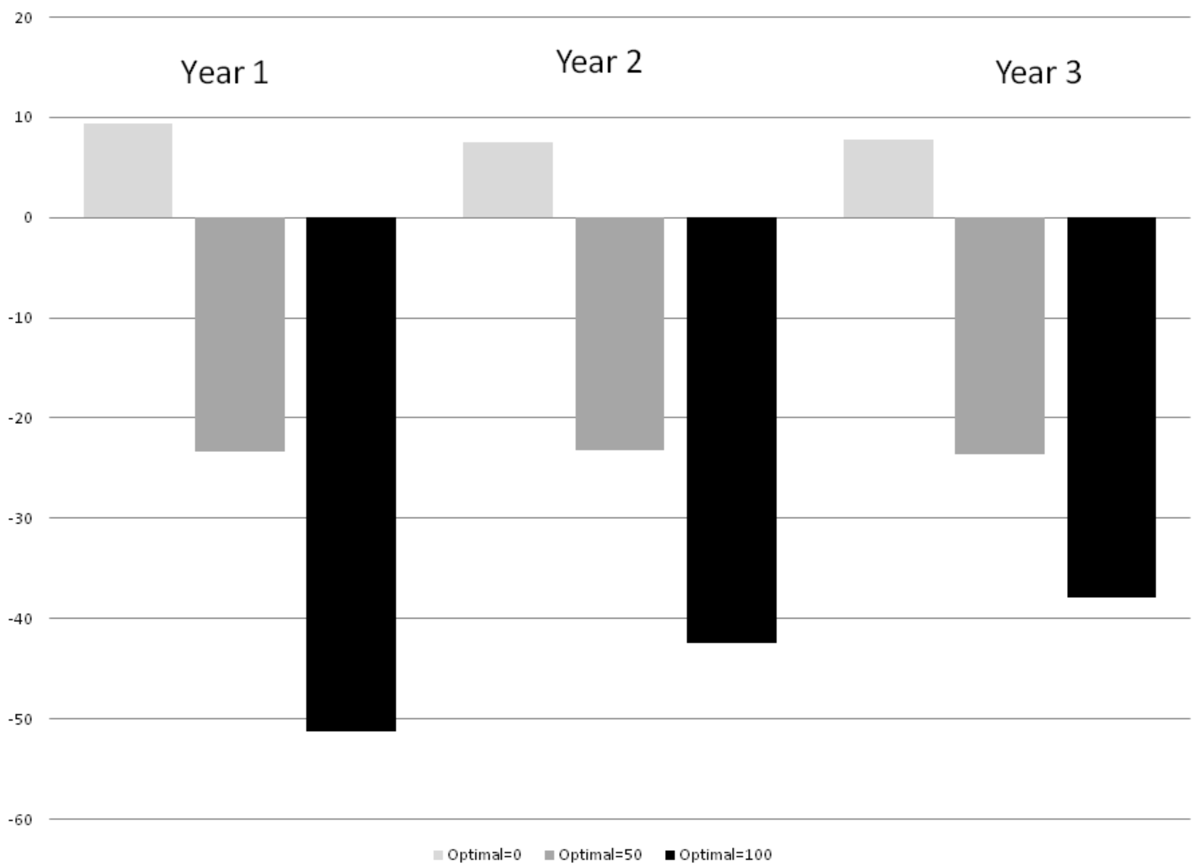

Figure 3: Actual minus optimal storm protection levels by simulation year. Black bars represent cases where the optimal protection level was 100, dark gray 50, and light gray 0.

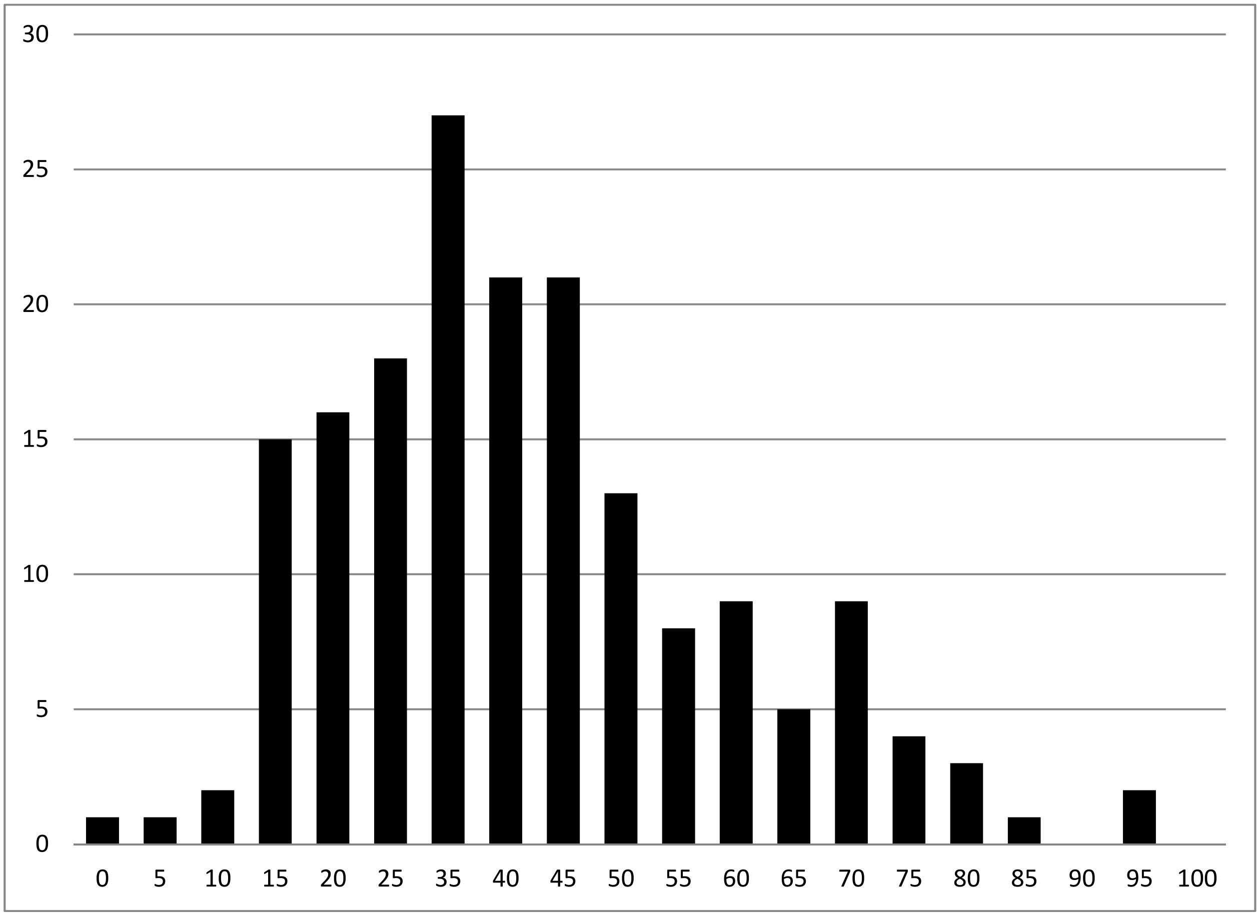

Figure 4: Distribution of mean under-investment errors across participants in cases where mitigation was optimal, Study 1.

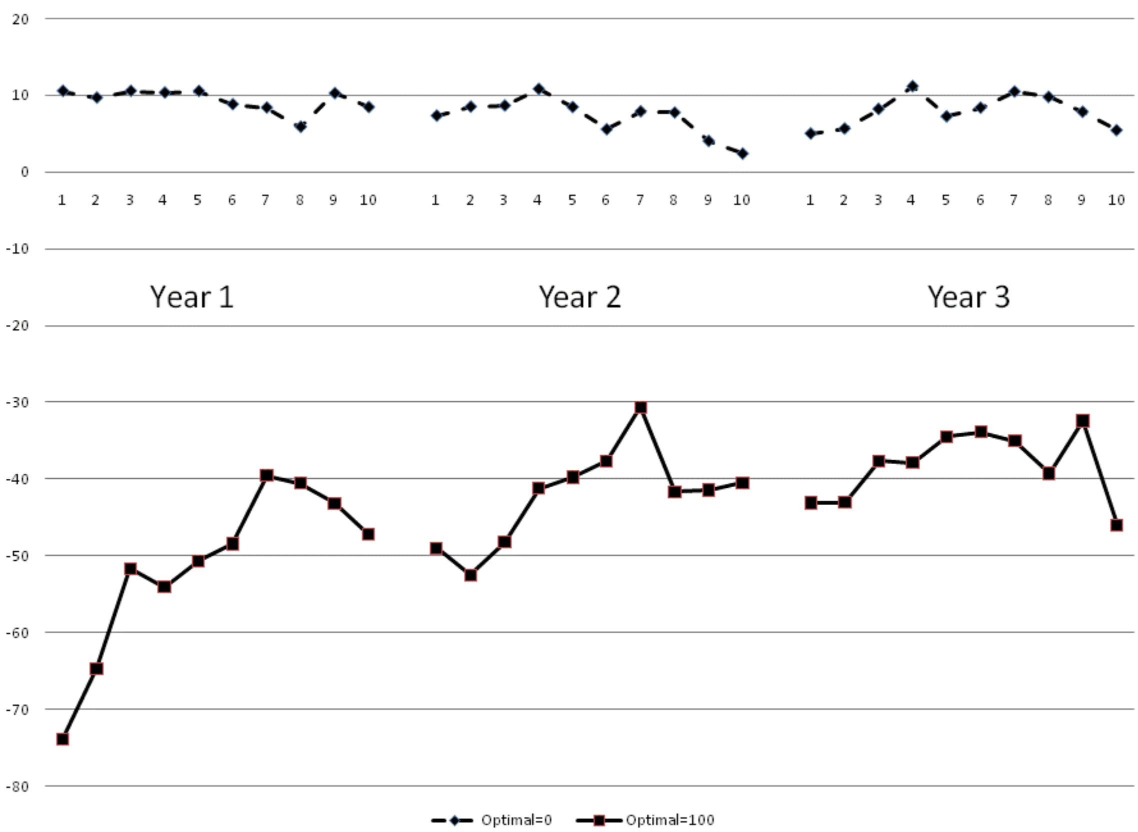

Figure 5: Actual minus optimal investments for 10 sequential hurricane encounters by year when optimal investment was either 100 (solid line) or 0 (dashed line), Study 1.

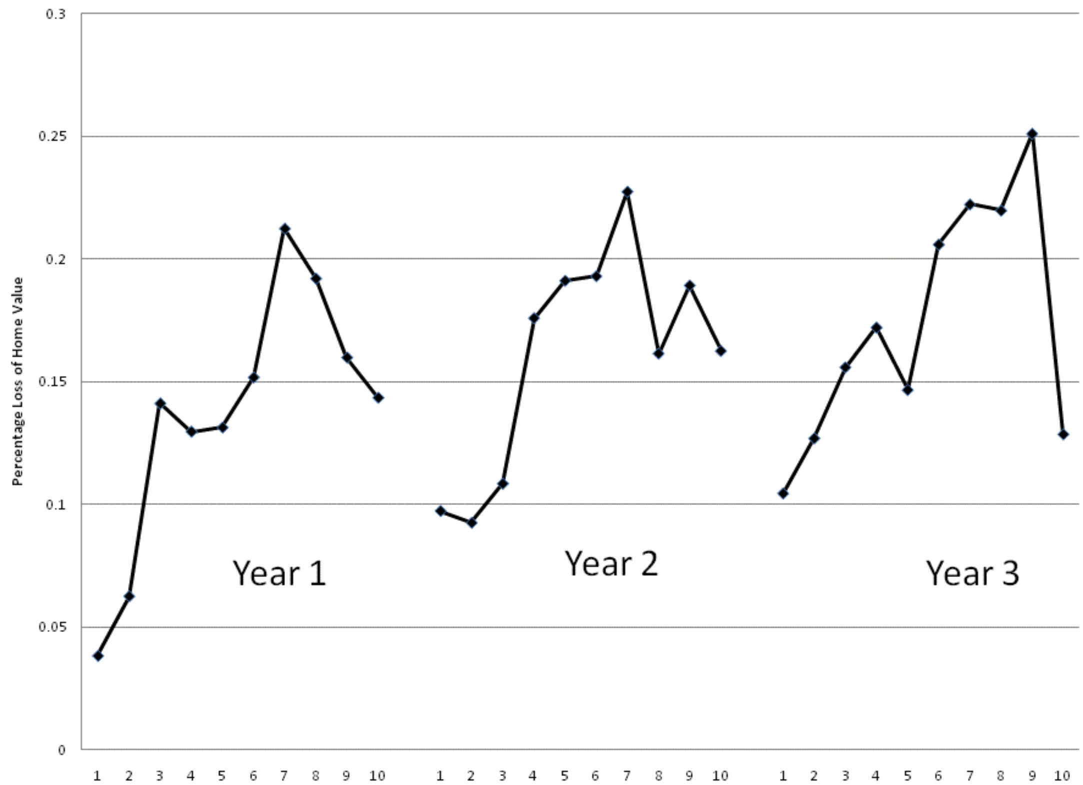

Figure 6: Mean losses that could have been avoided by following the optimal protection policy by storm and year, expressed as a percentage of total home value

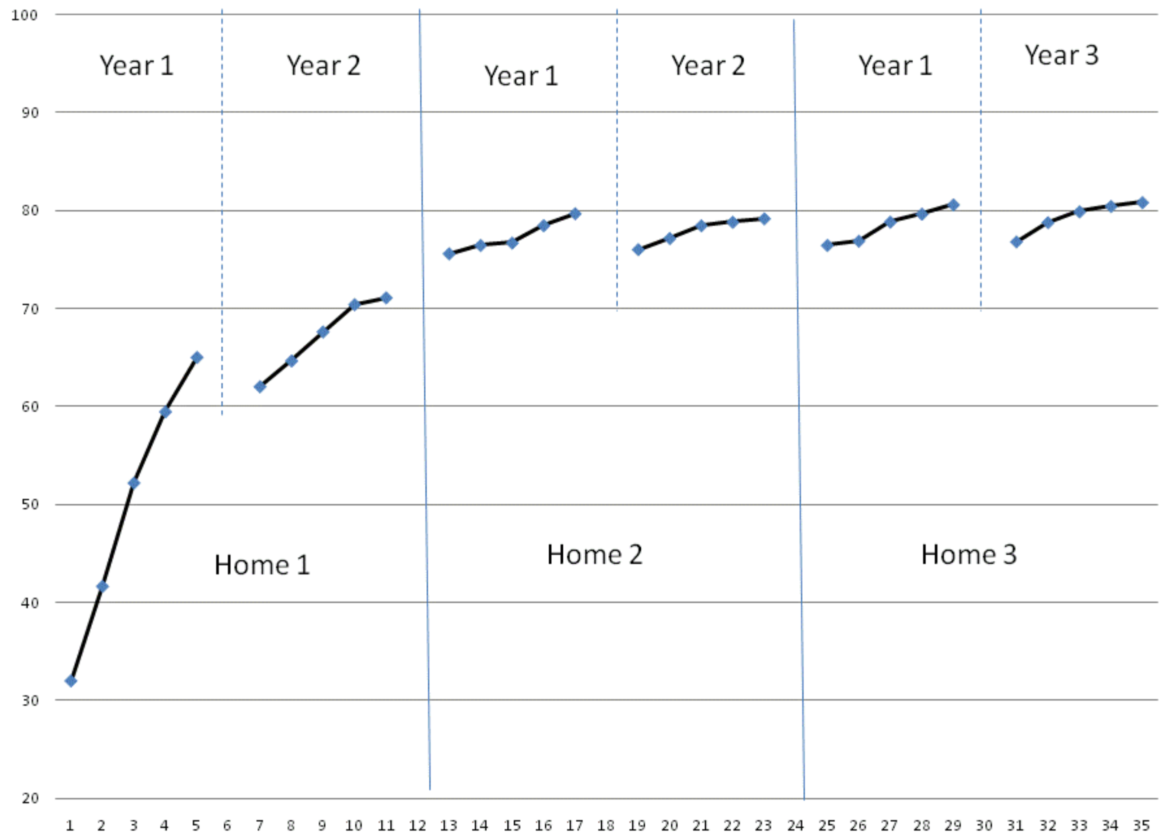

Figure 7: Mean levels of home protection over time in the long-term mitigation variation of the simulation. Optimal level of protection was 100.

Appendix: Screen Shots from the Hurricane Simulation, Study 1

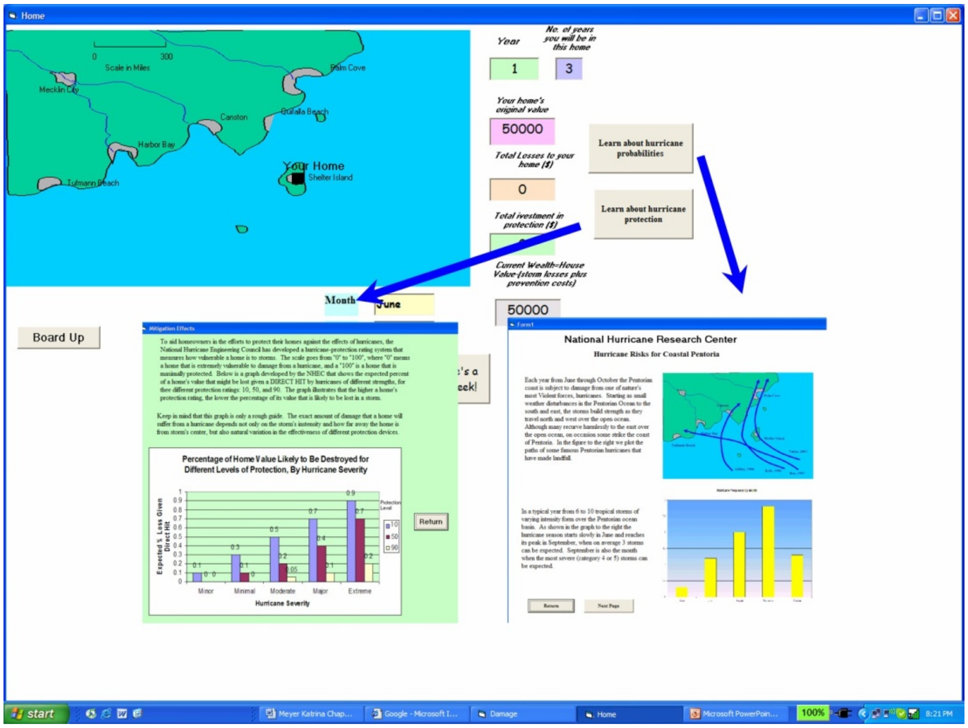

Figure A-1: Basic Interface and Information gathering

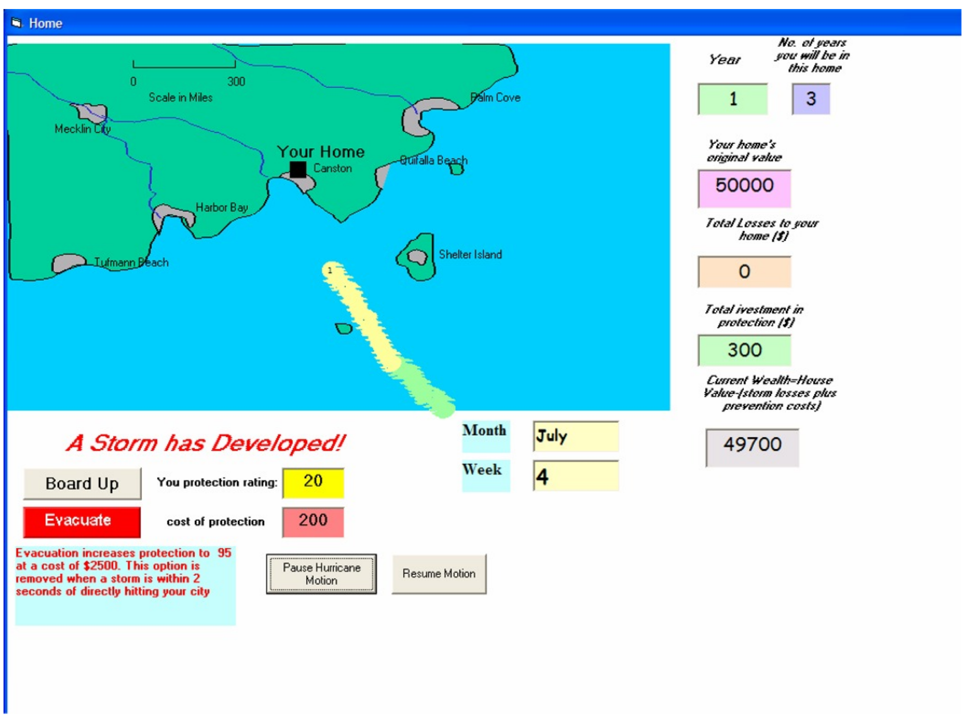

Figure A-2: Storm motion and mitigation decisions

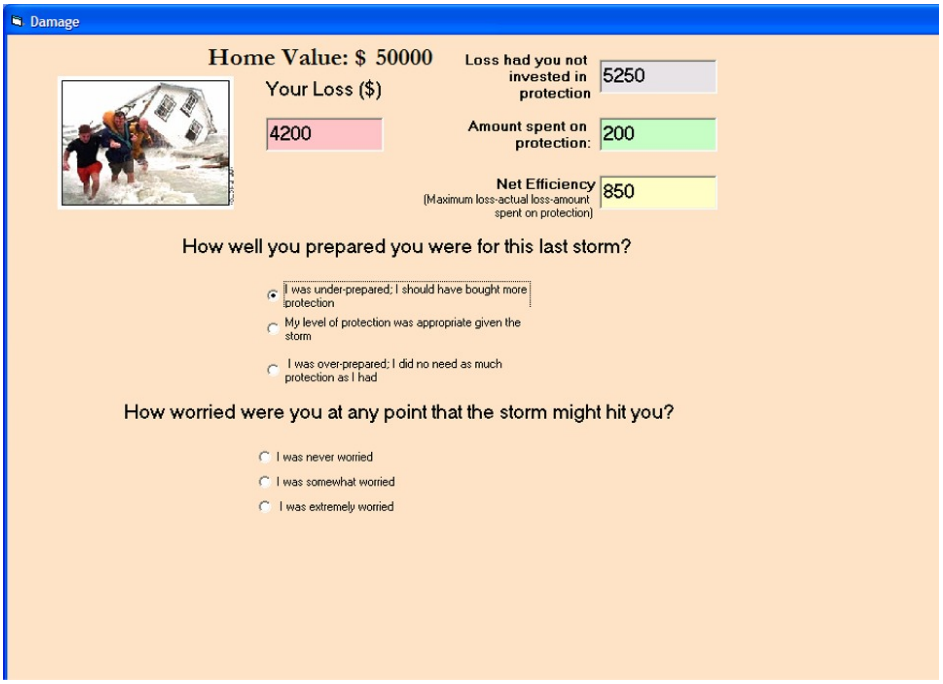

Figure A-3: Damage feedback and debriefing