2

Lessons from the Earthquake Lab: An Experimental Analysis of Learning from Experience about Natural-Hazards

Robert Meyer

WHEN IT COMES TO KNOWING how best to protect against natural hazards, experience does not always seem to be the best teacher. While enormous strides have been made over the past century in knowing how to mitigate losses from hazards, this knowledge often seems to be underused by the communities faced with such risks. For example, in the year prior to Hurricane Katrina, New Orleans underwent a full-scale dress rehearsal for the storm when Hurricane Ivan forced a large-scale evacuation of the city. During the hurricane, weaknesses in the city's preparedness plans, such as its inability to evacuate and shelter the city's immobile population, become apparent. Yet when Katrina hit a year later, New Orleans had taken few steps to remedy those weaknesses.[1] Similarly, despite the large losses suffered by U.S. coastal communities after the hurricane seasons of 2004 and 2005, the vast majority of coastal residents made no investments to improve the wind-resistance of their homes, or even to acquire hurricane-preparedness packages.[2]

But are such reports really evidence of slow learning, or are they simply examples of the difficulties inherent in making good decisions with limited information? While Hurricane Ivan may well have served to draw attention to the potential risks that New Orleans faced if hit by a major hurricane, the storm *did* bypass the city, which may have understandably built up the city's confidence in the low likelihood of a major flood disaster. Moreover, in the spring of 2005, long-run forecasts called for a comparatively quiet hurricane season, a prediction that may have encouraged planners to believe that they had ample time to seek a remedy for the evacuation problem.[3] Finally, as vivid as the losses were along the Gulf Coast from Hurricanes Katrina and Rita in 2005, many residents who were unaffected by the storms may not have been convinced that large capital investments in structural improvements would really be cost-effective for them in the long term given the low actuarial likelihood of such events affecting them in the future.

This chapter reports the results of research that focused specifically on earthquakes and was designed to understand how well individuals learn from experience about making mitigation decisions. Given the difficulty of studying this experiential process in field settings, where repeated exposure to earthquakes is rare, we conducted our investigation using laboratory experiments that simulated repeated exposure to such hazards. In these experiments, participants were provided with homes in virtual communities and, over a period of time defined by the experiment, had the opportunity to learn about effective mitigation both by directly experiencing earthquakes and by observing the decisions and experiences of others.

The findings from this research left us with a disquieting view of the limits of individual and social learning. Although participants were given ample opportunities to learn, investment patterns were marked by a long-term tendency to underinvest in mitigation when it was objectively cost-effective. This reluctance to invest in mitigation, however, was not universal; paradoxically, participants systematically *over*invested in mitigation when it was objectively cost-*in*effective and the role of mitigation was to that of a placebo.

TRIAL AND ERROR LEARNING IN THE EARTHQUAKE LAB

Much of what we know we learn by trial and error. We learn to play tennis by repeating motor movements that produce positive outcomes and avoiding those that produce negative outcomes, and we develop advanced skills at math by repeating those problem-solving techniques that most reliably produce correct answers.[4] But while trail-and-error learning heuristics may serve us well in most walks of life, one might conjecture that they will be far less useful in guiding us to policies for investing in mitigation against low-probability, high-consequence hazards. The reason is simple: the rarity with which we encounter natural hazards leaves us with few opportunities to observe the consequences of alternative strategies for mitigation: because instances in which investments are *not* needed vastly outnumber those in which they are, successful mitigation strategies, by definition, can be deemed effective only by the absence of negative consequences.

Overview

To study how the extent to which sparse experience with natural hazards might impede learning about mitigation, we used a simulation that allowed us to observe laboratory decision makers in a virtual community, and determine how well they could learn from experience the best way to invest in mitigation against earthquakes. The simulation was designed to provide experimental control over the actuarial risk of earthquakes while mimicking the informational environment that characterizes real-world investment decisions. Specifically, during the course of the simulation, participants had the opportunity to read simulated news articles that provided information about the nature of the earthquake risks they faced, purchase protection against earthquakes, observe the investment decisions made by other residents in the community, and witness the damaging effects of quakes.

The immediate goal of the simulation was to see whether participants, given enough experience, would discover a level of investment that offered the optimal balance between mitigation costs and benefits—a balance that was controlled by the experimenter. The experiment examined two optimal scenarios: one in which investments were highly cost-effective, leading to a high optimum investment level, and one in which investments were cost-ineffective, leading to a low optimum investment level.

Participants and Incentives

Eighty-seven undergraduate and graduate students volunteered to participate in response to a cash incentive. Experiments were run on a small-group basis in a behavioral research lab, with each participant seated in a partitioned cubicle and equipped with a personal computer. In return for their participation, each subject received a $10 show-up fee and was told that the participant who earned the highest score in the game (defined below) would be given a $200 cash reward.

The Simulation Procedure

Seated at their computers, participants were asked to imagine that they had just moved into a home in a hypothetical country that was prone to periodic earthquakes, where they would be living for five years. The central interface, reproduced in Figure 2–1, consisted of a map of the hypothetical country that showed the location of the participant's home as well as the location of other residents' homes, and updated information about the participant's current wealth, total losses, and current level of mitigation.

As time progressed in the simulation, buttons would appear on the interface enabling participants to navigate through four phases of decision making:

1. *Information search:* By clicking on a button that said "Learn about earthquakes and their dangers," participants could access a series of research reports that provided detailed information about the frequency with which they might expect to encounter earthquakes of varying severity, and about the damage that a quake of a given magnitude could cause to an unmitigated home conditional on the earthquake's location and strength. At the start of the simulation, all subjects were required to certify that they had read these reports before they were permitted to undertake mitigation decisions. Whether and how they used this information, however, was discretionary.
2. *Investments in mitigation:* By clicking on a button that said "Buy protection," participants entered an "earthquake protection store," where they could buy up to 100 "mitigation units" at a cost of $100 per unit. Participants could purchase mitigation units at any time during the five years of residence. The level of mitigation afforded by the purchases was cumulative, and once an investment was made, it could not be revoked within the five-year ownership cycle. Participants could also click on a button that said "See others' protection"; this produced a map that showed the most recent levels of mitigation purchased by other players.
3. *Earthquake determination:* After participants decided whether to access information and/or buy mitigation, they clicked on a "Ready" button. This brought up a series of six green or red buttons that indicated whether other players had finished making their decisions (Figure 2–1). When all the buttons turned green, subjects either viewed a message that said "no quake this year" or, if there was a quake, learned of its location and magnitude. A quake was manifested on the screen by an animation that showed a set of concentric circles emanating from its epicenter and a text message indicating its strength.
4. *Damage resolution.* If there was a quake, a button appeared on the screen labeled "View damage reports," which took participants to a new screen that showed the level of damage sustained by each player's home. On that same screen players were also able to view the levels of protection undertaken by the other "residents," so they could easily toggle between information about levels of protection and levels of sustained damage.

In the simulation, performance was measured by a "wealth score," which was defined as the participant's initial home value—in this case, $40,000—minus investments in protection and damage losses from quakes. If the home suffered quake damage, it was assumed to be rebuilt with the previous level of mitigation (if any) restored. Note that if the total amount spent on mitigation and repairs exceeded the initial home value, the participant would have a negative wealth score for that particular cycle of the simulation. There was no extraordinary penalty for having a negative wealth total (such as interest payments); it simply lowered the overall score a participant would realize over all rounds of the simulation. Participants could continue to buy protection and make repairs even when their wealth score dropped below zero. To simplify the task, participants faced no liquidity constraints, and funds not spent on mitigation or repair could not be externally invested.

Participants played eight independent replications of this five-year ownership game. After each five-year cycle, the financial slate would be wiped clean and another five-year cycle would begin. Each participant's overall score in the simulation was his or her cumulative net assets after the eight replications (forty total decision periods).

Figure 2–1. Main simulation interface.

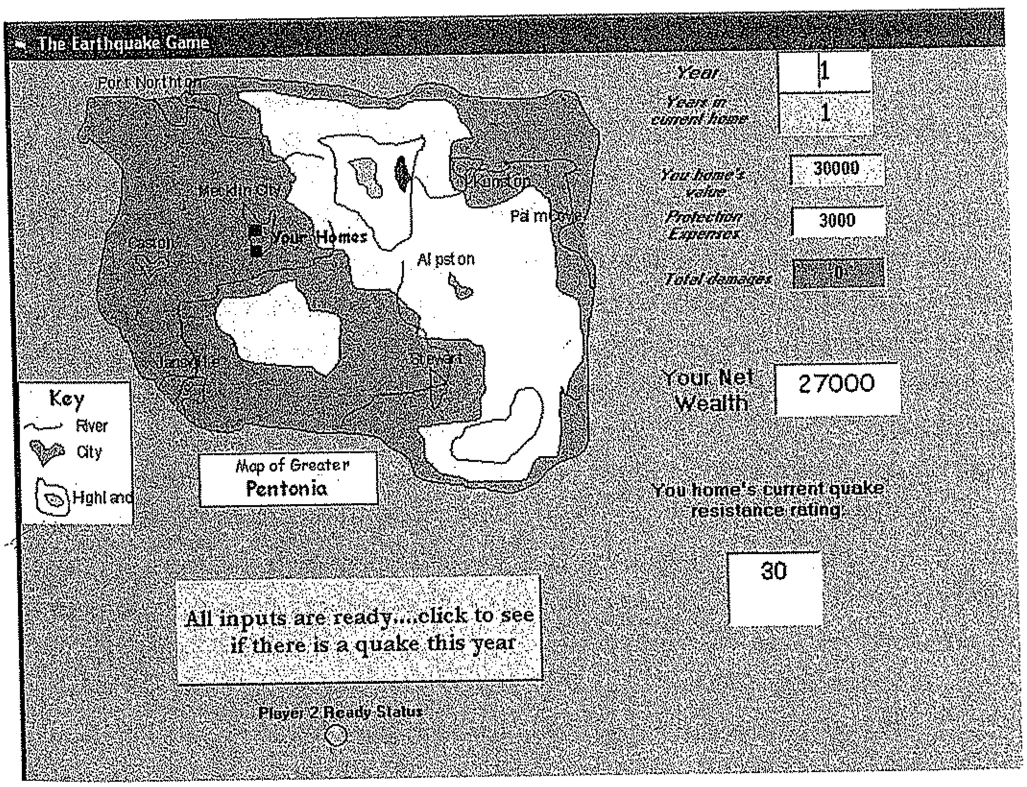

Simulation Parameters

The research reports provided to subjects at the start of the simulation informed them that in any given year, there was a 50 percent chance that an earthquake of some magnitude would occur somewhere in the country, with its epicenter being randomly determined. In addition, participants were informed that there were four possible levels of earthquake intensity, ranging from "minor" to "extreme," and that the probability that an observed quake would be of each of these magnitudes was .5, .30, .20, and .10, respectively. These probabilities were conveyed by way of a histogram that plotted the relative likelihood of an earthquake being of each magnitude.

The amount monetary damage *(L)* suffered by a home *(i)* if the quake's epicenter was at a specific location *(j)* was given by the formula

$L_{ij} = e^{-\alpha d_{ij}} SV(1\text{-}I)(1\text{-}m)$,

where e = base of the natural logarithm,

a = a scaling parameter,

d_{ij} = the Euclidean distance between the player's home *(i)* and the quake's epicenter *(j)*,

S = as the scalar measure of the quake's strength,

V = the starting monetary value of the player's home,

I = the percentage of possible mitigation units purchased by the player, and

m = a continuous scalar parameter bounded by the range [0,1] that captured the marginal effectiveness of improvements (described below).

Subjects were not given this formula, but its meaning was conveyed through a histogram that plotted the percentage of a home's value that would be lost conditional on the home's strength at two distances: (1) at the quake's epicenter and (2) at a maximum distance from the epicenter as shown on the map.

True-Effectiveness Manipulation

Participants were told that there is considerable disagreement among mitigation experts about whether investment in mitigation is worthwhile, with half claiming that it is highly effective and half claiming that it is ineffective. Subjects were told that the true value was something they would need to discover on their own through experience, and that there was a single "true" value of mitigation that applied to all residents in their community of players. This true value of m was determined at the start of a given set of eight replications, with mitigation being effective for half of all communities (where $m = .8$) and ineffective for the other half (where $m = 0$). The damage function was scaled so that subjects in the high-effectiveness conditions should have invested the maximum in protection (100 units) while those in the low-effectiveness conditions should have invested zero.

Social Feedback Manipulation

While participants were led to believe that the simulations were networked, they were, in fact, playing independently, with the information they received about the actions of other players being controlled by programmed agents. To make this manipulation convincing, participants were made to wait for varying amounts of time until the other "players" finished making their decisions. Waiting times were a function of the time that the participant took to make his or her choice (those making the decision very quickly had to wait longer), whether other players suffered damage in the last round (if there was no quake, the elapsed time was short), and the stage in the simulation (mean waiting times decreased as the game progressed).

In debriefings after the simulation, none of the participants indicated a suspicion that the decisions they saw being made by other players might have been computer generated. One-third of the subjects were assigned to a "solo play" control group, where no information was ever provided about the decisions being made by others. The others were given feedback that reflected one of two programming rules:

1. *Mirrored play:* Simulated players observed the number of mitigation units purchased by the participant in the current period, and purchased that same number of units (plus or minus a small random amount) in the subsequent period.
2. *Positive leaders:*[5] Simulated players gradually invested higher amounts in mitigation over the eight game replications regardless of its true effectiveness. The updating rule was $I_{it} = aD_{it-1} + \epsilon$, where D_{it-1} was the damage recorded by the simulated player if there was an earthquake in the previous period, t was time, and ϵ was a uniformly distributed random error. The slope parameter α was chosen to ensure that the mean investments by other players was 100 units (the maximum) by the eighth round of the simulation.

It should be noted that under the second process, the participant would observe other players investing increasingly larger amounts in mitigation over time, regardless of whether such investments were objectively cost-effective.

Results

How efficient were participants in learning the optimal levels of mitigation? In Figure 2–2 we plotted the average investment levels over time, pooled over social feedback conditions, when mitigation was truly effective (broken line) and truly ineffective (solid line). If participants were making decisions optimally, we would expect to see mitigation levels over time converging to one of two numeric extremes: 100 if the investments were objectively cost-effective or 0 if they were objectively cost-ineffective. While the data reveal some initial awareness of the true effectiveness of mitigation among participants—investments were 25 percent higher when mitigation was effective during the first two cycles of home ownership compared with when it was ineffective—the more salient feature of the data was the lack of subsequent movement toward the (100, 0) optimal levels. Indeed, if anything, the data show a tendency for the difference in average mitigation levels to diminish rather than grow over time, thereby dispelling the notion that participants were learning whether the earthquake protection was effective. Thus, the data most closely resemble the most pessimistic view of learning ability provided by the simulation: that plotted in Figure 2–2, showing that naïve agents were slow to update the beliefs that they formed on the basis of short-term feedback about apparent effectiveness.

To provide insights into the process that produced this behavior, in Figure 2–3 we superimposed plots of two features of the participants' mitigation decisions over time: (1) the relative frequency with which the participants purchased mitigation (of

Figure 2–2. Average observed investment levels over time by true mitigation effectiveness.

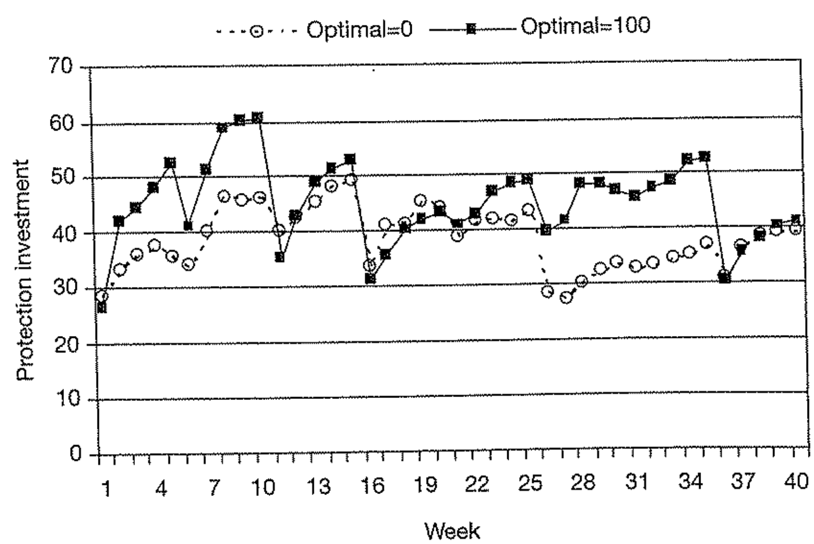

Figure 2–3. Changes in purchased mitigation at the start of each five-year cycle (solid dots), and probability of an additional purchase being made in each subsequent period conditional on a lagged experienced loss (lines).

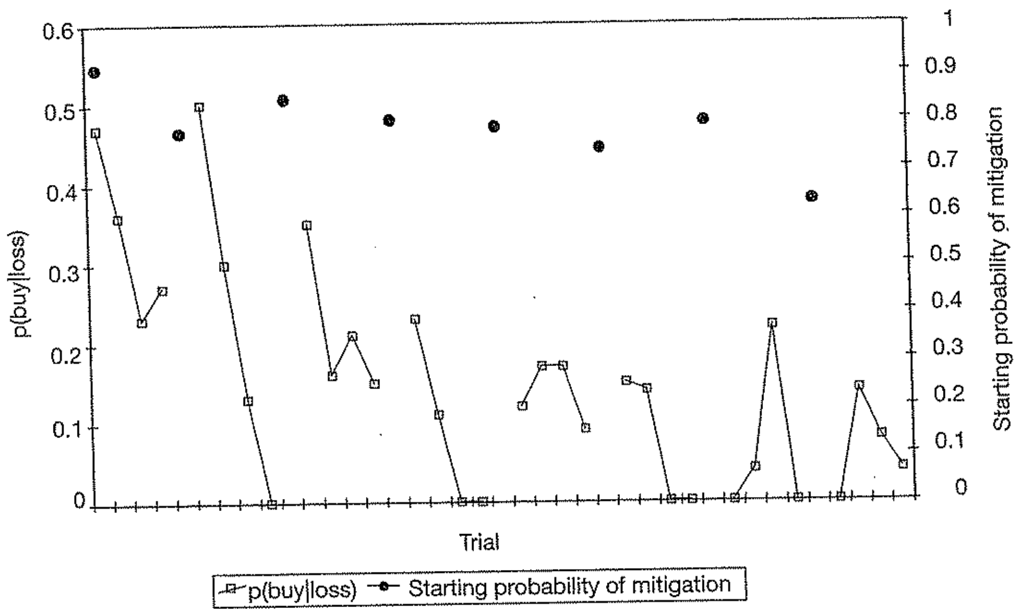

any quantity) in the first year of each five-year cycle of ownership and (2) the lagged effect of experienced earthquake damage on the purchase of additional mitigation in each subsequent year. Figure 2–3 suggests that during the first several cycles of home ownership, investments evolved over time through the following stylized anchor-and-adjustment strategy:

> Start each five-year cycle of home ownership by purchasing a limited buffer stock of mitigation. If earthquake losses are then experienced in that or subsequent periods, react by buying additional units of protection. As the horizon of ownership decreases, reduce the amount of additional units purchased.

To illustrate, in each of the first four cycles of ownership, 85 percent of participants purchased at least one unit of mitigation when first given the opportunity, but the amount was limited (a mean of 34 units of a possible 100) and not significantly related to true effectiveness (33 when it was ineffective, 35 when it was effective; $F[1,402] = .77$; $p = .38$). After that starting purchase, subsequent purchases were made in decreasing responsiveness to experienced earthquake losses (Figure 2–3), which might be interpreted as a rational response to the decreasing time horizon of hazard. For example, over the first two cycles, an average of 50 percent of participants bought more protec-

Figure 2–4a-c. Average observed investment levels by social feedback condition and true mitigation effectiveness over cycles.

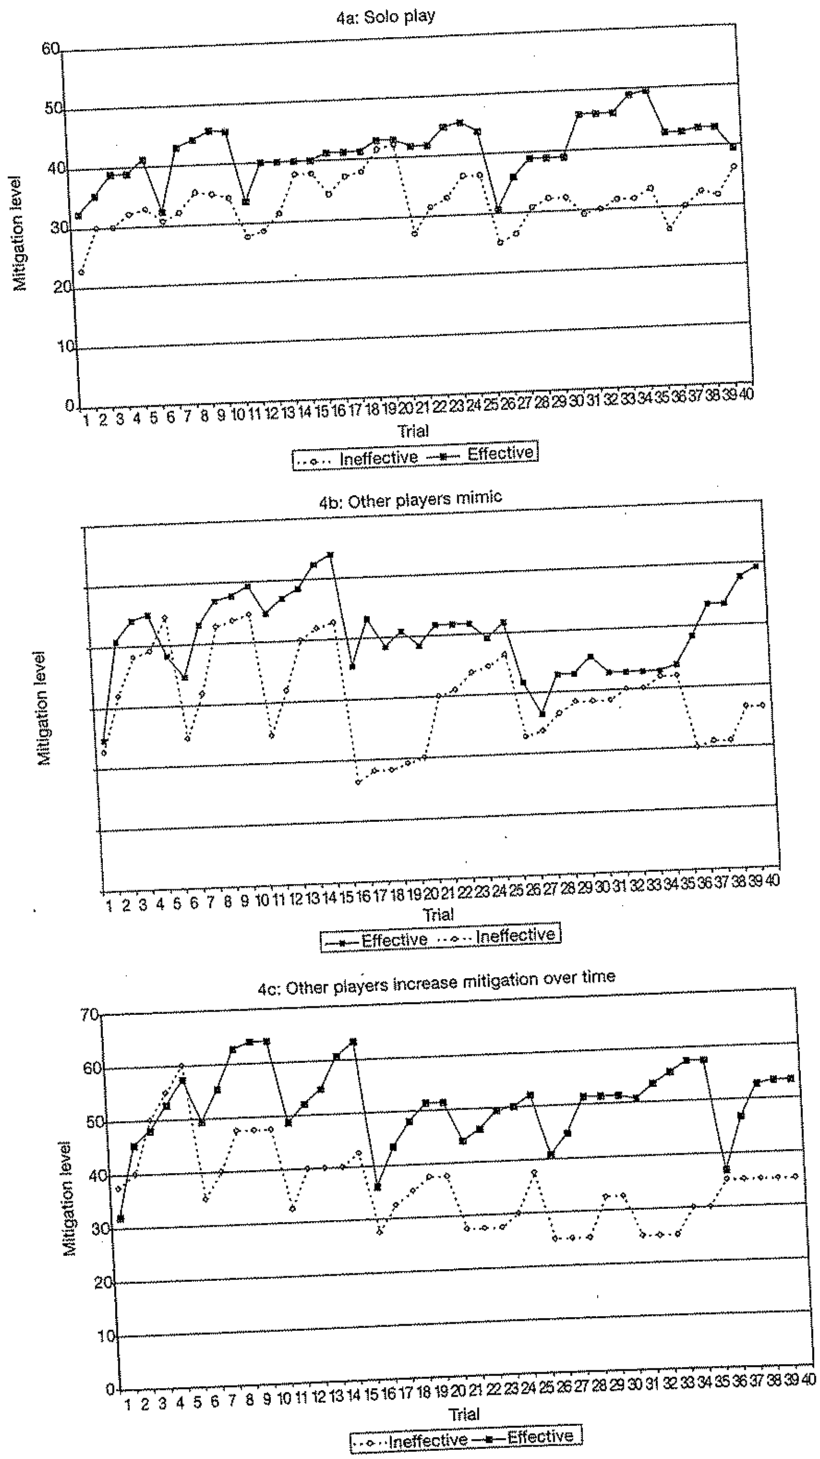

tion after experiencing a quake loss in the first period, but this decreased to an average of 13 percent when a quake loss was experienced in the penultimate period.

After four such cycles of investment dynamics, participants appeared to settle into a more stable process characterized by a fixed initial investment that was less likely to be revised despite subsequent earthquake losses. But this equilibrium adaptation was far from optimal: initial investments were far too high where mitigation was ineffective and far too low where it was effective. In addition, even after forty trials and eight cycles of home ownership, most subjects were still revising their initial mitigation levels (although less systematically) in response to short-term experienced losses; for example, in the last (eighth) cycle of home ownership, only 40 percent of participants made no changes in their investment levels for their five-year course of tenure. Thus, few participants seemed to grasp the normative insurance concept that the best strategy for mitigation is to make all purchases in the first period when they have their greatest time value—not later, *after* damage has occurred.

To explore how decisions about mitigation were affected by the ability to observe the decisions made by others, we plotted mean investment levels over time by each social feedback condition: where participants received no social feedback (Figure 2–4a), where simulated other players mimicked the investment decisions made by the participant (Figure 2–4b), and where simulated other players gradually concluded that mitigation was effective (Figure 2–4c). The figures suggest that the ability to observe the decisions made by others did little to either aid or inhibit learning. In all three cases, subjects invested more in mitigation when it was truly effective and less when it was not (as they should have), but average investments remained far from optimal, marked by underinvestment when it was optimal to invest and overinvestment when it was not.

Perhaps the clearest view of the limited role that social feedback played in influencing investment decisions can be seen in the case where participants observed other decision makers gradually investing more in mitigation over time, reaching 100 percent by the eighth home cycle—a behavior that was perfectly optimal for half of all participants and perfectly suboptimal for the other half. In this case there is some evidence that participants' decisions were affected by those made by others, but it is limited. Specifically, Figure 2–4c shows a pattern in which subjects who were in the ineffective mitigation condition (the broken line in the figure) rapidly increased their investments over the first few trials, an action that would have mimicked that being displayed by other decision makers. But over time, these participants learned to disregard this "advice"; by the fourth cycle of home ownership, their mitigation levels had fallen to that observed among participants in solo play and those whose decisions were being mimicked (Figures 2–4a and 2–4b). In contrast, there is no evidence of reciprocal learning among participants in effective mitigation conditions when the implied advice was optimal: while equilibrium investment levels in those cases were higher than those in either solo or mirrored play, they were still well below both the levels implied by the optimal policy and the levels they observed among others.

Why were participants not more influenced by peer actions? Insight into this question can be found in Figure 2–5, which plots the relative frequency with which participants looked at the mitigation levels of other players when making their own investment decisions. The figure provides a simple explanation: participants began the simulation actively looking at the mitigation investment decisions of others (over the first five cycles, an average of 70 percent of participants engaged in this behavior), but this interest rapidly diminished, averaging 35 percent over the last three cycles. Hence, subjects either arrived at a decision policy that they felt did not need revising after the first two or three rounds of play, or, after first being curious about the actions of others, quickly concluded that the information was of little value in making their own decisions. The reality, of course, was far from that: by comparing other players' mitigation levels with their own after they had sustained earthquake damage, participants might have been able, in theory, to quickly and conclusively discover whether mitigation was effective.

Figure 2–5. Mean relative frequency with which participants looked at the mitigation levels of other players by social feedback condition and simulation trial (1–40), aggregated in blocks of five trials.

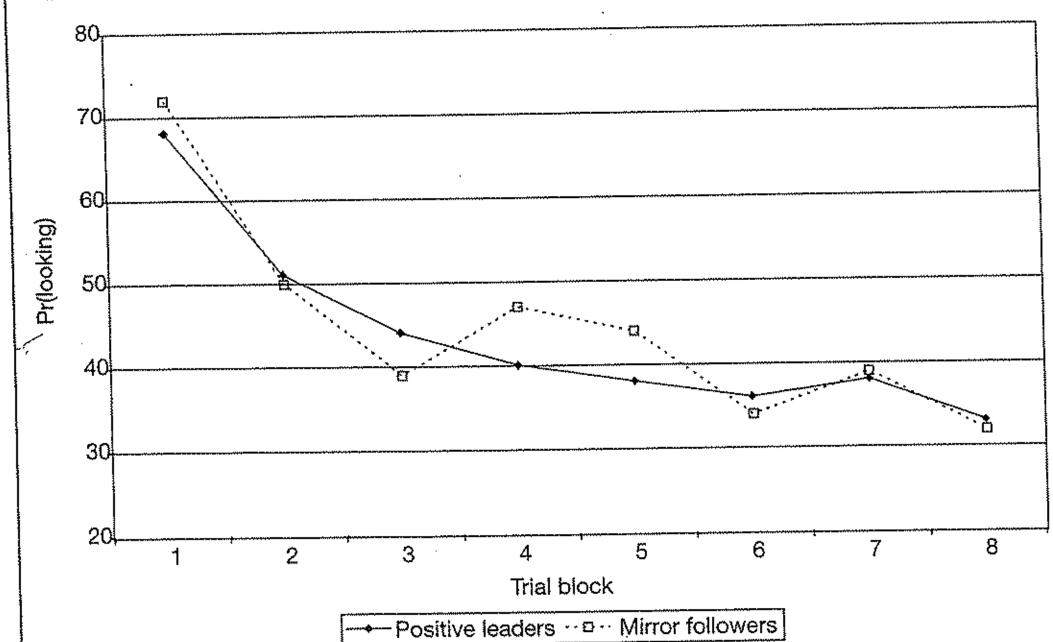

Finally, to provide a more rigorous statistical analysis of investments over time, we modeled observed mitigation levels over time as a linear function of the following seven sets of predictors:

1. Social feedback condition (two contrasts: positive leaders versus solo play, and mimicked decisions versus solo play)
2. Lagged level of mitigation
3. Year of tenure in the home
4. Home number (cycle of play)
5. True mitigation effectiveness
6. The six interactions between social feedback condition and predictors 3 through 5 above
7. The two interactions—between true effectiveness and lagged protection, and between true effectiveness and home number.

The results of this analysis, reported in Table 2–1, lend statistical support to the qualitative observations about learning offered above. Specifically, the analysis supports a modest positive main effect of true effectiveness, but supports no interaction between this effect and the decision-making experience (as measured by five-year home cycle)—thus statistically confirming the lack of convergence of investment levels toward the optima across conditions noted earlier. In addition, the data fail to support a significant "positive leader"-by-true-effectiveness interaction ($p = .129$), reinforcing our earlier finding that observing other players increase their investments over time did little to enhance (or deflate) learning of true effectiveness compared with such learning in solo play. Finally, the analysis reveals a significant negative interaction between mirrored play and true effectiveness ($p < .001$)—a result that implies that the difference in mitigation levels between the two true-effectiveness conditions was less pronounced (consistent with greater confusion) when other players were mimicking the actions of the decision maker.

DISCUSSION

Much of what is known about how to protect against natural hazards has been acquired through a costly process of trial and error. The 2004 Asian tsunami tragedy provides a compelling case in point: as tragic as the loss of life was, it prompted governments around the Indian Ocean to see the value in establishing a regional tsunami warning system, a preventive measure long in place around the Pacific Rim to the east. On the other hand, the fact that such a system was *not* in place in 2004 underscores how ineffective learning can sometimes be. Although scientists had been making repeated calls for tsunami warning systems to be established outside the Pacific in the years preceding 2004, such calls had gone unheeded, presumably because of a lack of recent direct experience with such events.

Are there inherent limits to our ability to learn about the effectiveness of mitigation measures from past experience? Our research examined this question by reporting how experimental participants made repeated decisions about whether to invest in mitigation in a dynamic earthquake simulation. In the simulation there was an optimal policy for mitigation that was unknown at the start but that could be partially discovered over time, either by direct experience or by observing the experiences of other players.

On average participants grossly underinvested in mitigation when it was truly effective and overinvested when it was ineffective. There was little evidence of investments converging toward optimal levels over time regardless of whether one was able to observe, or was aided by being able to see the consequences of, mitigation decisions being made by others. The failure to converge to optima is consistent with previous research showing that human decision makers are poor at learning from feedback in complex noisy systems.[6] Although the mechanism that drove the damage from earthquakes was deterministic, the complexity of the function would have made it difficult for participants

Table 2–1. Ordinary Least Squares Regression of Mitigation Levels over Time

Parameter	Estimate	Standard error	T	p(T)
Intercept	25.33581	2.854601	8.88	<.0001
True effectiveness (TE)	3.176299	1.60772	1.98	0.0483
Replicate	-0.0201	0.463314	-0.04	0.9654
Year in home	-5.80111	0.314845	-18.43	<.0001
Lag protection	0.896883	0.034832	25.75	<.0001
Positive leaders (PL)	1.222339	3.716493	0.33	0.7423
Mirrors player (MP)	7.037851	2.963405	2.37	0.0176
TE*Replicate	0.01902	0.267601	0.07	0.9433
TE*Lag Protection	-0.027	0.020446	-1.32	0.1867
PL*TE	2.574409	1.69781	1.52	0.1295
PL*Replicate	-0.41285	0.32422	-1.27	0.203
PL*Year	-0.46462	0.520328	-0.89	0.372
MP*TE	-5.44446	1.392549	-3.91	<.0001
MP*Replicate	-0.5511	0.30193	-1.83	0.0681
MP*Year	0.817815	0.487448	1.68	0.0935

Overall model: $F(14,3187) = 485.1$, $p < .001$; $R^2 = .68$.

to discern the extent to which causality from a given damage episode (e.g., whether damage was low because of mitigation) was effective or the quake was ineffective.

A participant who is thinking about the long term should be willing to actively experiment with mitigation investment levels; from such experimentation, the participant would learn that significant damage might still be incurred, even with a maximum investment in mitigation. Our study shows, however, that this insight about the value of experimentation seemed to elude participants—a bias similar to that found in other studies of sequential decision making.[7] There were no cases of subjects purchasing 100 percent of available protection at the start of the task to test a hypothesis about mitigation effectiveness. Rather, the modal strategy was to purchase a limited amount (e.g., 25–30 percent)—a quantity that would be insufficient to provide significant protection if mitigation was effective, or, in turn, be informative as to the extent to which it could be effective.

The data also suggest that participants focused on what they discovered about mitigation in the early rounds and simply tired of the task of learning. Decisions appeared to be characterized by a simple anchor-and-adjustment policy: participants started each round of decision making by investing in a moderate amount of mitigation, and they then bought more if they experienced a loss. But after two to three cycles of applying this policy, they seemed to abandon further attempts to update their strategy. Reactions to experienced losses diminished, as did their interest in viewing the mitigation decisions made by other players.

LIMITATIONS AND FUTURE WORK

To what degree do the limits to learning observed in the laboratory mirror what is likely to be observed in the real world? On the one hand, circumstances of learning in the experiments were far more favorable than they would be in a real-world setting. Subjects had an explicit scoring rule tied to a monetary incentive, and they had access to far greater amounts of both direct and indirect experiential information than would be available in the real world. On the other hand, they lacked many of the aids to decision making that often arise in practice, such as the ability to talk to true experts. And perhaps most important, they faced only hypothetical losses.

Several of the learning biases observed in the simulation are reminiscent of errors in mitigation decisions that have been noted in real-world settings. For example, there is considerable empirical evidence of the overriding importance of direct encounters with hazards as a basis for perceptions of risk[8] and the undertaking of mitigation measures.[9] What is perhaps most surprising about the findings reported here is that these biases will not be cured by replication, nor can they be reduced simply by putting more information into the hands of decision makers.

Finally, we see this work as highlighting the potential value of dynamic laboratory simulations as a tool for gaining a better understanding of human response to natural hazards. Our knowledge to date of how individuals and households learn to adapt to hazards has been limited simply because nature offers us few data points, and almost never a natural experiment. While laboratory experiments will never emerge as a replacement for field studies, they may serve as a useful complement by providing both a means for testing hypotheses about hazard response that may emerge from fieldwork as well as pointers for what to look for in future empirical studies. The research reported here offers a simple illustration of this potential, and we hope it will foster additional applications in the future.

Endotes

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

2 Abby Goodnough, "As Hurricane Season Looms, State Aim to Scare," *New York Times*, 31 May 2006.

3 William M. Gray and Philip J. Klotzbach, "April 1 Extended Range Forecast of Atlantic Seasonal Hurricane Activity and U.S. Landfall Strike Probability for 2005" (Fort Collins, Colo.: Department of Atmospheric Science, Colorado State University, 2005), hurricane.atmos.colostate.edu/forecasts/2004/dec2004/ (accessed February 23, 2008).

4 Robert J. Meyer and J. Wesley Hutchinson, "Bumbling Geniuses: The Power of Everyday Reasoning in Multistage Decision Making," in *Wharton on Making Decisions*, ed. Stephen J. Hoch and Howard Kunreuther (New York: John Wiley, 2001), 37–62.

5 A "negative leader" condition was not run because subject resources were limited.

6 John Sterman, "Misperceptions of Feedback in Dynamic Decision Making," *Organizational Behavior and Human Decision Processes* 43, no. 3 (1989): 301–335.

7 See, for example, Robert Meyer and Young Shi, "Sequential Choice under Ambiguity: Intuitive Solutions to the Armed-Bandit Problem," *Management Science* 41 (May 1995): 817–834.

8 Michael K. Lindell and Ronald W. Perry, "Household Adjustment to Earthquake Hazard: A Review of Research," *Environment and Behavior* 32, no. 4 (2000): 461–501.

9 Howard Kunreuther et al., "High-Stakes Decision Making: Normative, Descriptive, and Prescriptive Considerations," *Marketing Letters* 13, no. 3 (2002): 259–268.