1-16-2005

An Experimental Analysis of Learning from Experience about Natural-Hazards

Robert Meyer
University of Pennsylvania

Howard Kunreuther
University of Pennsylvania

An Experimental Analysis of Learning from Experience about Natural-Hazards

Abstract

The ability of individuals to learn optimal strategies for mitigation against infrequently-occurring natural hazards is explored. We report the results of two experiments in which participants are faced with the problem of learning the most cost-effective means of protecting against earthquake losses. The experiments utilize dynamic computer simulations in which participants are endowed with homes in virtual communities that are prone to periodic impacts by earthquakes. Participants can invest in measures that potentially mitigate losses from quakes but the effectiveness of these measures is initially uncertain. Over time participants have the opportunity to learn about true effectiveness both by direct experience with simulated earthquakes and by observing the decisions and experiences of other players. The data offer a pessimistic view of learning abilities; not only do participants persist in investing in mitigatin instruments that, in fact, have no ability to lower damage, but they also fail to fully invest in instruments that are highly effective. Among the mechanisms that appeared to impede learning was a tendency to mimic local group norms in investment levels (which are suboptimal) and to prematurely terminate attempts to learn. The paper concludes with a discussion of the implications of the work for both basic research on decision making in low-probability, high-consequence settings as well as prescriptive research in natural-hazard mitigation.

Disciplines

Behavioral Economics | Business | Emergency and Disaster Management | Marketing | Real Estate | Recreation Business | Urban Studies and Planning

Comments

This is an unpublished manuscript.

An Experimental Analysis of Learning from Experience about Natural-Hazards

Robert Meyer
Howard Kunreuther

January 16, 2005

Robert Meyer is the Gayfryd Steinber Professor of Marketing at the Wharton School of the University of Pennsylvania, Philadelphia, PA 19010, meyerr@wharton.upenn.edu. Howard Kunreuther is the Cecilia Yen Koo Professor of Operations and Information Management and Business and Public Policy at the Wharton School, kunreuth@wharton.upenn.edu. The work was supported by resources provided by the Alfred West Learning Laboratories of the Wharton School, University of Pennsylvania. the Wharton Risk Management and Decision Processes Center, and Singapore Management University

Abstract

The ability of individuals to learn optimal strategies for mitigation against infrequently-occurring natural hazards is explored. We report the results of two experiments in which participants are faced with the problem of learning the most cost-effective means of protecting against earthquake losses. The experiments utilize dynamic computer simulations in which participants are endowed with homes in virtual communities that are prone to periodic impacts by earthquakes. Participants can invest in measures that potentially mitigate losses from quakes but the effectiveness of these measures is initially uncertain. Over time participants have the opportunity to learn about true effectiveness both by direct experience with simulated earthquakes and by observing the decisions and experiences of other players. The data offer a pessimistic view of learning abilities; not only do participants persist in investing in mitigation instruments that, in fact, have no ability to lower damage, but they also fail to fully invest in instruments that are highly effective. Among the mechanisms that appeared to impede learning was a tendency to mimic local group norms in investment levels (which are suboptimal) and to prematurely terminate attempts to learn. The paper concludes with a discussion of the implications of the work for both basic research on decision making in low-probability, high-consequence settings as well as prescriptive research in natural-hazard mitigation.

How skilled are individuals and communities at learning how to protect themselves against natural hazards? On the one hand, it is clear that learning from experience can, at times, be fast and efficient. Residents in low-lying communities that endure major floods quickly learn where it is *not* safe to build homes and businesses; likewise, hurricanes can offer a vivid lesson of the merits of building of wind-resistant structures. Consistent with this, prior empirical work on how individuals make mitigation decisions has consistently shown that direct encounters with hazards not only serve to heighten individuals' awareness of the dangers they pose, but also induce active investments in protection (e.g., Lindell and Perry 2000; Peacock 2003; Russell, Goltz, and Bourque 1994).

On the other hand, one might also point to cases where learning from experience seems surprisingly slow and spurious. Despite the enormous damage caused in California by the great San Francisco earthquake in 1906, for example, it was not until the Long Beach earthquake of 1933--where there was considerable damage to schools--that California established a building code for public structures (Andrus 1952). Similarly, mitigation experts in the Midwest fight an annual struggle to debunk the widely-held belief that the best way to prevent a home from collapsing during a tornado is to open its windows (to equalize pressure)—a practice that, in fact has long been known by engineers to *enhance* the chance of structural failure (Pendergrass 1999)

But are such anecdotes really evidence of slow learning, or simply examples of the difficulties of making good decisions under limited information? Most of the damage in the 1906 San Francisco earthquake was due to fire so that there may not have been a connection made between designing better structures and reducing the likelihood of direct losses from an earthquake. Similarly, the misperception that open windows mitigate tornado damage presumably persists because unambiguous counter-factual evidence is rarely observable; if a house that leaves its windows open is severely damaged by a tornado, one cannot observe the damage that *would* have occurred had the windows been closed. What makes resolving these two views of learning efficiency especially difficult is the limited opportunity we have to test hypotheses about the dynamics of mitigation decisions in field settings due to the rarity of natural hazards.

The purpose of this paper is to examine how well individuals learn to make mitigation decisions over time in laboratory environments that simulate repeated exposures to natural hazards. In these experiments participants are endowed with homes in virtual communities that are prone to earthquakes. Over time they have the opportunity to learn about effective mitigation strategies both by direct experience as well as by observing the decisions and experiences of others.

We emerge from these studies with a disquieting view of the limits of individual and social learning. Although participants are endowed with ample opportunities to learn, investment patterns were marked by a tendency to over- invest in mitigation when it was normatively ineffective and under-invest when it was highly effective. These biases did not vanish with experience. Among the apparent causes was a tendency to make decisions by simply mimicking the investments made by other participants, and to cease attempts to actively learn early in the task.

The paper is organized as follows. The next section characterizes the dynamic mitigation problem that forms the focus of our experimental work and shows how it would be solved by optimal agents. Section 3 explores why actual behavior might depart from this benchmark based on previous research on limits to inductive learning. Section 4 examines the actual ability of individuals to learn optimal earthquake mitigation strategies in the context of two laboratory experiments. The concluding section interprets these findings and suggests future research to be undertaken on learning from experience.

1. Inductive Social Learning of Mitigation Strategies

We consider how individuals solve mitigation problems that have the following structure:

> A decision maker is endowed with a home and wealth W in a community prone to earthquakes. At each moment in time t there is a known constant probability p_q that an earthquake will occur at a random location in the region, the intensity of which is an independent random draw from a known density $f(q)$. If a quake occurs, each member of the community must pay an amount for repairs that is an increasing function of the magnitude of the quake and proximity to the epicenter. To mitigate such losses, each decision maker can use a portion of their wealth in time t to make non-recoverable investments I_t in permanent structural improvements up to some maximum amount I_{max} . This investment reduces potential losses by a fraction mI_t, where m is the marginal effectiveness of a dollar spent on mitigation. These investments—as well as earthquake damages—are commonly observed, and wealth *not* spent on mitigation cannot be externally invested. The marginal effectiveness of mitigation, m, is uncertain at the start of the sequence of decisions. Specifically, there is a known probability p_m that mitigation will have a high marginal effectiveness m^H and a probability $1\text{-}p_H$ that it will have a low marginal effectiveness m^L. Each decision maker makes a series of decisions about how much to invest in mitigation so as to maximize total wealth over a finite time horizon.

If this dynamic game against nature were to be repeated several times, is there a "right" way for each participant to learn about the proper amount to invest in mitigation? There is, but it might not be one that would be either intuitive or palatable to decision makers. To see this, first note that an implication of the assumption that mitigation investments has a constant marginal effectiveness (m) is that the optimal investment strategy will be all-or-nothing: if a homeowner concludes that it is worthwhile to invest a limited amount I to protect against earthquakes, he or she should observe that even better returns could be had by investing a successively higher amounts I^+ up to the maximum,

I_{max} . Hence, if a homeowner were to know the true marginal value of mitigation (*m*) the optimal policy would have the following form: if $E(L \mid p_q, I, m, T)$ is the expected cumulative loss over a *T*-period time horizon in a setting with an earthquake risk p_q, with investment *I* and marginal effectiveness *m*, the optimal mitigation strategy is:

$$\text{Invest: } \begin{cases} I_{\max} & if \quad E(L \mid p_{q,} I_{\max}, m, T) - E(L \mid p_q, 0, T) < \mathrm{I}_{\max} \\ 0 & otherwise \end{cases} . \qquad (1)$$

Note that under such an analysis the wisdom of investing will be strictly decreasing in *T*; hence, if it is deemed *not* to be worthwhile to invest in mitigation in the first year of tenure, it cannot be worthwhile in any later year.

What complicates this problem, however, is that expression (1) will *not* be the optimal investment policy for a decision maker who is uncertain about *m* and has the opportunity to repeatedly revisit the decision. In such cases the task of deciding how much to spend in each period becomes a problem in stochastic-dynamic programming, where each decision maker would be presumed to be able to look to the future over several rounds of homeownership (and rebuilding) and make the investment decision that maximizes long-term expected wealth assuming that what is learned about the effectiveness of mitigation (*m*) in each period *t* is used to make more informed decisions in period $t+1$, $t = 1 \ldots T-1$ (Meyer and Hutchinson 2001). For example, a rational multi-round investment strategy might involve undertaking a high initial experimental investment level (I_1*), and then observing how experienced damages compare to those that would have been predicted under the hypothesis that the true effectiveness is either m^H or m^L given the quake's magnitude and proximity. This discovery would then be used to guide investment decisions in subsequent trials and games (Meyer and Shi 1995). For low probability events, such as earthquakes, it may be difficult to learn about the effectiveness of mitigation given that few quakes will occur in ones lifetime.

The nature and effectiveness of heuristic solutions

It is unlikely, of course, that real decision makers would make mitigation decisions over time in a way that mirrors the above normative process. Even if one had a good working knowledge of the principles of optimal dynamic decision making, the enormous complexity of the risk function in this case would likely thwart attempts at analytic solution. Specifically, note that a household's cumulative expected losses will be a function not just of the amount invested in mitigation in each period t (I_t), its true effectiveness (the uncertain parameter m), and the incidence rate of quakes (p_q), but also the joint density of severity and distance—all integrated over the expected horizon of home ownership.

How accurate will homeowners be in making intuitive decisions about mitigation? The literature is somewhat ambiguous in its guidance. On the one hand, laboratory research on intuitive decision making under uncertainty has repeatedly shown human decision makers to perform quite poorly compared to optimal benchmarks, encumbered by biases such as a tendency for people to be myopic in their thinking (Kunreuther, Onculer, and Slovic 1998), underutilize information about probabilities (Kunreuther, *et al*,2002) and be sensitive to normatively-irrelevant frames of reference (such as a bias toward choosing status-quo actions; e.g., Samuelson and Zeckhauser, 1988). Likewise, there is little empirical evidence suggesting that individuals know how to optimally update beliefs about mitigation effectiveness in light of accumulated observations (e.g., Grether, 1980), or strategically choose investments so as to maximize future information value (e.g., Meyer and Shi 1995).

On the other hand, there is also evidence that shows that when individuals are allowed to learn by trial-and-error they can often display behaviors that closely resemble those prescribed by complex optimal models, even though decisions makers have little insight into the mathematics that underlie the optima (e.g., Meyer and Hutchinson 2001). Specifically, all that may be required is that people operate in an environment where optimal policies yield outcomes that are observably better than suboptimal ones, and where there are opportunities to repeatedly see these outcomes (e.g., Fudenberg and Levine 2000; Kalai and Lehrer 1993; Meyer and Hutchinson 2001).

Would naïve learning lead decision makers to optimal mitigation strategies in out task? To investigate this, consider a version of our task in which investments in mitigation (I_t) are scaled over the unit interval such that $1- I_t$ implies the percentage of damage from a quake that could *potentially* be avoided if mitigation was perfectly effective given the investment I_t. Specifically, if a homeowner decides to invest I_t in protection and experiences a quake that would normally render the damage d^*_t given no mitigation, he or she would experience the fractional damage $d_t = d_t^* (1-mI_t) + \varepsilon$, where m is a (0,1)-bounded increasing measure of the *true* marginal effectiveness of mitigation (initially unknown by the decision maker), and ε is an independently-distributed random variable with mean 0 and variance Var(ε).

While a large number of rules might characterize how a decision maker might try to learn in such a setting, one process that is often posited in the literature is simple reinforcement learning (e.g., Erev and Roth 1998). Specifically, applied here the a homeowner would be assumed to try to learn the optimal level of I_t to purchase over time by using the following trial-and-error rule:

1. Choose an initial investment level I_1 at random from the interval (0, 1), and make no changes in investment until an earthquake is observed.
2. Given a quake, compare the damage, d_t, to a guess about the amount that *would* have occurred had there been zero investment in mitigation. This guess, $d_t^{'}$, is an independent random draw from a distribution centered about the true unmitigated damage d_t*; i.e., $d_t^{'} = d_t^{*} + \delta$, where δ is an independently-distributed random variable with mean 0 and variance $Var(\delta)$.
3. If the observed damage is less than $d_t^{'}$, conclude that mitigation is effective, and increase investment by an amount z. If not, decrease it by z.

It is easy to show that such a naïve learning process could lead a decision maker a optimal mitigation policy, but only under two limited conditions: mitigation must be truly effective (b must be greater than 0) and there cannot be excessive noise in either the conjectures about unmitigated damage (the error δ) or the process that determines damage from quakes (the error ε). To see this, note that the decision maker will choose to invest more in mitigation after an earthquake (decide $z>0$) if the observed damage (d_t) is less than that which they imagine would have occurred had they not invested (d_t'). Such conclusions, in turn, would be increasingly likely to occur as:

1) The decision maker's real level of protection against quakes increases (given by the quantity mI_t); and

2) The odds that they erroneously observe $d_t < d'_t$ decreases; that is, the more accurate the conjectures are about unmitigated damage.

This idea can be stated more formally as follows. Let $Pr(z_t>0)$ be the probability that the decision maker will choose to invest in mitigation at time *t*. By the assumption that conjectures about unmitigated damage are unbiased estimates of the true unmitigated damage, this likelihood can be expressed as the following function of mitigation's true effectiveness and the joint distribution of the errors ε and δ:

$$\begin{aligned}\Pr(z_t > 0) &= \Pr(d_t < d_t^{'}) \\ &= \Pr[(d_t^{*}(1 - mI_t) + \varepsilon) < (d_t^{*} + \delta)] \\ &= \Pr[(\varepsilon - \delta) < mI_t d_t^{*}\end{aligned} \qquad (1)$$

where mI_t is the proportional level of protection actually possessed by the decision maker and d_t^{*} is the true level of damage that would be observed given no mitigation. Because, by definition, $E(\varepsilon\text{-}\delta)=0$, expression (1) implies that the likelihood of investing will be monotonically increasing in the product $mI_t d_t^{*}$ regardless of the assumptions one makes about the distribution of errors*;* that is, on average a decision maker will be more likely to (correctly) conclude their investments have paid off when size of the quake is large (d_t^{*} is large), they have invested substantially in mitigation (I_t is large) and these investments are truly effective ($m \to 1$).

But note that under this process, the *speed* of learning may not be particularly fast nor uniquely converge to the maximum (optimal) investment. One impedance will be the variance in the joint distribution of errors; the greater the variance, the better the odds on a given trial that the decision maker will mistakenly conclude that $d_t \geq d'_t$ (mitigation is ineffective) and hence reduce or discontinue her investments. In the extreme, as $Var(\varepsilon,\delta) \to \infty$ equilibrium investments will converge to 50%--that which would arise under a random guessing policy. Likewise, it should be clear that decision makers could also never learn to *stop* investing in mitigation when it is, in fact, truly *in*effective (m=0). In that case investments would act as a random guessing policy, since on each trial the decision maker would be just as likely to falsely conclude that mitigation was effective when it was not (regardless of the amount that had been invested) as correctly conclude that it was ineffective (i.e., in this case $\Pr(z>0)=\Pr(\varepsilon-\delta)>0=.5$ for all I_t, d_t^*).

Can learning be improved by observing others?

An important limitation of the above analysis, however, is that it overlooks an important feature of real-world mitigation decisions: the ability to observe and learn from the behavior of others. The communal nature of mitigation decisions provides individuals with two potential aids to learning:

1. The ability to observe and analyze instances of damage caused by hazards among households that have made varying investment in mitigation; and

2. The ability to imitate the behavior of households who are believed to hold superior experiential knowledge about hazards.

While it may indeed be the case that hazards are rarely directly experienced, one could easily learn about the effectiveness of mitigation strategies by observing behavior in other settings. But would individuals naturally undertake such inferential analyses?

Empirical evidence from related research in other task settings would not seem encouraging in this regard. There is a large body of work showing that individuals often see the negative experiences of others as holding limited relevance to their own personal judgments about risk (Weinstein 1980; Weinstein and Klein 1996). Specifically, judgments about personal risk have often been found to be marked by a *social optimism bias*, that is, individuals tend to believe that their personal risk of succumbing to a probabilistic hazard (such as a disease) is less than that faced by others (Weinstein and Klein 1996). Hence, for example, while witnessing a distant neighbor's house collapse as a result of an earthquake may well trigger the belief that *the neighbor* faces heightened risk, it may do little to alter *personal* beliefs about the need for mitigation.

Even if homeowners *did* try to learn from the experiences of others, past work on the quality of intuitive inference suggests that data which is inconsistent with prior beliefs will be likely be either overlooked or ignored (see., e.g.., Klayman and Ha 1987). In much the same way that a smoker might point to the example of a 100 year old man who smokes a pack of cigarettes a day as proof that the hazards of smoking are overstated, homeowners who believe that mitigation is not cost-effective might turn to unmitigated homes that survive disasters as proof that there is no need for them to invest in protection. In short, the fact that people can observe the experiences of others provides no assurance that they will effectively learn from these data.

There is, of course, a far simpler way to learn from observation: simply mimic the behavior of other homeowners who are believed to hold superior knowledge about mitigation strategies (e.g., Coleman, Katz, and Menzel 1966). Indeed, as noted by a number of authors (e.g., Edwards 1993; Kunreuther, 1978; Lindell and Perry 2000), most individual decisions about investing in hazard mitigation tend to be made by observing and following community norms rather than by making independent assessments of costs and benefits. As noted by Bikhchandani, Hirshleifer, and Welch (1992), such a heuristic approach could be rational if one believes that strategies that perform better are more likely to be adopted than those that perform worse. In other words, following the herd can be an effective—and highly efficient—strategy for making mitigation decisions.

If, however, the initial set of decisions made in a community are themselves misguided, norm-based processes might do more to reinforce and perpetuate suboptimal actions than alleviate them. One of the greatest fears of emergency management officials is that individuals refuse to evacuate after a disaster warning because they see none of their neighbors are leaving their homes.

2. Empirical Analysis

Overview

Two experiments investigate the ability of individuals to learn optimal mitigation strategies over time by their own experience and that of their neighbors. Participants face the task of discovering the most cost-effective way of protecting their home against the risk of earthquakes. Over time they have the ability to gain this knowledge both by experimenting with mitigation investments and directly observing the resulting damage to their homes from an earthquake as well as indirectly observing the investment decisions and experiences of others living nearby. Earthquakes were chosen because the subjects (graduate and undergraduate students at a Northeastern University) would have had limited prior decision-making experience and we wished to insure that behavior reflected knowledge gained in the course of the simulation rather than previously-formed beliefs.

The two experiments had somewhat different objectives. The first was designed to provide an assessment of learning abilities in a controlled environment where participants were free to make their own investment decisions over time, but where the observed actions of others were generated by programmed agents. The second experiment examined learning abilities in a much more realistic—but less controlled—setting where participants made mitigation decisions in real community of acquaintances.

Participants and Procedure

87 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, each equipped with a personal computer. In return for participating in the experiment, all subjects received a $10 show-up fee and were told that the participant who earned the highest score in the game (defined below) would be given a $200 cash reward.

Participants were asked to imagine that they had just moved into a home in a hypothetical country that was prone to periodic earthquakes, and would be living there for 5 years. This home was valued at $40,000, and after the 5 years their performance in the task would be defined by a "wealth score", which was this initial house value ($40,000) minus the costs of repairs from experienced earthquakes and total investments in mitigation. To simplify the task, participants faced no liquidity constraints and funds not spent on mitigation or repair could not be externally invested. Participants thus could continue to buy protection and make repairs even when their wealth score dropped below zero, and there was no extraordinary penalty for negative scores. Finally, when an earthquake arose participants were automatically charged for the cost to repair the home to a like-new state, with the cost of the repair being the percentage of a home that had been damaged (hence, if a participant's home was 100% destroyed, they would be billed $40,000).

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 5-

year cycle would begin. Each subject's overall score in the simulation was their cumulative net assets after the eight replications (40 total decision periods).

The central interface, reproduced in Figure 1, consisted of a map of the hypothetical country that displayed the location of the participant's residence as well as that of other players, updated information about their current wealth, their total losses, and their current level of mitigation. As play progressed in each period (or year), buttons would appear on the interface enabling participants to navigate through four phase of decision making:

1. *Information search*. By clicking on a button that said, "learn about earthquakes and their dangers", subjects were taken to a series of research reports that provided detailed information about the frequency with which they might expect to encounter earthquakes of varying severity, the damage that a quake of a given magnitude could impose on an unmitigated home conditional on its 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. The use of this information was then discretionary.
2. *Investments in mitigation*. By clicking on a button that said, "Buy Protection" subjects were taken to an "earthquake protection store" where they could buy up to 100 "mitigation units" at a cost of $100 per unit. Investments in mitigation were cumulative and non-revocable within any 5-year cycle. In addition, by clicking a button that said, "See others' protection" a map displayed the most recent levels of mitigation undertaken by other players.

3. *Earthquake determination*. After a person made decision whether or not to gather information and/or buy mitigation, he or she clicked on a "ready" button. They were then shown a series of six green or red buttons that indicated whether or not other players had finished making their decisions (Figure 4). When all buttons turned green subjects either viewed the message, "no quake this year", or, if there was a quake, 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, as well as a text message indicting 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 experienced by each player's home. On that same screen players were also given the opportunity to view the levels of protection undertaken by each player, so they could easily toggle between information about levels of protection and levels of experienced damage.

If the participant's home suffered damage from an earthquake, the dollar loss was immediately deducted from his wealth total for that cycle of the simulation. In other words, the participant was charged the cost of repairing the damage. The home was then assumed to be rebuilt with the previously level of mitigation (if any) restored. Note that if the total amount spent on repairs and mitigation exceeded the initial home value, the the participant would have a negative wealth total 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.

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% chance that an earthquake of some magnitude would occur somewhere in the country, with its epicenter being randomly determined. In addition, there were four possible levels of earthquake intensity ranging from “minor” to “extreme”, with conditional probabilities of .5, .30, .20, and .10, respectively conveyed by means of a histogram.

The percentage of a home *i*’s value that would be lost if the quake’s epicenter was at location *j* occurred (P_{ij}) was given by the formula $P_{ij} = e^{-\alpha d_{ij}} SV(1-I)(1-m)$, where d_{ij} was the Euclidean distance between the quake’s epicenter j and the home, S was defined as the scalar measure of the quake’s strength, V was the value of the player’s home, I was the percentage of possible mitigation units purchased by the player, m was a continuous scalar parameter bounded by the range [0,1] that captured the marginal effectiveness of improvements (unknown to subjects), and α was a scaling parameter. Subjects were not given this formula, but were conveyed its meaning by being shown a histogram that plotted the percentage of a home’s value that would be lost conditional on its strength at two distances: one whether the home was at the quake epicenter and one where it was at a maximum distance from the epicenter as shown on the map.

Subjects were told that there was considerable disagreement among mitigation experts about whether investments were worthwhile, with half claiming that it was highly effective and half claiming that it was ineffective. Subjects were told that the true value was something they would need to discover on their own by experience, and 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 (the case where $m=.8$) and mitigation being ineffective for the other half (the case where $m=0$). The damage function was scaled such that subjects in the high-effectiveness conditions should have invested the maximum in protection (100 units) while those in the low-effectiveness condition should have invested zero.

The Social Feedback Manipulation. While participants were led to believe that the simulations were networked, they, in fact, were playing independently, with the information they received about the actions of other players being controlled by programmed agents. To make this manipulation convincing participants waited varying amounts of time for other "players" to finish making their decisions.

Waiting times were stochastic and a function of the time 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 belief 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 investment decision made by the participant in period *i*, and implemented the same decision with random noise in period *i+1*.
2. *Positive leaders*[1]. Simulated players gradually invested higher amounts in mitigation over the eight game replications regardless of its true effectiveness. The updating rule was $I_{it}=\alpha D_{it-1}+\varepsilon,$ where Dh_{it-1} was the damage recorded by the simulated player if there was an earthquake in the previous period, and ε was a uniform random error. The slope parameter α was chosen such that mean investments by other players was 100 units (the maximum) by the eighth round of the simulation.

Note that under the second process, simulated community decisions evolved toward perfect optimality in cases where mitigation was effective but perfect suboptimality when mitigation had no benefits.

Results

Overall learning efficiency. As noted above, in the simulation there were two asymptotic mitigation optima that depended on the true effectiveness of mitigation; subjects who discovered that mitigation was effective should have purchased 100 units of mitigation, while those who discovered that it was ineffective should have purchased 0. Of course, since effectiveness was uncertain *ex-ante*, realistically we might only hope to see *convergence* to these optima over time, as subjects learn from the experience of their own investments and those made by other (programmed) players. The analysis of simple learning rules that we presented at the outset, however, suggests that this convergence may be asymmetric: trial and error might properly lead subjects to invest more in learning

[1] A "negative leader" condition was not run due to limitations in subject resources.

when it is effective, but have more difficulty discouraging subjects from investing when it is ineffective.

Figure 2 plots the average investment levels over time when mitigation was truly effective (broken line) and truly ineffective (solid line). While the data reveal some initial awareness of the true effectiveness of mitigation among participants—investments were 25% higher when mitigation was effective during the first two cycles of home ownership compared to when it was ineffective. The data reveal that following these two cycles there was no movement toward *either* of the normative equilibria (100%, 0%). Indeed, if anything, the data show a tendency for the difference in mean mitigation between these conditions to deteriorate over time with those who were learning that mitigation was effective investing less.

To provide insights into the degree to which the mean investment levels plotted in Figure 2 were driven by individual differences in decisions about whether mitigation was worthwhile at all, in Figure 3 we plot the percentage of subjects who decided to purchase *any* protection in the first year of homeownership (when investment had their greatest value), and the mean protection conditional on purchase. The figure suggests that subjects actually approached the simulation with rather optimistic priors about value of investing in at least some mitigation; 94% purchased some mitigation when first given the opportunity—

when they had the first (mitigation had its greatest

To provide a more rigorous statistical analysis of investments over time we modeled observed mitigation levels over time as a linear function of fourteen predictors:

1. True migration effectiveness

2. Year of tenure in the home (1-5);
3. Game replication (or tenure cycle; 1-8);
4. Lagged level of mitigation;
5. Social feedback condition (2 contrasts: positive leaders v. solo play and mimicked decisions v. solo play);
6. The interactions between true effectiveness and lagged protection level and game replication
7. The interaction between social feedback condition and predictors 1, 2 and 3

The results of this analysis, reported in Table 1, lent statistical support to the qualitative observation about learning offered above. The analysis supports a modest positive effect of true effectiveness ($t=3.176$; $p=.048$), but no interaction between this effect and game replicate—implying a a lack of convergence toward the optima with increased game experience. On the other hand, the analysis suggests that subjects at least grasped the normative idea that investments in mitigation were less worthwhile as the time horizon of ownership decreased—as evidenced by a significant negative effect of year in the home ($t=-18.43$, $p<.0001$).

Did having access to decisions made by others aid learning? The analysis suggests that it did not. First, the data fail to support a significant "positive leader"-by-true effectiveness interaction ($p=.129$), implying that observing other players increase their investments over time did little to enhance (or deflate) learning of true effectiveness compared with solo play. Second, while the analysis supports a significant interaction between mirrored play and true effectiveness ($p<.001$), the coefficient of the interaction is negative in sign—implying that as game experienced increased investments by subjects in the mirrored-play social condition displayed *less discrimination* between the two true effectiveness conditions. In other words, in this case having access to social feedback appeared to hurt learning rather than enhance it.

Process analyses: why didn't participants learn? To provide deeper insights into the processes that drove participants' mitigation decisions in the task, Figure 3 superimposes plots of two features of participants' mitigation decisions over time, pooling over social feedback conditions: the relative frequency with which participants purchased mitigation (of any quantity) in the first year of each 5-year cycle of ownership, and the lagged effect of experienced earthquake damage on whether additional mitigation was purchased in each subsequent year *(p(buy|loss)*. The figure suggests that during the first several cycles of home ownership investments evolved over time through the following stylized anchor-and adjustment strategy:

> *Start each 5-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 a decreasing function of the time horizon of ownership remaining in that cycle.*

To illustrate, in each of the first four games of the simulation (each with 5 years of home ownership) 85% 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; $F(1,402)=.77; p=.38$). As shown in Figure 3, after that initial purchase, on average subsequent purchases were made in decreasing responsiveness to experienced earthquake losses, which might be interpreted as a rational response to the decreasing time horizon of hazard. For example, in the first two games an average of 50% of participants bought more protection after experiencing a quake loss in the first year of home ownership, but this decreased to an average of 13% when a quake loss was experienced in the fourth year of any cycle (the last year there would have been an opportunity to make a purchase).

After four such cycles of investment dynamics, participants then appeared to settle into a more stable process characterized by a fixed initial investment that was less likely to be revised given subsequent earthquake losses. Specifically, in the first four games when quake damage was experienced in the first year of homeownership 39% of the time, on average, this event was followed by a decision to purchase more protection. In the last four games, however, earthquake damage in the first year was followed by a decision to purchase more protection only 7% of the time (13% if the damage occurred in the second year). As shown in Figure 3, this was not due to subjects buying more initially and simply having less of an ability to buy later (the series of black dots in the figure). Rather, the data are more consistent with subjects simply concluding that their initial protection purchase was not one that they could particularly improve upon by buying more after an earthquake—though some subjects (about 10%) were still trying to adaptively learn in this manner until the very end of the task.

Friends and Neighbors effects. To further explore how decisions about mitigation were affected by the ability to observe the decisions made by others, in Figures 4a-4c we plot mean investment levels over time for each social feedback condition by true mitigation effectiveness. The Figures provide two insights about how social feedback affected decisions over time that clarify the statistical analyses previously reported in

Table 1. First, it shows that the significant negative interaction between true effectiveness and mirrored-play did not reflect globally diminishing learning abilities in this one condition, but rather poor performance by participants in one intermediate phase if the task: their sixth and seventh homes (trials 25-35). Hence, while having access to this kind of feedback did not aid learning, it also did not appear to harm it.

The Figure also suggests that subjects who were placed in communities where other decision makers gradually invested more in mitigation over time (reaching 100% by the eighth home cycle), *did* imitate this behavior, but the effect was quite small, and not conditioned by whether the high investments that other participants were making was optimal or not. For example, in the eight game replicate—when participants would have seen their peers investing 100% in mitigation—investment levels averaged of 52% in the case where the 100% was optimal (compared to 42% and 45% for the solo and mirrored-feedback conditions, respectively), and 35% when it was suboptimal (compared to 32% and 24% for the solo and mirrored-feedback conditions, respectively).

Why were participants not more influenced by peer actions? Insight into this is provided in Figure 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 mitigation of others (over the first 5 trials the mean percentage was 70%), but this interest in the actions of others rapidly diminished, averaging 35% over the last three cycles. Hence, subjects may have arrived at a strategy after the first 2-3 rounds of play that they felt no need to change. Alternatively after first being curious about the actions of others, they may have 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 experienced damage in would have, in theory, been possible for participants to quickly and conclusively discover whether mitigation was effective or not.

Discussion

The failure of subjects to learn the optimal mitigation policies in the simulation is at least somewhat surprising in light of the plethora of information subjects *could* have used to more accurately guide their decisions. In particular, within each round of play subjects had the opportunity to statistically infer the true effectiveness of mitigation by studying not only their experiences but also those of others. Likewise, they were provided with actuarial information about the levels of damage that they *should* have observed given quakes of various magnitudes and distances, something that, in principle, would have greatly aided assessments of effectiveness. Finally, most critically, they had the opportunity to learn through several rounds of play. Yet, subjects persistently underinvested in mitigation when it was optimal to heavily invest, and over-invested when it was optimal *not* to invest.

What explains this apparent learning failure? The data suggest that the inefficiencies accrued, at least in part, to a failure by subjects to *attempt* to learn during the task. Specifically, they declined to make full use of the experiential information that was available to them on each trial, and revealed a deteriorating interest in learning as the task was replicated. As time progressed subjects made fewer attempts to look at how the mitigation decisions of others was related to the damage they suffered after quakes, and their own decisions about investing were decreasingly sensitive to experienced losses over time. Hence, long-term investment levels reflected what was learned early in the task--a brief period where those in high-effectiveness conditions learned to invest more than those in low-effectiveness conditions, but far short of the level that would be optimal.

This finding, however, comes with the caveat that the failure of subjects to utilize social feedback to aid learning might have accrued to way community interaction was simulated in the task. The social setting subjects faced was a highly stylized one in which the decisions they viewed were simulated rather than real, and decisions were made in a lock-step fashion where the actions taken by others were observable only *after* all decisions had been made in a given round and submitted for play.

Subjects who sensed this fact (even if they did not articulate it) would naturally have been skeptical about its value. Hence, if a subject was unsure whether to purchase mitigation in a given round, he or she could only see what other players had done in the *previous* round—when the circumstances may well have been quite different (e.g., there would have been a longer time horizon, and there may or may not have been an earthquake). Either one of these factors could have caused subjects to see the social feedback as being much less useful for learning than would be the case in a more natural setting,

To address this concern, we designed a second experiment that allowed us to examine whether the findings of the first study generalize to a more realistic simulation where decisions were made by a real networked community of players, and where players could simultaneously observe and react to the decisions made by others.

Experiment 2

Participant and Procedure

The subjects were 109 undergraduate and graduate students who volunteered to participate in a simulation exercise similar to Experiment 1, but with two major differences:

1. The simulation was programmed in real time allowing subjects to continuously observe and respond to the changes in the investment decisions being made by other players;
2. The observed decisions were actually those of other players rather than programmed agents.

Several other enhancements were designed to increase both the likelihood that subjects would attend to the decisions being made by other players as well as enhance the task's overall realism. At the start of the simulation all house icons took the form of uncolored outlines of houses. Home-protection decisions now took the form of discrete construction improvements that had a cumulative maximum investment score of 100. These included structurally sounder chimneys, foundations, roof, walls, and/or windows and changed the color of this part of the house to orange. The current level of mitigation held by each player was displayed on the map screen both numerically and graphically in the form of a house icon where the color of different components reflected the level of mitigation. The mitigation levels of all players were in constant view throughout the simulation, and new purchases of mitigation and earthquake damages were instantly updated on the map screen.

Because the simulation was run in real time, subjects could make purchases of protection at anytime they wished during the simulation without having to wait for other players. Participants were told that each had a starting wealth of $50,000, of $40,000 reflected in the value of the home and $10,000 cash that could either be invested at a 10% rate of return (compounded and paid a fixed number of times during the course of the simulation) or used to purchase protective improvements.

Earthquakes could occur at any moment. As in the first experiment, when an earthquake occurred it was showed in an animated fashion by a set of concentric circles emanating from its epicenter. If a home was damaged by the quake, the house icon for that player momentarily changed to that of a half-collapsed house, with the percentage of damage being displayed numerically. To simulate the need to rebuild, the icon then changed to a house under construction with the word "repairing" flashing next to it. The repair period lasted between five to ten seconds depending on the amount of damage, during which time no additional protection improvements could be purchased.

The simulation was administered by having all groups play one 5-minute warm-up round followed by three 10-minute full-length rounds of the simulation. All 10-minute games utilized the same parameters, and involved an average of 4 earthquakes per game. To encourage attention to the decisions being made by others, the homeowners in each simulated community were introduced to each other prior to the start of the simulation, and their names were depicted next to each respective house icon on the game map. While verbal communication among the players was not permitted during the simulation, it was allowed in the intervals between each 10 minute simulation.

All other design aspects of the simulation were the same as in the first experiment. Subjects could learn about earthquake probabilities and consequences by clicking on a "research" button, and each group of subjects were randomly assigned to one of two mitigation-effectiveness conditions: one where mitigation was marginally effective (yielding an optimum of complete mitigation) and one where it was marginally ineffective (yielding an optimum of zero mitigation).

Results

Did playing the simulation within a more realistic social community aid learning? The answer was a strong “no”. In fact, the above changes degrade rather than enhance the degree to which subjects were able to discover the mitigation optima. Across the three 10 minute simulations, subjects who were placed in high-effectiveness mitigation environments ended each game with a mean protection level of 34 (median 26) units even though the optimum was 100. Those in low-effectiveness mitigation environments ended each game with a mean protection level of 35 units (median 24) instead of the optimum of 0. Note that these mean investment levels are lower than those observed in the first experiment, and are here not even directly consistent with the optimal levels of investment

To provide a more rigorous analysis of the factors that drove investment decisions, we regressed each player’s mean ending protection level against six predictors: the mean level of protection purchased by others in the player’s community, the number of players in the community, the true effectiveness of mitigation (a 0-1 indicator), game replicate, the amount of time a subjects spent reading research reports, and the interaction between true effectiveness and experience with the game.

The results of this analysis, reported in Table 2, provide a straightforward explanation for the lack of learning that was observed in the task: investment decisions were almost exclusively driven by a single predictor: the level of investment observed being made by others in the community ($b=.5159$; $t=7.19$; $p<.001$). In words, for each unit mean unit of increased or decreased investment observed among other players, subjects adjusted their own investments in the same direction by roughly 50%. In contrast, we see no main effect of true effectiveness (and the coefficient is nominally negative), and no interaction between effectiveness with game experience, implying no evidence of learning. The only other predictor that approached significance was that of the time spent reading research reports; those who read more—perhaps reflecting a greater concern about avoiding damage invested more, regardless of whether mitigation was effective or not.

General 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 due to a lack of recent direct experience with such events (http://www.noaanews.noaa.gov/stories2004/s2358.htm).

Are there inherent limits to our ability to learn about the effectiveness of mitigation measures from past experience? This research examined this question by reporting how samples of experimental subjects made repeated decisions about whether to invest in mitigation in two dynamic earthquake simulations. In the simulations 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 over invested when it was ineffective. There was little evidence of investments converging toward optimal levels over time whether or not one was able to observe, or were 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 (e.g., Sterman 1989). Although the mechanism that drove damage from earthquakes was deterministic, the complexity of the function would have made it difficult for participants to discern the extent to discern causality from a given damage episode (e.g., whether damage was low because mitigation was effective or the quake was ineffective.

A participant who is thinking long-term should be willing to actively experiment with mitigation investment level so that they can learn that one still incurs significant damage given a maximum investment in mitigation Consistent with prior work (e.g., Meyer and Shi 1995), this is insight seems to elude participants. There were no cases in either study of subjects purchasing 100% of available protection at the start of the task to test a hypothesis about effectiveness. Rather, the modal strategy was to purchase a limited amout (e.g., 25-30%)—a quantity that would be insufficient to provide significant protection if mitigation was effective, or, in turn, be informative as to whether it is effective or not.

Participants also made limited use of the most readily-accessible source of information about effectiveness at their disposal, the experiences of other decision makers. In Experiment 2, where the decisions they were observing were those of other real decision makers, investment decisions appeared herd-like, with the single greatest driver of investment decisions being the modal investment being made by others. Sub-optimality thus ended up being self-reinforcing. As more participants imitate the behavior of others, there is less opportunity for them to learn. The evidence from Experiment 1, however, suggests that herd instincts have their When individuals saw controlled feedback that revealed to them that others were following suboptimal policies such as investing in mitigation when they shouldn't, the tendency to make decisions by imitation vanished. Even in this case individuals were did not discover the optimal mitigation strategy.

The data also suggest that subjects focus on what they discovered about mitigation in the early rounds and simply tired of the task of learning. The first experiment provided the clearest view of this effect. Decisions appeared to be characterized by a simple anchor-and adjustment policy: participants started each round of decisions by investing in a moderate amount of mitigation, then bought more if they experienced a loss. After 2-3 cycles of applying this policy, participants seemed to abandon further attempts to update this strategy. Reactions to experienced losses diminished, as did their interest in viewing the mitigation decisions made by other players.

Limitations and Future work

Whether the results reported here serve as a good model of limits to learning that that might arise in real settings is unclear. 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 that would arise in the real-world. On the other hand, subjects 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 biases observed in the simulation experiments appear consistent with errors in mitigation decisions that have been noted in real-world settings. For example, there is considerable empirical evidence of herd behavior in decisions about whether to take preparedness actions in the face of hazards (e.g., Baker 1991; Edwards 1993), and the overriding importance of direct encounters with hazards as a basis for perceptions of risk (e.g.,, Lindell and Perry 2000) and undertaking mitigation measures. (Kunreuther 1978) . What is perhaps most surprising about the findings reported here is that replication will not cure these biases nor that can they be reduced simply by putting more information at the hands of decision makers.

Finally, we see this work as also highlighting the potential value of dynamic laboratory simulations as a tool for gaining a better understanding of human response to natural hazards. To date our knowledge 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 compliment by providing a means for testing hypotheses about hazard response that may emerge from field work as well as providing 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.

References

Andrus, F.M. (1952), "Earthquake design requirements of the Uniform Building Code." In Proceedings of the Symposium on Earthquake Blast Effects on Structures. Berkeley, 1952. pp. 314-316.

Atwood, L. Erwin and Ann Marie Major (1998), "Exploring the 'Cry Wolf' Hypothesis," *International Journal of Mass Emergencies and Disasters*, Vol. 16(3), 279-302.

Camerer, Colin, and Howard Kunreuther (1989), "Decision Policies for Low Probability Events: Policy Implications", *Journal of Policy Analysis and Management*, 8, 565-592.

Coleman, J. S., Katz, E., & Menzel, H. (1966). <u>Medical Innovation: A Diffusion Study.</u> New York: Bobbs-Merrill

Baker, E. J. (1991). Hurricane evacuation behavior. *International Journal of Mass Emergencies and Disasters*, 9, 287-310.

Bikhchandani, Sushil, Hirschleiffer David, and Ivo Welch (1992), "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades", Journal of Political Economy, 100 (5), 992-1026

Brehmer, Berndt (1980), "In one word: not from Experience", *Acta Psychologica*, 45, 223-241.

Edwards, M. L. (1993). Social location and self-protective behavior: Implications for earthquake preparedness. *International Journal of Mass Emergencies and Disasters*, II, 293-304.

Erev, Ido, and Alvin Roth (1998), "Predicting how people learn in games: reinforcement learning in experimental games with unique, mixed strategy equilibria", American Economic Review, 88(4)-881.

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

Grether, David (1980), "Bayes Rule as a Descriptive Model: The Representativeness Heuristic", *Quarterly Journal of Economics*, 95 (November) 537-557.

Kahneman, Daniel and Lovallo, Dan (1993). "Timid choices and bold forecasts: A cognitive perspective on risk taking." *Management Science*, 39, 17-31.

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., R. J., Meyer, R. Zeckhauser, P. Slovic, B. Schwartz, C. Schade C., M.F. Luce, S. Lippman S., D. Krantz, and R. Hogarth R. (2002), "High-Stakes Decision Making: Normative, Descriptive, and Prescriptive Considerations", *Marketing Letters,* 13(3), 259-268.

Kunreuther, Howard, Onculer, A., and Paul Slovic (1998), "Time Insensitivity for Protective Investments", *Journal of Risk and Uncertainty*, 16, 279-299.

Lindell, Michael K., and Perry, Ronald W. (2000), Household Adjustment to Earthquake Hazard: A Review of Research, *Environment and Behavior*, 32(4), 461-502.

Lindell, M. K., & Perry, R. W. (1990). Effects of the Chernobyl accident on public perceptions of nuclear plant accident risks. *Risk Analysis*, 10, 393-399.

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

Meyer, Robert J., and Young Shi (1995), "Intuitive Solutions to the Armed-Bandit Problem", *Management Science,* 33, 234-267.

Mulilis, J-P., Duval, T. S., & Lippa, R. A. (1990). The effects of a large destructive earthquake on earthquake preparedness as assessed by an earthquake preparedness scale. *Natural Hazards*, 3, 357-371.

Nachbar, J.H. (1997), "Prediction, Optimization, and Learning in Games", *Econometrica* 65, 275-309

Peacock, Walter (2003), "Hurricane Mitigation Status and Factors Influencing Hurricane Mitigation Status Among Florida's Single-Family Homeowners", *Natural Hazards Review*, 4(3), 149-158.

Pengergrass, Crystal (1999), "Tornado Effects on Residential Structures", unpulblished manuscript, RDU program, University of Alabama, Birmingham.

Russell, L., Goltz, J. D., & Bourque, L. B. (1995). Preparedness and hazard mitigation actions before and after two earthquakes. *Environment and Behavior*, 27, 744-770.

Samuelson, William and Richard Zeckhauser. (1988), "Status Quo Bias in Decision Making," *Journal of Risk and Uncertainty* 1, 7–59

Sterman, John (1989), "Misperceptions of Feedback in Dynamic Decision Making", Organizational Behavior and Human Decision Processes", 43, 301-335.

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

Weinstein, Neil D, William M. Klein (1996). "Unrealistic optimism: Present and future," *Journal of Social & Clinical Psychology*,15(1), 1-8.

Table 1: OLS Regression of mitigation levels over time, Experiment 1

Parameter	**Estimate**	**Std 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: $F(14,3187)=485.1$, $p<.001$; $R^2=.68$

Table 2: OLS Regression of ending mitigation levels, Experiment 2

Parameter	**Estimate**	**Std Error**	**T**	**p(T)**
Intercept	.1832	.0620	2.96	.0034
True Effectiveness	-.0239	.0449	-.53	.5963
Replicate	-.0064	.0173	-.37	.7106
Group Size	-.0039	.0117	-.33	.7424
Time spent on Research	.0004	.0028	1.31	.1910
Mean investment by other players	.5159	.0718	7.19	<.0001
Effectiveness*Replicate	.0435	.0617	.71	.4812

Overall model: $F(6,257)=9.853$, $p<.001$; $R^2=.187$

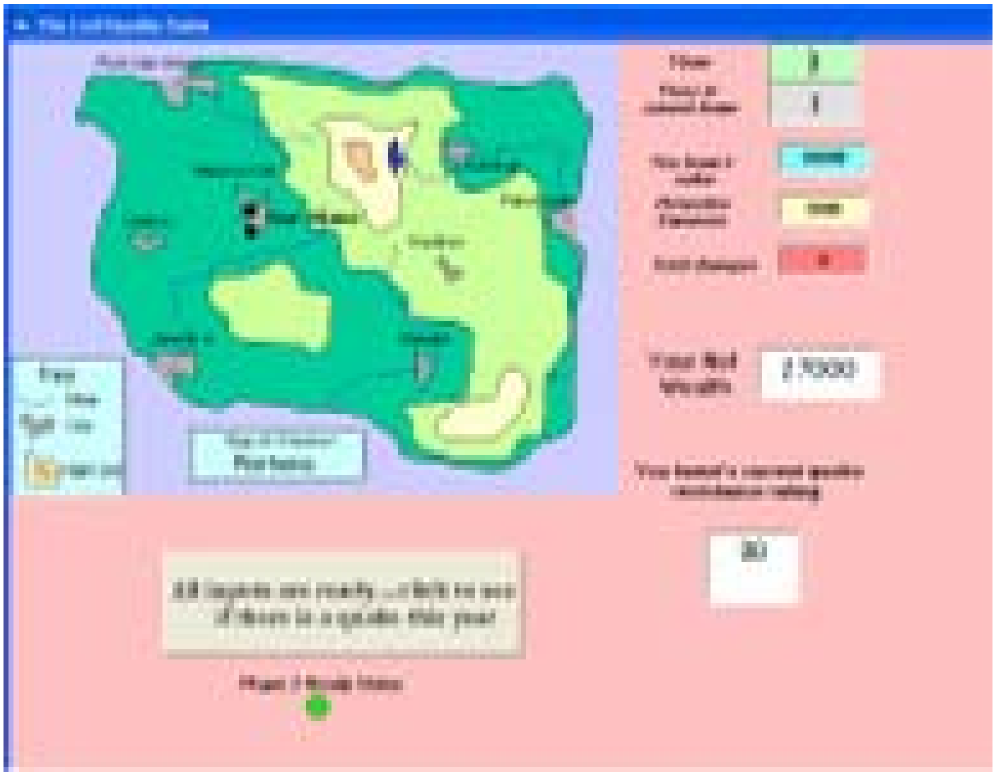

Figure 1: Main simulation interface used in Experiment.

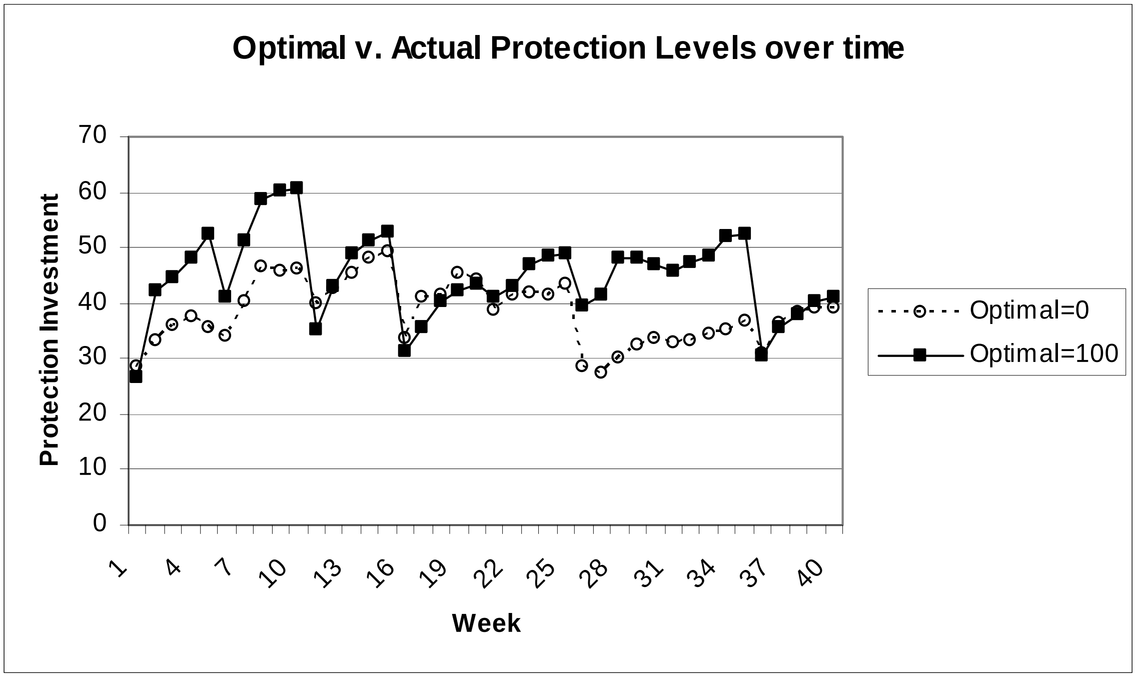

Figure 2: Average observed investment levels over time by true mitigation effectiveness. Oscillations accrue to increasing investments over 8 five-year ownership blocks, Experiment 1.

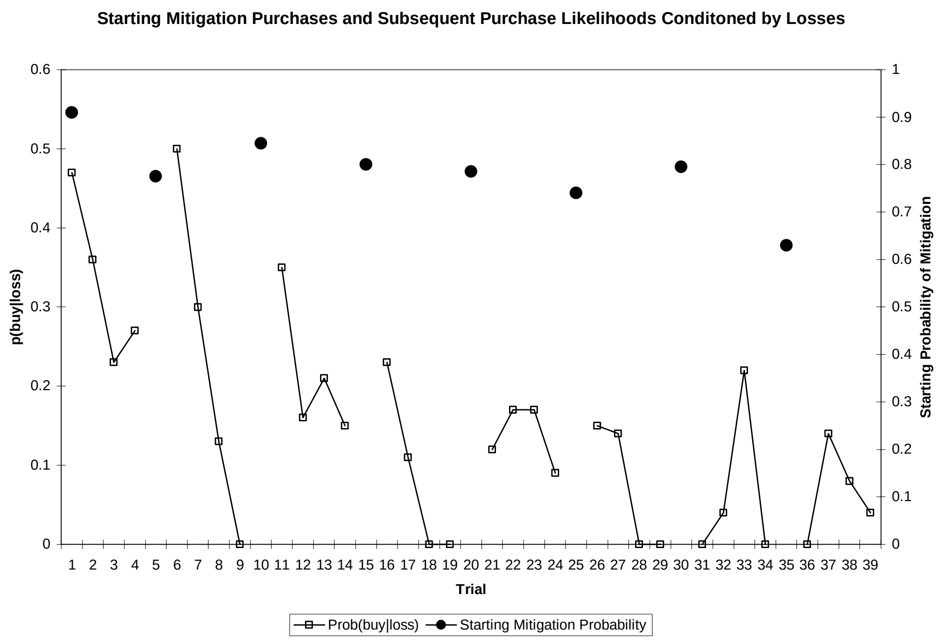

Figure 3: Changes in purchased mitigation at the start of each 5-year tenure (solid dots), and the probability that an additional purchase was made in each subsequent period conditional on a lagged experienced loss (lines) , Experiment 1.

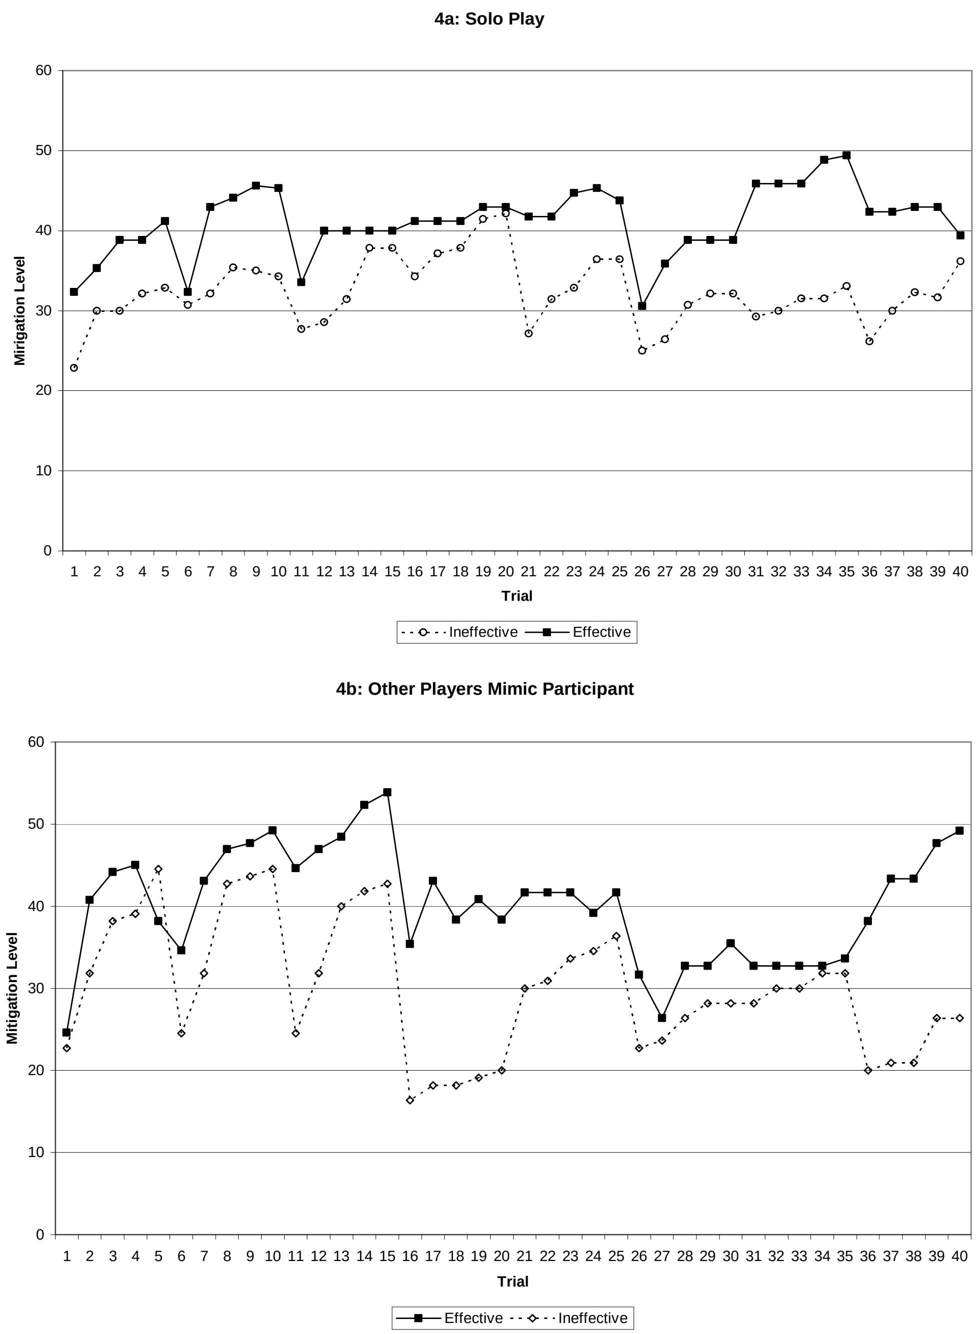

Figure 4 Average observed investment levels by social-feedback condition and true mitigation effectiveness over trials, Experiment 1.

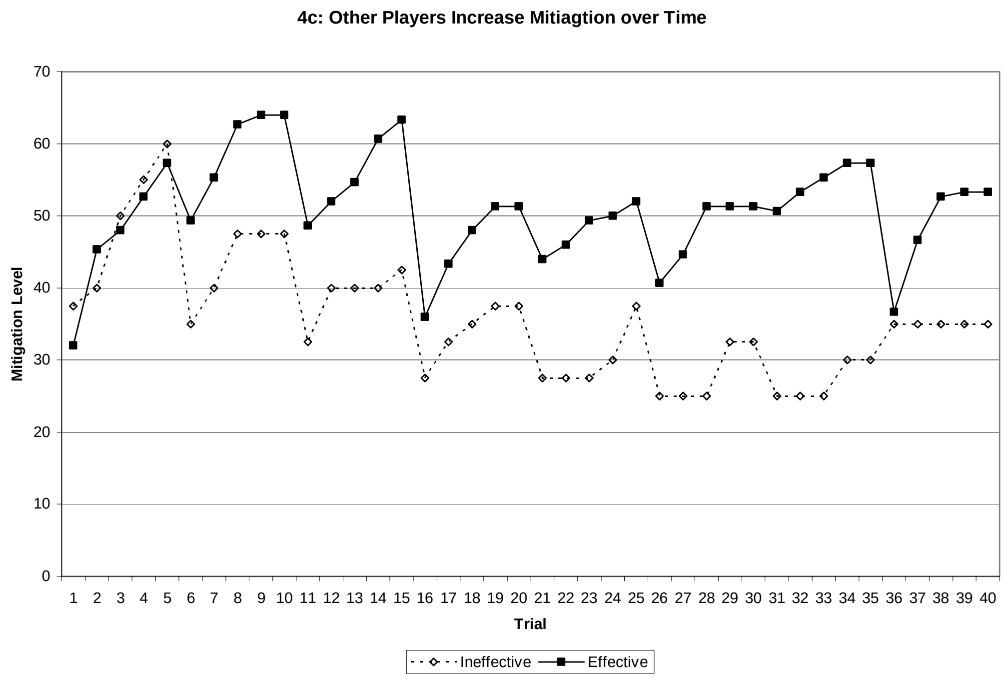

Figure 4 (continued): Average observed investment levels by social-feedback condition and true mitigation effectiveness over trials, Experiment 1.

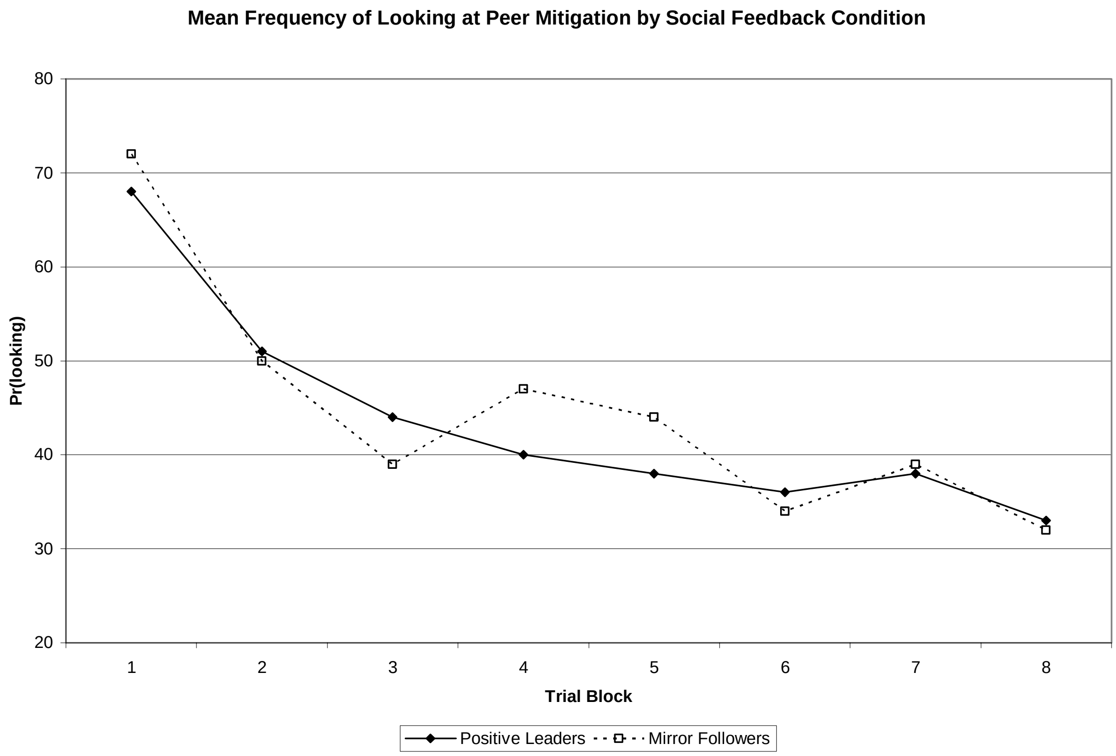

Figure 5. Mean relative frequency with which participants looked at the mitigation levels of other players by social feedback condition and trail block, Experiment 1