THE TREE OF EXPERIENCE IN THE FOREST OF INFORMATION: OVERWEIGHING EXPERIENCED RELATIVE TO OBSERVED INFORMATION *

Uri Simonsohn
University of Pennsylvania

Niklas Karlsson
Göteborg University

George Loewenstein
Carnegie Mellon University

Dan Ariely
MIT

* An earlier version of this paper circulated with the title "Overweighing personal relative to observed information: evidence from repeated games"

THE TREE OF EXPERIENCE IN THE FOREST OF INFORMATION: OVERWEIGHING EXPERIENCED RELATIVE TO OBSERVED INFORMATION

Abstract:

Standard economic models assume that the weight given to information from different sources depends exclusively on its diagnosticity. In this paper we study whether the same piece of information is weighted more heavily simply because it arose from direct experience rather than from observation. We investigate this possibility by conducting repeated game experiments in which groups of players are randomly rematched on every round and receive feedback about the actions and outcomes of all players. We find that participants' actions are influenced more strongly by the behavior of players they directly interact with than by those they only observe.

Keywords: experiential learning, observational learning, repeated games, prisoner dilemma, weak-link, behavioral economics.

"Experience is a dear teacher, but fools will learn at no other" Benjamin Franklin

1. Introduction

Information often comes from multiple sources. Standard economic models assume that people weigh different sources according to their diagnosticity. Holding diagnosticity constant, how a specific piece of information was obtained should be irrelevant.

One important distinction, when it comes to the process leading to acquisition of information, is whether the information was obtained through personal experience – i.e., in a process that had or could have had direct consequences for oneself -- or only by observing the experience of others. We refer to the former as "experienced information" and the latter as "observed information." [1]

It has long been hypothesized that experienced information, by virtue of being more emotional, concrete and proximate has a greater impact on behavior than observed information (Nisbett and Ross, 1980). As we summarize in some detail later, however, existing research has either failed to find support for such a difference, or it has found it but with research designs that are subject to confounds such as differences in informational content, format or relevance. In this paper we present results from repeated game experiments that vary the dimension of experience while maintaining content, format and relevance constant. We find that experienced information does have a greater impact on behavior than observed information. Indeed, in two of our three experiments we find that *only* experienced information has a statistically significant influence on behavior, even when observed information is equally informative.

[1] Although the vast majority of information people obtain can be classified as either experienced or observed, some cannot, such as new information that arises from mentally processing exiting information, be it experienced or observed, or information which is obtained from fictional situations, or from sampling without payoffs. In this paper we focus exclusively on the contrast between observed and experienced information

All three experiments employed the same basic paradigm. Groups of subjects participated in several rounds of a game, the Weak-link in experiment 1 and the Prisoner's Dilemma in experiments 2 and 3. Subjects were randomly assigned to subgroups that played independently of one-another. In every round, each of the subgroups played the game once. After every round, each subject observed the actions of all participants in all subgroups. Subjects were then reassigned to new subgroups (again randomly formed) for the next round. This process was repeated until the end of the experiment.

The key feature of these experiments is that they vary only whether information is experienced or observed, maintaining content, format and relevance of information constant. All subjects see the actions taken by all other subjects, presented in the same format, and because of the random rematching from round to round, all actions are equally relevant to everyone -- equally diagnostic of what they can expect to encounter in the future.

Because subjects played in subgroups, however, a given action was directly experienced by some subjects (those in the same subgroup) but only observed by others (those in other subgroups). If experienced information has a greater impact on behavior than observed information, then subjects' actions should respond more strongly to the actions of subjects in their previous subgroup than to those in other subgroups.

Consistent with this prediction, we find that in the Weak Link experiment, subjects responded roughly twice as strong to the actions of players in their previous group than to actions of players in other groups, while in the Prisoner Dilemma experiments we find that only experienced information has a significant impact on behavior.

Two possible mechanisms could account for a behavioral premium on experienced information. People may simply not attend to observed information (i.e., it is essentially not

really *observed*), or they may attend but not respond to it (i.e., it gets weighted less). We attempt to distinguish between these two explanations in experiments 2 and 3. In experiment 2, after subjects repeatedly played a prisoner dilemma game with different partners, but were informed of the actions of other players, we asked them to recall how many of the subjects they did not interact with defected. Their answers were highly accurate, suggesting they attended to the actions of players they were not directly paired with. However, their own subsequent play was influenced only by the person they had just interacted with. In experiment 3 after each round subjects wrote down the actions of all subjects on a piece of paper. This guaranteed that subjects attended to the observed information and yet, as in study 2, only experienced information was a significant predictor of behavior. Taken together these experiments strongly suggest that the behavioral premium placed on experienced information is caused by placing a greater weight on it, rather than to lack of attention to observed information.

In studying the dynamics of behavior in experimental games, our experiments are related to the learning literature within experimental game theory. This literature seeks to explain the sequential choices of people playing strategic games (for a review see chapter 6 in Camerer, 2003). It traditionally distinguishes between two main types of learning: simple-reinforcement (where the tendency to take a future action depends on how profitable it has proven in the past) and belief-based (where future actions are influenced by how profitable different actions *would have been*, had they been taken in the past).

Camerer and Ho (1999) proposed a model called "Experience-Weighted Attraction" (EWA), which, by combining both types of learning in a parsimonious way, successfully accounts for behavior across a broad range of games (see also Camerer, Ho and Chong, 2002). Particularly relevant to this paper, the authors used EWA to estimate the relative impact on

learning of actual payoffs vs. that of forgone payoffs. They find, across three different sets of games, that the impact of the former is roughly twice that of the latter (Camerer & Ho, 1999). This premium probably arises because simple-reinforcement has a greater emotional component than belief-based learning (Camerer et al., 2005).

Note that both simple-reinforcement and belief-based learning rely on experienced information, the former on decisions' actual consequences, and the later on their counterfactual ones. In this paper we propose that it is important to contrast the behavioral consequences of this source of learning (experienced information) with another source, that which occurs by considering the actual and counterfactual consequences of the actions of others (observed information).

In what follows, section 2 summarizes earlier related research, sections 3, 4 and 5 present results from the three experiments and section 6 concludes.

2. Related research

This section reviews several independent lines of research that are related to the question of how people weigh directly experienced as compared with merely observed information. We begin with the literature that utilizes field data to estimate the impact of personal experience on preventive behavior. We then review experiments from the psychology literature that attempt to contrast the influence of experienced and observed information, and finally we summarize the experimental economics literature that studies the impact on behavior of observing other players play the same game.

2.1 Field evidence of personal experience and self-protective behavior

A large literature has examined the impact of personal experience on a variety of self-protective behaviors. The findings from this literature vary across domains. For example, there is no evidence that personal experience with accidents influences seat-belt use (see e.g. Svenson et al., 1985; Weinstein 1987), there is mixed evidence as to whether direct experience with crime leads to an increase in crime-preventive behaviors, such as leaving doors locked or purchasing an alarm, (DuBow et al., 1979; Hough, 1985; Parker et al., 1980; Skogan, 1987; Svenson et al., 1985; Tyler, 1980; Weinstein1987) and there is robust evidence that direct experience with a natural disaster increases preparedness for future disasters (see e.g. Halpern-Felsher et al., 2001; Kunreuther, 1978; Lindell and Perry, 2000; Schiff, 1977).

In the most comprehensive review of this literature, Weinstein (1989) – himself a contributor -- points out that much of the existing research has been subject to serious methodological problems, such as crude all-or-none measures of past experience, and the common use of self-report measures of behavior as the dependent variable. Most importantly, however, these studies only collect data on experienced information (e.g. whether or not somebody was the victim of a crime) and not on observed information (e.g. whether they are aware of the crime rate in the neighborhood where they live). Even if these field studies were free of the methodological limitations pointed out by Weinstein, therefore, they could document *an* effect of personal experience on behavior, but they lack the necessary comparison to determine whether a *premium* is placed on experience.

2.2 Psychology experiments contrasting the impact of experienced and observed information

A few papers have contrasted risk preferences when people receive descriptions of the distribution of outcomes with when they experience a sequence of outcomes drawn from the same distribution (see e.g. Barron and Erev, 2003; Hertwig et al., 2004; Yechiam et al., 2005). These papers find that, consistent with prospect theory, choices based on descriptions are made as if people overweighed small probabilities, but that choices made based on experience underweight them. In these studies, unlike ours, information is not held constant across conditions because subjects in the description-based condition received the objective distribution of outcomes, while subjects in the experienced-information condition rely on their (incompletely) sampled information. Indeed, differences in information across the conditions is cited by the authors as one possible explanation for why people react differently to the two types of information.

Lind, Kray and Thompson (1998) compared the effects of personal experience of unfair treatment with reports of unfair treatment by others. Groups of three subjects interacted with a "supervisor," a confederate, who in one condition behaved unfairly once with each of the three members of the group, and in the other condition three times with the same member. The authors found that groups with an equal distribution of unfairness (i.e. where all members had experienced unfairness once) rated the supervisor as less fair. Lind et al. interpret this as suggesting that "people place greater weight on their own experience of injustice than on that of others" (page 17). However, the experiential and observational conditions also differed in ways other than just on whether information was experienced or not; subjects who received the unfair treatment knew more about it than subjects who only learned of it through a short conversation (groups of three subjects had 2 minutes total to share their experience during a 20 minute

experiment). In addition, the distribution of unfairness may be in itself informative: subjects may have inferred that if only one of them felt the supervisor was unfair, the problem may lay with the subject rather than the supervisor (in fact, this is one of the explanations proposed by the authors).

Braver and Rohrer (1978) examined the degree of cooperation in a Prisoner's Dilemma game by participants who had previously played 20 rounds of the game against a confederate who always played *tit-for-tat*, and by participants who had observed the same 20 rounds. Observers were slightly *more* likely to cooperate than the actual players (in a new game), but this difference was significant only under somewhat arbitrary specifications (overall cooperation rate did not differ). However, this is an extremely weak test of the difference between observed and experienced information. Unless the participant playing against the confederate switched actions sufficiently often, neither the actor nor the observer would have been able to figure out that the confederate was playing a tit-for-tat strategy.

Olsson and Phelps (2004) compared skin-conductance reactions to the presentation of a target stimulus between subjects who had previously received electric shocks immediately following the presentation of that stimulus and subjects who had observed a *confederate* receive such shocks. They found a slight directional effect towards greater "learning" (i.e. greater increase in skin conductivity) in the experiential learning condition, but this difference did not approach significance. As was the case in previous experiments, informational content was not maintained constant across conditions, however, since subjects experiencing the electric shock (which was calibrated to be "uncomfortable but not painful") knew more about what that shock felt like than subjects who only observed a confederate receive such a shock, and hence (non-

significant) differences in skin conductivity may have resulted from differences in informational content rather than in the process by which information was obtained. [2]

2.3 Observational learning studies in experimental economics

A few experimental economics papers have studied the impact of observing others play on learning. Schotter and Merlo (2003) had one set of subjects engage in multiple attempts to solve a task with one correct answer, while another set of subjects watched. All subjects then participated in a similar task (the only difference was greater stakes). The performance of observers in the high-stakes task was actually better than that of those who had practiced. Schotter and Merlo speculated that their results may be due to the fact that players only pay attention to their realized payoffs while observers pay more attention to counterfactual payoffs because they are engaged in a more abstract processing mode (this interpretation of their findings is consistent with earlier results by the same authors, Merlo and Schotter, 1999).[3]

Duffy and Feltovich (1999) compare the behavior of dyads playing *ultimatum* and *best-shot* games either in private or while observing (and being observed by) another dyad. They obtained mixed results with respect to the impact of observing the behavior of others on behaving closer to Nash equilibrium predictions (further from equilibrium in the ultimatum game, no effect in the best-shot game). They speculate that differences in the *content* of information across conditions may explain the mixed nature of their results. Since in the condition where subjects observed others they themselves were being observed, differences

[2] Olsson and Phelps were interested on whether learning occurred at all across different conditions and hence they did not report the significance of the *difference* between experiential and observational learning in their paper. Through personal communication with Andreas Olsson we learnt that this difference was not significant.
[3] Note that if this was the *only* difference between experiential and observational learning we would not expect subjects who are engaging in the same task (as they do in our experiments) to weigh experienced and observed information differently.

across conditions (and experiments) may be driven by the impact of being observed (e.g. reputation concerns) rather than by that of observing. Ballinger, Palumbo, and Wilcox (2003) ran experiments where overlapping "generations" of subjects made saving decisions. They found that later generations, which observed (and received advice from) earlier ones, accumulated savings that were closer to the optimum.

In sum, although several independent lines of research within both economics and psychology have investigated issues that are related to the question at hand, most have concentrated on studying whether experienced information influences behavior at all (in the self-protective behavior literature) or whether observed information influences behavior at all (in the experimental economics literature). The few studies that have contrasted experienced and observed information, in the psychology literature, have not maintained all other aspects of information constant across conditions, and hence their (mixed) results do not address the question of interest here: does *the same* piece of information influence behavior differently depending on whether it was experienced or observed?

3. Experiment 1 – Weak-link game

The first experiment employed the Weak-link (WL) game, first studied by (Van Huyck et al., 1990). The WL is a coordination game in which all players simultaneously choose a number between 1 and 7. Payoffs increase with the minimum of all the numbers chosen in a group, and diminish with the distance between an individual's choice and the minimum. Equilibrium arises when all subjects choose the same number, because, conditional on everyone *else* choosing the same number, no individual can be better off by choosing a different one. Players are not, however, indifferent between the seven equilibria since the higher the number they coordinate

on, the higher is the payoff to everyone. Table 1 shows the payoff structure of the game we used, mapping the relationship between an individual's choice, the minimum choice in the subgroup and the individual's payoff.

*** Table 1 ***

The WL game resembles several real-life group decision making situations in which the member of a group exerting the least effort affects the outcomes of all members. A common example is a lunch meeting that cannot begin until all participants have arrived. No one wants to arrive before the last person arrives, but everybody is worse off waiting. Any arrival time can be an equilibrium, as long as all participants manage to coordinate on simultaneous arrival, but everyone arriving on time is the preferred equilibrium.

Several papers have looked at how people play the WL game (see e.g. Camerer and Cachon, 1996; Camerer and Knez, 2000; Schotter and Merlo, 2003; Weber, 2006) . These studies find that most people initially choose high numbers, but that groups of three or more players converge towards lower ranked equilibria rather quickly, presumably as a consequence of the fact that a few participants choose low numbers early on.

Because WL game participants appear to make current choices dependent on the behavior of others in previous rounds, we thought it was an ideal game to use in a study comparing experienced with observed information. Consistent with the notion that people place a premium on experienced information, our prediction is that subjects' choices will be more strongly influenced by the choices of their previous counterparts than by the choices of subjects assigned to other groups, even though, due to random rematching in each round, the behavior of

all subjects is equally relevant in terms of what subjects could expect to encounter in the next round.

3.1 Method

We ran the WL experiment at the Göteborg University in Sweden. Participants were recruited from a pool of undergraduates who had signed up for participation in experiments. For each session, we arranged with 16 subjects to come to the lab. The first 12 subjects to arrive were part of the experiment and latecomers were asked to do another task in a separate room. The six sessions we ran included a total of 72 subjects. Subjects sat in front of computer terminals where they could not see other subjects' screens. The experiment was web-based, run on a specially programmed PHP script (instructions are included in the appendix).

After reading the instructions, subjects were presented with two hypothetical outcomes for WL games, and were asked to calculate the payoffs that these games would have produced. The experiment began only after everyone had correctly calculated these payoffs, showing that they understood the game.

WL games were played independently by three subgroups consisting of four subjects each. A player's payoffs for any given round depended exclusively on the behavior of other players in the same subgroup. At the end of every round, subjects were randomly re-assigned to a new subgroup formed from the 11 other people in the room (including, but not limited to, the people they had just played with), with their identities kept hidden. Figure 1 shows a screenshot of the webpage on which subjects chose a number in the 1-7 interval by selecting the appropriate digit from the pulldown menu.

Figure 1

In each round, after all participants had chosen their number, they all saw everyone's actions and were reassigned to a new subgroup. This process was repeated until the end of the

experiment. Figure 2 shows a screenshot of the webpage that subjects saw between rounds. After all l0 rounds were completed, one was randomly chosen to count for real money.

** Figure 2 **

3.2 Results

Figure 3 shows the distribution of choices for the 72 subjects through the 10 rounds. Consistent with prior WL experiments, lower numbers become more common as the rounds progress. Table 2 shows the average minimum per-round, which starts at 3.72 and ends at 2.61.

Figure 3 *

Table 2

Our key hypothesis is that subjects put more weight on the behavior of subjects in their previous subgroup than on the behavior of subjects in other subgroups. To test this hypothesis, we conducted two different analyses of the data. The first consisted of running regressions in which the dependent variable was the choice of each subject on each round and the key predictor was the minimum in the subject's own-group in the previous round, controlling for the minima in each of the three groups. It assesses whether people put more weight on their own group by testing if the coefficient of the minimum of a subject's own group is significant.

Although this strategy has the advantage of being simple and straightforward, it has the disadvantage of utilizing only the minima in each group; it does not directly take into account the choices of 8 out of the 11 other players.

The second strategy consisted of regressing subjects' choices not on the minima of the groups, but on a function of the subjective probabilities associated with each possible minimum in the subject's next group, probabilities which in turn depend on the choices of all eleven players. The premium placed on experienced information is assessed by estimating the relative weighting of the 3 choices in the subject's own group, relative to the 8 choices in the other two groups, that leads to subjective probabilities that best match subjects' actual choices.

In what follows we present the results from both strategies. The estimated premium on experienced information is approximately the same in both: subjects behave as if they weighted the behavior of subjects in their own group around twice as heavily as the behavior of subjects in other groups.

3.2.1 Empirical strategy 1 – Using minima in previous round to predict behavior

We regressed the choice of player *i* in round *k* on the minimum choice in each of the three subgroups in round *k-1*, plus a slope dummy for the group in which player *i* participated in round *k-1*, and we controlled for the individual's previous choice and for the round number in the experiment. This regression is represented by equation 1 below,

$$b_{i,k} = \alpha_0 + \alpha_1 b_{i,k-1} + \alpha_2 \text{Min}_{1,k-1} + \alpha_3 \text{Min}_{2,k-1} + \alpha_4 \text{Min}_{3,k-1} + \alpha_5 \text{OwnGroupMin}_{i,k-1} + \alpha_6 \text{Round} + \alpha_7 \text{Group}_1 + \alpha_8 \text{Group}_2 + \varepsilon_{i,k} \quad (1)$$

Where:

- $b_{i,k}$ is the number chosen by individual *i* in round *k*, (1-7)
- $Min_{j,k}$ is the minimum number in subgroup *j*, in round *k* (with j=1,2,3)
- $OwnGroupMin_{i,k-1}$ is the minimum number in the group where subject *i* participated in round *k-1*.
- Round corresponds to the round number (2-10), and
- $Group_1$ and $Group_2$ are dummies which take the value of 1 if the subject was in groups 1 and 2 respectively, and 0 otherwise. These dummies are added because $OwnGroupMin_{,k-1}$ can be thought of as an interaction between having played in a specific subgroup and the minimum in that group, and hence the regression must include the main effects as well.

If subjects do not make a distinction between the choices of players in their previous own group and of those in other groups, α_5 should be 0. This would mean that once we know the minima in each of the three groups, it makes no difference which minimum was experienced by the subject. If, on the other hand, subjects place a higher weight on their own subgroup, α_5 should be positive.

Since the dependent variable is discrete and bounded we estimated Ordered Probit regressions (we obtain qualitatively identical results via OLS). Because the regression includes several choices for each individual we report clustered (by subject) standard errors. Table 3 shows the results from estimating this regression.

Table 3

Column 1 only includes the minimum choice in the group where a subject played on the previous round, and controls for the subject's previous choice and for round number. The effect of the own-group's minimum is estimated as positive and significant ($\alpha_5 = 0.119$, $p < .0001$). This result is analogous to the field studies reviewed previously, documenting *an* influence of personal experience on preventive behavior.

Only after controlling for the minima in all three groups, however, can we interpret α_5 as indicative of subjects weighting experienced information *more heavily* than observed information. In column (2), therefore, we control for the minima in all three groups and hence the parameter estimate of α_5 corresponds to a slope dummy testing whether individuals weight the minimum of their own subgroup more heavily. As predicted own-group remains positive and significant after controlling for observed information ($\alpha_5 = .066$, $p = .042$).[4] The coefficient estimate suggests a sizeable impact of experienced information. According to the estimates from column 2, individuals are nearly twice as sensitive to the minimum choice in a group they previously participated in, than to the ones on which they did not.[5] For completeness, column (3) presents the results if the slope dummy for own-group is excluded.

An interesting question is whether subjects' excessive reliance on experienced information diminishes with experience. To address this question we estimated the regression presented in column (2) separately for rounds 2-5 and rounds 6-10. The resulting point estimates

[4] For those individuals who chose the lowest number in their subgroup in the previous round, we use the second lowest value in their group as the key predictor, in order to avoid using their own past actions to predict their behavior. Our results do not change if we use the second lowest choice in all subgroups for all subjects ($\alpha_5=.093$, $p=.018$).

[5] It is surprising that the point estimate of α_3 is so much larger than those of α_2 or α_4. This may be in part the consequence of group 2's results being presented in the center of the screen. This highlights the potential role of attention in our findings, which we address in experiments 2 and 3.

of α_5 were .054 and .085 respectively, suggesting that, if anything, subjects *increased* the weighting of their own group compared to that of others, although the change is not significant. Alternatively, if the regression is estimated with an interaction term between round and the minimum in a subject's own group, the point estimate of such interaction is positive (B = .013), though not statistically different from 0 (p = .28).

3.2.2 Empirical strategy 2 – Using subjectively expected minimum in next round to predict choices

As noted, although strategy 1 is simple and straightforward, only 3 out of every 11 choices of other players are incorporated directly in its analysis. Strategy 2, in contrast, uses all of them. The second strategy builds upon the logic that subjects' choices should depend on what they expect their groups' minima to be in the next round, beliefs which, in turn, should be a function of *all* 11 choices in the previous round.

The second strategy consists of estimating the hypothesized premium on experienced information by computing the implicit weights placed on experienced and observed actions that lead to subjective probabilities that best fit subjects' actual choices.

In particular, suppose that, to estimate the probabilities associated with their group minimum in the next round being equal to each of the possible seven values, participants weighted each of the choices by the eight subjects not in their previous group (i.e. observed information) by a weight of 1, and those of subjects in their own group (i.e. experienced information) by a weight of θ. Strategy 2 consists of estimating the θ that leads to subjective probabilities that best explain the observed actions. If subjects weighted experienced

information and observed information equally, then θ=1 would maximize fit, while if subjects only responded to experienced information, θ would be estimated as diverging towards infinity.

Let $p_{i,k}(m_{i,k}|\theta,H_{i,k-1})$ represent subject *i*'s subjective probability that the minimum in her group in round *k*, will equal *m*, with:

$m \in [1,7]$, and

$H_{i,k-1}$ consisting the full history in the previous round (i.e. the actions of the other 11 subjects in the experiment in round k-1).

Intuitively subjects can be thought of as generating $p_{i,k}$ by drawing 3 balls (for the corresponding 3 other players that will form their subgroup) from an urn where there are 8 balls representing the 8 actions of subjects in the other two subgroups, and 3*θ balls representing the experienced subjects.

For example, suppose that in one round, of the 11 other players, 10 chose the number seven, and 1 chose the number one. Subjects from the group where a player chose one will form subjective probabilities as if drawing from an urn where there are (8+2θ) balls with the number seven and θ balls with the number one, while subjects from the other groups will draw from an urn with (7+3θ) balls with the number seven and just 1 ball with the number one. Obviously, if θ = 1 all subjects would assign the same subjective probabilities to each minima regardless of which group they were part of in the previous round (in particular

$p_{i,k}(1 \mid \theta = 1, H_{i,k-1}) = 1 - \frac{10}{11}\frac{9}{10}\frac{8}{9} = .27$, for all *i*).

If θ > 1, on the other hand, subjects who were in the group were a player chose the number one would have a higher subjective probability of the minimum in the next round being equal to one, than subjects who were not in that group. For example, with θ=2, if subject *i*

participated in the group where the number one was chosen and subject *j* did not, their subjective probabilities would be $p_{i,k}(1 \mid \theta = 2, H_{i,k-1}) = 1 - \frac{12}{14}\frac{11}{13}\frac{10}{12} = .395$ and

$p_{j,k}(1 \mid \theta = 2, H_{j,k-1}) = 1 - \frac{13}{14}\frac{12}{13}\frac{11}{12} = .214$ respectively.

To estimate the θ that best fits subjects' responses, one needs a function that maps these imputed subjective probabilities into a [1-7] number. We estimated θ for three different possible such functions:

(i) the best response to the subjective probabilities ($b^*_{i,k}$) ,

(ii) the expected minimum ($E[m_{i,k}]$), and

(iii) the seven probabilities themselves ($p_{i,k}(1|\theta,H_{i,k-1})$, $p_{i,k}(2|\theta,H_{i,k-1})\ldots p_{i,k}(7|\theta,H_{i,k-1})$).

In particular, to estimate θ we computed, for different values of θ, the R^2 from regressions in which the dependent variable was a subject's actual choice in each round, and the predictors were the three transformations of $p_{i,k}(m_{i,k}|\theta,H_{i,k-1})$ listed above. We chose as our estimated true value of θ that which obtained the highest R^2. The general specification of these regressions can be characterized by equation (2):

$$b_{i,k} = \beta_0 + \beta_1 \text{Round} + \beta_2\, b_{i,k-1} + \beta_2\, G[p_{i,k}(m|\theta,H_{i,k-1})] + \varepsilon_{i,k} \qquad (2)$$

Where

$b_{i,k}$ is the number chosen by individual *i* in round *k*, (with $b_{i,k} \in [1,7]$)

G is one of the aforementioned three functions that map $p_{i,k}(m_{i,k}|\theta,H_{i,k-1})$ into $b_{i,k}$

The specific regressions, then, corresponded to the following three versions of equation (2):

$$b_{i,k} = \beta_0 + \beta_1 \text{Round} + \beta_2\, b_{i,k-1} + \beta_3\, b^*_{i,k} + \varepsilon_{i,k} \qquad (2i)$$

$$b_{i,k} = \beta_0 + \beta_1 \text{Round} + \beta_2\, b_{i,k-1} + \beta_3\, E[\min]_{i,k} + \varepsilon_{i,k} \quad (2ii)$$

$$b_{i,k} = \beta_0 + \beta_1 \text{Round} + \beta_2\, b_{i,k-1} + \beta_{3,1}\, p_{i,k}(1|\theta,H_{i,k-1}) + \ldots\ \beta_{3,7}\, p_{i,k}\,(7|\theta,H_{i,k-1}) + \varepsilon_{i,k} \quad (2iii)$$

Figure 4 shows the R^2 for each of these specifications, obtained for different values of θ. The y-axis depicts the R^2 and the x-axis the θ associated with that R^2.

Figure 4

The R^2 curve for specification (2iii) is the highest, which is not surprising given that it uses 6 additional degrees of freedom. Somewhat surprisingly, curve (2ii) is above curve (2i), suggesting subjects are not choosing the action that provides the highest expected payoff conditional on their implied subjective probabilities.

Most importantly, consistent with the results obtained with strategy 1, the θ that maximizes R^2 of all three specifications is around 2. This means that subjects behave as if they weighted the behavior of subjects in their group twice as heavily as that of subjects in other groups.

This second approach also allowed us to estimate the negative consequences of placing a premium on experienced information. In particular, one can compute the expected payoffs that a subject would have obtained if she behaved according to θ=1 and compare this figure to her actual payoffs. According to our calculations, on average, if we allow one subject to behave

according to θ=1 while maintaining the actions of all other players constant, the expected payoffs would have been 82.43 Swedish crowns, compared to the actual average of 72.85.[6]

4. Experiment 2 – Repeated Prisoner's Dilemma

In experiment 1 we analyzed a repeated WL game, and found support for the hypothesis that people overweight experienced information. In experiment 2 we seek to replicate the results using a different game, the PD. Consistent with the notion that people place a premium on experienced information, our prediction is that subjects' decision to defect or cooperate will be more strongly influenced by the behavior of their previous counterpart, than by the behavior of other subjects.

The PD possesses two advantages over the WL game. First, because subgroups in the WL game had four players, each choosing from among seven possible actions, there is some ambiguity in terms of what the relevant information was. This is highlighted by the need for two different empirical strategies to analyze the results. In the PD game, in contrast, participants choose between only two possible actions (cooperate or defect) and groups consist of only 2 players; experienced information in this case consists unambiguously of the binary choice of previous opponents, while observed information consists of the choices of all other players.

The simpler structure of the PD game also allowed us to ask subjects to recall the outcomes of the game. We used their answers to evaluate the possibility that our results are driven by attention -- i.e. that subjects place greater weight on their own experience because they do not pay attention to the outcomes of the other groups.

[6] Very similar results are obtained regardless of how subjects are assumed to respond to $p(m|\theta)$, i.e. the greater payoffs are not arising as a result of assuming subjects best-respond, but rather, from the equal weighting of information.

4.1 Method

A total of 54 subjects, recruited at MIT, participated in the experiment in exchange for payment. There were a total of eight experimental sessions, three run with eight subjects, and five with six subjects. Subjects were assigned to a personal computer on which they read the instructions for the experiment. After everybody had finished reading the instructions, the first round begun.

The procedure closely resembled that of the WL experiment. In each round subjects chose to cooperate or defect (labeled A and B respectively). Once everyone had made a decision, all subjects saw everybody's actions and were then reassigned to a new randomly determined counterpart. Figures 5 and 6 show screenshots of the web-pages on which subjects entered their choices and observed the outcomes of each round, respectively. As shown in figure 6, subjects learned of the choices of others by seeing the relevant cells in the payoff matrices highlighted on their screens. For example figure 6 shows that in group 1, player 1 chose B, while player 2 chose A. Note that the choices of every subgroup, including the subject's own, are displayed in the same format.

** figures 5 and 6 **

Figures 5 and 6 also show the payoff structure for the game: $8 for mutual cooperation, $9 for defecting on a cooperator, $5 for cooperating with a defector, and $6 for mutual defection. Experiments lasted at least 6 rounds, one of which was randomly chosen to count.[7] At the end of all rounds, students were asked to recall the total number of times choice B had been selected

[7] Due to a problem with the PHP code, some experimental sessions ran for 8 rounds. We present the analyses including only the first 6 in order to include the same number of rounds per experimental session. The qualitative nature of the results is not affected if we include all rounds.

across all subgroups. They were also asked to recall how many times their own past counterparts had chosen B. We use their answers to these questions to see if subjects were paying as much attention to the actions of subjects in other subgroups as to the action of the individual they were directly paired with on a particular round.

4.2 Results

Table 4 shows the percentage of subjects who cooperated by round (it also presents parallel figures for experiment 3, which will be discussed in the following section). As is the case in most PD experiments, the proportion of subjects who cooperate drops as the experiment progresses. Figure 7 shows the proportions of people who cooperate in round *k*, in experiment 2, as a function of their previous action and the action of their previous opponent. There is a clear interaction between the individual's previous choice and the opponent's: previous cooperators are more likely to cooperate after interacting with another cooperator (86%) than with a previous defector (64%) but previous defectors are less likely to cooperate after interacting with a cooperator (15%) than after interacting with a defector (27%).

** Table 4**

** Figure 7 **

With this interaction in mind, we tested the hypothesis that individuals place more weight on the behavior of the player they faced by estimating the following regression:

$$\text{Own}_{i,k} = B_0 + B_1\ \text{Own}_{i,k-1} + B_2\ \%\ \text{of Cooperators}_{i,k-1} + B_3\ \%\ \text{of Cooperators}_{i,k-1} * \text{Own}_{i,k-1} + \quad (3)$$

B_4 Experienced$_{i,k-1}$ + B_5 Experienced$_{i,k-1}$ * Own$_{i,k-1}$

All variables refer to whether the relevant individual(s) cooperated in a given round, in particular:

- $Own_{i,k}$ takes the value of 1 when individual *i* cooperates in round *k*, 0 otherwise
- $\%\ of\ Cooperators_{i,k-1}$ is the percentage of subjects in other groups who cooperated in the previous round, and
- $Experienced_{i,k-1}$ takes the value of 1 when individual *i* experienced a cooperator in round *k-1*, 0 otherwise

As was the case in the previous experiment, we use clustered standard errors by subjects to correct for the fact that the data include multiple decisions by each subject.

** Table 5 **

Table 5 reports the parameter estimates for equation (3) estimated with a Probit regression. Column 1 shows that defectors who faced a cooperator were less likely to cooperate than defectors who faced a defector (B_4 = -0.427, p = .119), while cooperators who faced a cooperator were more likely to cooperate than cooperators who faced a defector (B_5 = 1.159, p = .003). The joint test that both B_4 and B_5 are 0 is rejected with p = .016. These results demonstrate *an* effect of experienced information, however, only after controlling for aggregate behavior can we interpret them as indicative of subjects weighting directly experienced information *more heavily* than observed information.

Column 2 shows that controlling for the aggregate cooperation rate, and its interaction with a subject's previous decisions, directly experienced information still matters, indeed, the size of the coefficients is virtually unchanged. Cooperators who faced a cooperator are more likely to cooperate in the next round (B_5 = 1.150, p = .004), and defectors are somewhat less so (B_4 = -0.44, p = .105).[8] More importantly, the joint test for experienced information is rejected (p = .021) while that for observed information is not (p = .542). For completeness, column 3 shows the influence of observed information, without controlling for experienced information. It is noteworthy that even without controlling for experienced information, observed information does not significantly predict subjects' choices (p = .301).

4.2.1 Attention.

As mentioned earlier, at the end of each experimental session of the PD game, subjects were asked to recollect:

(i) how many of their partners cooperated, and

(ii) how many of all subjects had cooperated

Their answers to these questions are useful for discriminating between two possible mechanisms behind our results. One possibility is that subjects collect all the relevant information and at the time of making a decision place a premium on the information they experienced. Alternatively, subjects may simply not pay attention to the outcomes of other groups and, since they don't know what others do, their behavior is not affected by it. Both of these explanations violate standard assumptions about how people use information, albeit through different mechanisms.

[8] As was mentioned earlier, some experimental sessions run for 8 instead of 6 rounds. If rounds 7 and 8 of those session are included in the regression the results remain qualitatively unchanged (B_4 = -.183, p=.012; B_5=.396, p<.001).

To evaluate whether subjects paid attention to the outcomes of other groups we conducted two simple tests. The first consists of estimating the following regression:

$$Estimate_i = \gamma_0 + \gamma_1 \, Observed_j + \gamma_2 \, Own_i + \gamma_3 \, Experienced_i + \varepsilon_{i,} \qquad (4)$$

where:

$Estimate_i$ is the estimated total number of cooperators by subject *i* (i.e. their answer to question ii). [9]

$Observed_j$ is the actual number of observed cooperators (i.e. not faced by the subject) in session *j*.

Own_i is the actual number of times subject *i* cooperated

$Experienced_i$ is the actual number of times subject *i*'s counterparts cooperated

If subjects were perfectly accurate in their recollection of what others did, γ_1 in equation (4) would equal 1, since for every subject (not-paired with them) that defected, their total recollection of defection should increase by 1. If, on the other extreme, subjects completely neglected what others did, their total estimate of defection should be uncorrelated with the actual defection rate and γ_1 would equal 0.

The parameter estimates for γ_1, γ_2 and γ_3 were 0.714, 0.742 and 0.921 (SEs = .414, .334 and .318, respectively). The R^2 was 37.2%. Consistent with the hypothesis that participants *do* pay attention to the actions of others, the null that $\gamma_1 = 1$ cannot be rejected ($p > .4$) but the null that $\gamma_1 = 0$ can ($p < .1$). In other words, the null that participants do not pay attention to the actions of other players *is* rejected, and the null that they are perfectly calibrated *is not*.

[9] Note that each subject makes only one prediction because they participate in only one experimental session.

The second test we employed to assess whether subjects paid attention to the choices of subjects with whom they did not interact, consists of comparing the accuracy of their answers to questions (i) and (ii). If subjects did not pay attention to the actions of others we would expect their estimates of aggregate cooperation rates to be further off than their estimates of cooperation rates from their previous partners. We measured accuracy by calculating the average absolute error, in percentage terms. The average absolute error for question (i) was 11.5%, only slightly lower than the 12.3% for question (ii). To assess whether these two measures were statistically different, for each subject we subtracted the absolute error for (i) from the absolute error for (ii) and tested whether this difference was significant (note that this difference is computed at the subject level controlling therefore for subject-specific differences in overall recall). The test fails to reject the null that subjects are equally accurate for answering (i) and (ii), $p = .940$.

Since subjects are quite accurate in the recollection of aggregate cooperation rates, it seems that the impact of experienced information on behavior is not the result of simply ignoring non-experienced information, but rather, of placing a premium on information that has been experienced. Directionally speaking, however, the results presented above do suggest that subjects were less calibrated with the actions of subjects they observed vs. those of which they experienced and, hence, it remains a possibility that the premium placed on experienced information is caused by differential attention.

In our third and final experiment we further assessed the potential role of attention by asking subjects to write down the actions of all players in all groups after each round, guaranteeing that they had attended to them.

5. Experiment 3– Repeated Prisoner's Dilemma with subjects writing down all actions

5.1 Design

The design of this experiment was identical to that of study 2 with the following two differences: all sessions were run with 6 subjects at a time, and subjects were asked to write down the actions of all 6 players in the 3 dyads after every round. By asking them to write the actions down, we not only guarantee that they looked at them, but that they paid enough attention to reproduce them on a piece of paper.

5.2 Results

Average cooperation by round is reported in Table 4 (together with the parallel values from experiment 2). Cooperation rates were extremely similar across the two experiments. To estimate the impact of experienced and observed information we conducted the same regressions used for experiment 2. The results are presented in Table 6. Again, they are highly consistent with the findings from study 2 (Table 5). Most importantly, column (2) shows that experienced information (i.e. the action of the previous opponent) has a significant impact on a subject's likelihood to cooperate (p = .0004) but observed information does not (p = .2667), even though subjects did pay attention to it, since they were able to accurately write them down.

The fact that we are able to replicate our findings when attention is no longer a plausible explanation strongly suggests that experienced information is actually weighted more heavily than observed information, rather than simply attended to more carefully.

6. Conclusions

In this paper we present results from three experiments in which subgroups of subjects played several rounds of the Weak-link (WL) and the Prisoner's Dilemma (PD) game. Subjects

were given information about the decisions of subjects both in their own sub-group (experienced information) and in the other subgroups (observed information). Because subgroups were randomly determined in each round, the actions of all subjects in the previous round were equally relevant and hence should have received equal weight. In contrast to this normative prediction, in both games we found that subjects placed significantly more weight on the actions of subjects they had experienced than on those they had simply observed.

In the WL game, subjects did respond to the actions of subjects in groups in which they did not participate, but they responded roughly twice as much to subjects in their own previous subgroup. In the PD, we found that subjects responded (at least significantly) only to the actions of their previous counterpart. Suggesting that experienced information is weighted, rather than simply attended, more than observed information, in experiment 2 they were able to accurately recall observed actions and in experiment 3 they wrote down the actions of all players in all dyads.

The fact that different sources of information are weighed in non-optimal ways has important implications for economics. In many situation agents both observe and experience information. For example, people receive information about the returns of stocks they invest in and about stocks they do not invest in; they learn the consequences of diseases they experience and of those they learn about from friends and relatives; they experience the consequences of risks they take personally and observe the effects of risk-taking by other people. If they systematically overweight directly experienced information, their decisions are likely to systematically deviate from standard models of choice which assume they do not.

For example, although there has been ample experimental and theoretical work on the *existence* of social preferences, this research provides little guidance in predicting how people

decide what specific cause to provide assistance to in the real world. Anecdotal evidence suggests that personal experience plays a large role: people appear much more likely to adhere to social causes they have experienced directly. Although the anecdotal evidence could be caused by confounds between social causes people experience and do not experience, recent experimental work by Small and Simonsohn (2005), inspired in part by the results from this paper, provides evidence consistent with the notion that personal experience causally shapes preferences over the outcomes of others. They randomly assigned subjects to befriend a "victim" (fellow subject who lost $10) or a non-victim, and found that subjects matched with victims acted more generously towards *other* victims.

Another important instance in which a premium placed on experienced information could have important economic consequences is when market conditions change. If people react more strongly to experienced than to observed information, they may be slow to react to market changes even if they are informed about these changes, as they may need to experience them before being willing to act. For example, a seminal paper by Lucas (1973) proposes that inflation can affect output if producers confuse changes in the aggregate level of prices with changes in relative prices. One mechanism for this type of error (particularly in light of the fact that inflation data is freely available) is that producers weigh price changes they experience (those of their own product) too much, and those they learn about indirectly too little.

Consistent with such suboptimal weighting of information, (Simonsohn and Loewenstein, 2006) and (Simonsohn, 2006) find that movers between cities base their decision of how much to spend on rent, and how long to commute in a city they have just moved to, respectively, at least partially on the conditions they experienced on the previous city. When people move again

within their new city, however, after having experienced local prices, they readjust fully undoing the impact that previous conditions had.

As a final application, consider the common tendency of workers to load their retirement portfolios with stocks of their own company – a strategy that makes little sense in terms of risk diversification given the likely coincidence of job loss and poor performance by the company. Choi, Laibson and Madrian (2005) studied the impact of the debacles at Enron, Global Crossing, and Kmart on the investment of US workers at other firms in their employers' stock, and found it to be negligible.

The results presented in this paper show that people tend to weigh experienced information more heavily than observed experience. The recent evidence about how people choose charities, and the results from our PD experiments suggest that probably the premium given to experienced information is mediated by preferences.

This impact on preferences is probably mediated by emotions. Indeed, we have difficulty imagining any other plausible explanation for such a differential reaction when informational content is held constant. Future research could productively examine more directly the role of emotions in producing differences in learning from experienced and observed information.

Table 1. Payoff structre for weak-link experiment

	Minimum choice of other players						
Player's Choice	**7**	**6**	**5**	**4**	**3**	**2**	**1**
7	130	110	90	70	50	30	10
6	120	120	100	80	60	40	20
5	110	110	110	90	70	50	30
4	100	100	100	100	80	60	40
3	90	90	90	90	90	70	50
2	80	80	80	80	80	80	60
1	70	70	70	70	70	70	70

Payoffs are in Swedish Crowns. When the experiment was conducted 1US$ = 8 SEK aprox.

Table 2. Average group-minimum by round

Round	Average	SD
1	3.72	0.45
2	3.33	0.91
3	3.22	0.60
4	2.83	0.74
5	2.94	0.92
6	3.00	1.15
7	2.44	0.90
8	2.89	1.18
9	2.78	1.00
10	2.61	0.98

Table 3. Weak-link game. Choice in *k* as a function of minima in round *k-1* (Ordered Probit)[a]

Dependent Variable: Own choice (1-7) in round *k*

		(1) Only Experienced	(2) Experienced and Observed	(3) Only Observed
Own choice in previuos round	(α_1)	0.402***	0.382***	0.374***
		(0.055)	*(0.054)*	*(0.054)*
Minimum in Group 1 in previous round	(α_2)	--	0.040	0.062**
		--	*(0.034)*	*(0.031)*
Minimum in Group 2 in previous round	(α_3)	--	0.102***	0.122***
		--	*(0.033)*	*(0.028)*
Minimum in Group 3 in previous round	(α_4)	--	0.027	0.051*
		--	*(0.033)*	*(0.029)*
Minimum in own-group in previous round	(α_5)	0.119***	0.066**	--
		(0.023)	*(0.032)*	--
Round	(α_6)	-0.021	-0.016	-0.018
		(0.013)	*(0.014)*	*(0.013)*
Was subject in group 1 in previous round? (1 or 0)	(α_7)	0.113	0.105	--
		(0.097)	*(0.097)*	--
Was subject in group 2 in previous round? (1 or 0)	(α_8)	0.029	0.032	--
		(0.106)	*(0.105)*	--
Number of observations		648	648	648

Clustered (by subject) standard errors reported in parenthesis below parameter estimates

*,**,*** significant at the 10%,5% and 1% level respectively

[a] Six intercepts are estimated for Ordered Probit, but not reported in table

Table 4. Cooperation Rates by Round in Experiments 2 and 3

Round	Experiment 2	Experiment 3
1	61.1%	63.3%
2	55.5%	51.7%
3	51.8%	55.0%
4	44.4%	55.0%
5	44.4%	40.0%
6	43.7%	36.6%

Table 5. Regression analysis of cooperation in round *k for Experiment 2 (Probit)*

Dependent variable: $Cooperated_k$ (1 if self cooperated in round *k* , 0 otherwise)

		(1) Only Experienced[a]	(2) Experienced and Observed[b]	(3) Only Observed
Intercept	(B_0)	-0.612**	-0.963**	-1.071**
		(0.295)	*(0.421)*	*(0.429)*
$Cooperated_{k-1}$ (1 if yes, 0 if no)	(B1)	0.961***	1.155***	1.415***
		(0.247)	*(0.447)*	*(0.448)*
% of $cooperators_{k-1}$	(B_2)	--	0.636	0.562
		--	*(0.623)*	*(0.629)*
% of $cooperators_{k-1}$ * $Cooperated_{k-1}$	(B_3)	--	-0.418	0.064
		--	*(0.794)*	*(0.759)*
Experienced $Cooperator_{,k-1}$ (1 if yes, 0 if no)	(B_4)	-0.427	-0.444	--
		(0.274)	*(0.274)*	--
Experienced $Cooperator_{k-1}$*$Cooperated_{k-1}$	(B_5)	1.159***	1.150***	--
		(0.391)	*(0.400)*	--
Round	(B_6)	0.001	0.011	-0.013
		(0.050)	*(0.051)*	*(0.049)*
Joint Tests				
Experienced information (B_4 =0 and B_5=0)		*p=.016*	*p=.021*	--
Observed information (B_2 =0 and B_3=0)		--	*p=.542*	*p=.301*
Number of observations		270	270	270

p-values for two tailed tests are presented below parameter estimates (obtained with *Clustered Standard Errors*).

[a] Experienced information consists of the action of the player with whom the subject played on the previous round (1 if cooperated, 0 otherwise)

[b] Observed information consists of the percentage players in *other* groups who cooperated in previous round

*,**,*** significant at the 1%,5% and 10% level respectively

Table 6. Regression analysis of cooperation in round *k for Experiment 3 (Probit)*

Dependent variable: $Cooperated_k$ (1 if self cooperated in round *k* , 0 otherwise)

		(1) Only Experienced[a]	(2) Experienced and Observed[b]	(3) Only Observed
Intercept	(B_0)	-0.070	-0.505	-0.594*
		(0.270)	*(0.377)*	*(0.355)*
$Cooperated_{k-1}$ (1 if yes, 0 if no)	(B1)	0.330	0.741*	1.297***
		(0.237)	*(0.398)*	*(0.361)*
% of $cooperators_{k-1}$	(B_2)	--	0.801	0.759
		--	*(0.494)*	*(0.479)*
% of $cooperators_{k-1}$ * $Cooperated_{k-1}$	(B_3)	--	-0.778	-0.816
		--	*(0.650)*	*(0.610)*
Experienced $Cooperator_{,k-1}$ (1 if yes, 0 if no)	(B_4)	-0.160	-0.198	--
		(0.230)	*(0.238)*	--
Experienced $Cooperator_{k-1}$*$Cooperated_{k-1}$	(B_5)	0.927***	0.965***	--
		(0.311)	*(0.318)*	--
Round	(B_6)	-0.095*	-0.092*	-0.085
		(0.049)	*(0.054)*	*(0.053)*
Joint Tests				
Experienced information (B_4 =0 and B_5=0)		*p=.0004*	*p=.0004*	--
Observed information (B_2 =0 and B_3=0)		--	*p=.2667*	*p=.2836*
Number of observations		300	300	300

p-values for two tailed tests are presented below parameter estimates (obtained with *Clustered Standard Errors*).

[a] Experienced information consists of the action of the player with whom the subject played on the previous round (1 if cooperated, 0 otherwise)

[b] Observed information consists of the percentage players in *other* groups who cooperated in previous round

*,**,*** significant at the 1%,5% and 10% level respectively

Figure 1. Screen-shot of page where subjects make their [1-7] decision in each of 10 rounds of the Weak-link experiment

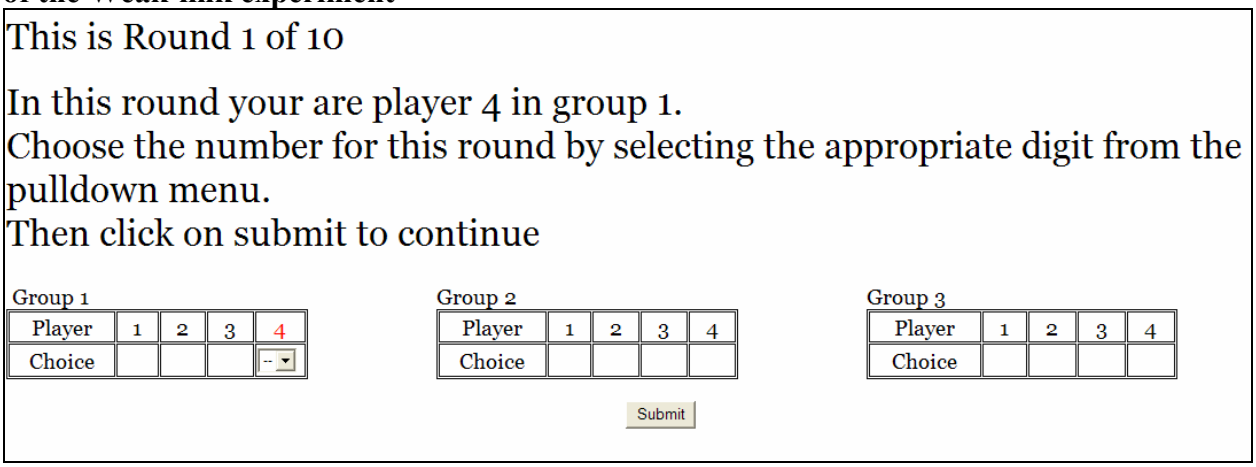

Figure 2. Screen-shot of page where subjects observe the decisions of all 12 players in each round of the Weak-link experiment.

Results for round 1

In this round your are player 4 in group 1.
The tables below show the results for all players in all groups.
You can see your payoffs in this round by looking at the results in your table.
Your payoffs are marked in red.

Click on the button to continue to the next round

Group 1

Player	1	2	3	4
Choice	7	6	4	7
Payoff	70	80	100	70

Group 2

Player	1	2	3	4
Choice	7	4	7	4
Payoff	70	100	70	100

Group 3

Player	1	2	3	4
Choice	6	4	6	5
Payoff	80	100	80	90

Go on to next round

Figure 3. Distribution of chosen numbers in WL games, rounds 1-10.

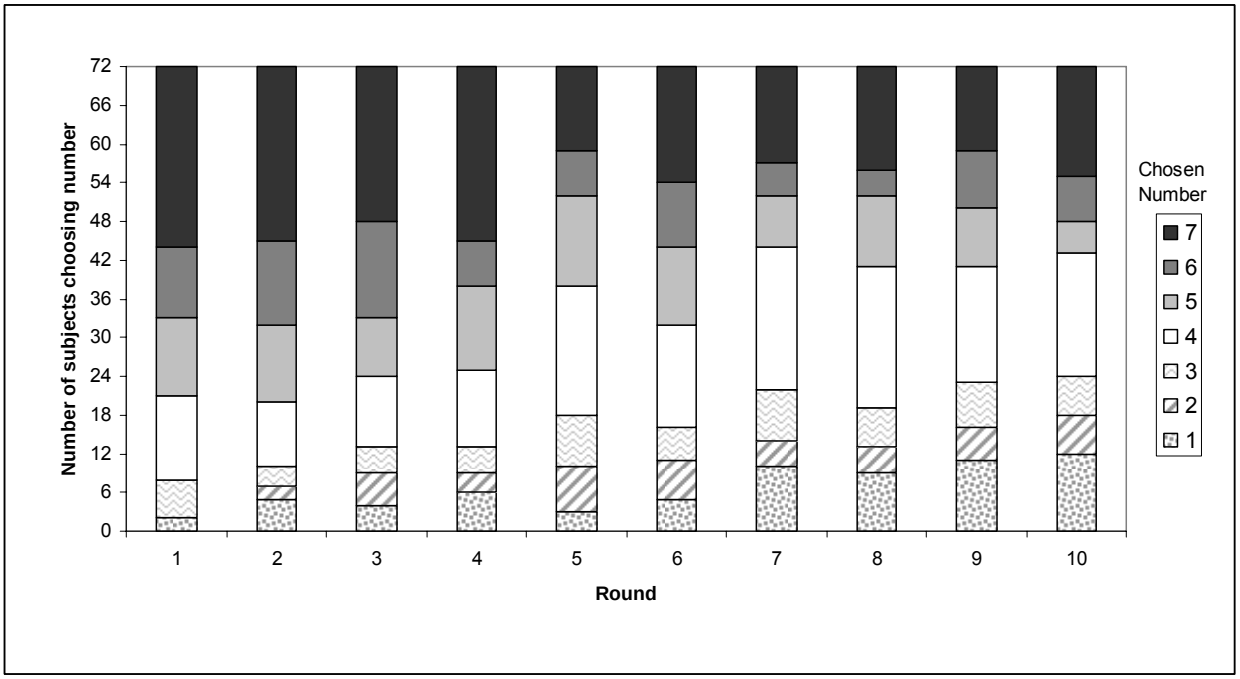

Figure 4. R^2's obtained from regressing subjects' actual choices on predicted choices from three functions of imputed subjective probabilities, for the minimum in the next group, as a function of the relative weighting of experienced and observed information (θ).

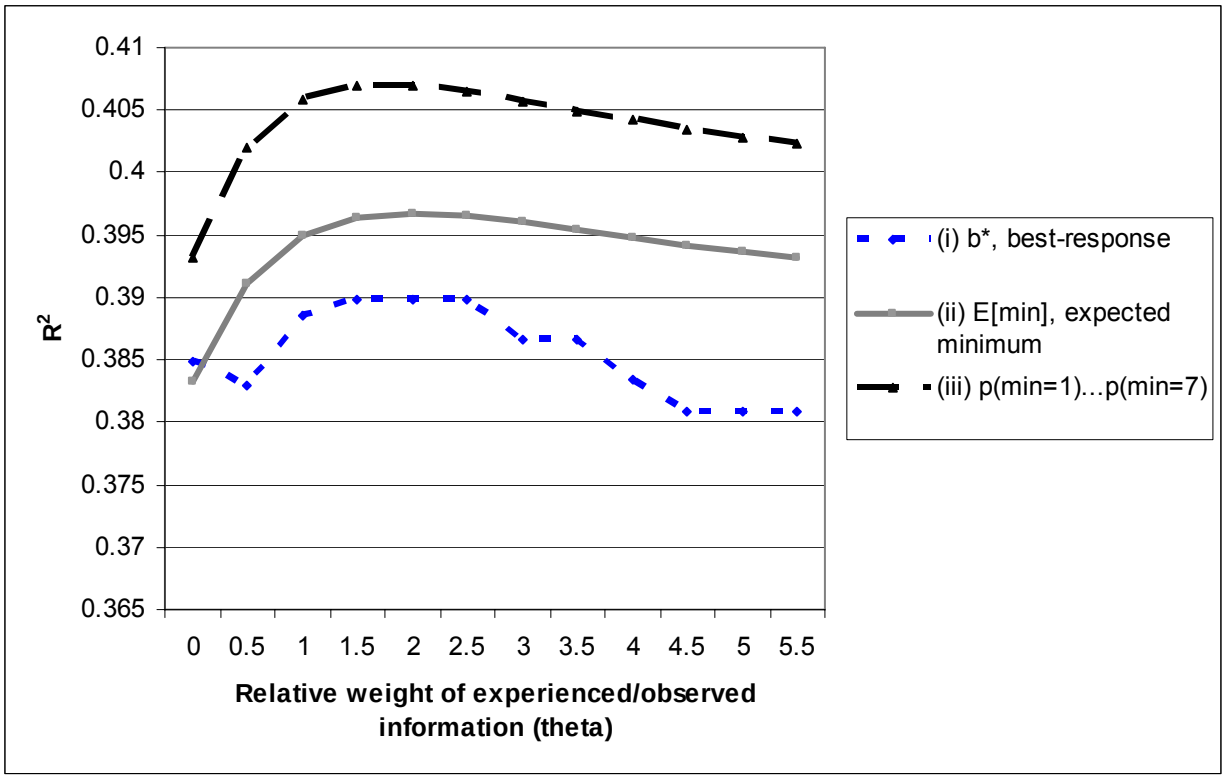

Figure 5. Screen-shot of webpage where subjects decide in the PD game.

Your ID is 1. This is Round 1. In this round you are Player 2 in Group 3.
Select A or B in the appropriate group by clicking on a button.

Group 1

		PLAYER 1	
		A	B
PLAYER 2	A	Player 1: $8 Player 2: $8	Player 1: $9 Player 2: $5
	B	Player 1: $5 Player 2: $9	Player 1: $6 Player 2: $6

Group 2

		PLAYER 1	
		A	B
PLAYER 2	A	Player 1: $8 Player 2: $8	Player 1: $9 Player 2: $5
	B	Player 1: $5 Player 2: $9	Player 1: $6 Player 2: $6

Group 3 (Yours)

		PLAYER 1	
		A	B
PLAY (You) ER 2	A	Player 1: $8 Player 2: $8	Player 1: $9 Player 2: $5
	B	Player 1: $5 Player 2: $9	Player 1: $6 Player 2: $6

Group 4

		PLAYER 1	
		A	B
PLAYER 2	A	Player 1: $8 Player 2: $8	Player 1: $9 Player 2: $5
	B	Player 1: $5 Player 2: $9	Player 1: $6 Player 2: $6

Figure 6. Screenshot of webpage where subjects see the outcomes of the previous round.

Your ID is 1. This is Round 1.
In this round you are Player 2 in Group 3.
The payoffs and outcomes are shown below. Your payoff is in bold text.

Group 1

		PLAYER 1	
		A	B
PLAYER 2	A	Player 1: $8 Player 2: $8	Player 1: $9 Player 2: $5
	B	Player 1: $5 Player 2: $9	Player 1: $6 Player 2: $6

Group 2

		PLAYER 1	
		A	B
PLAYER 2	A	Player 1: $8 Player 2: $8	Player 1: $9 Player 2: $5
	B	Player 1: $5 Player 2: $9	Player 1: $6 Player 2: $6

Group 3 (Your group)

		PLAYER 1	
		A	B
PLAY (You) ER 2	A	Player 1: $8 **Player 2: $8**	Player 1: $9 Player 2: $5
	B	Player 1: $5 Player 2: $9	Player 1: $6 Player 2: $6

Group 4

		PLAYER 1	
		A	B
PLAYER 2	A	Player 1: $8 Player 2: $8	Player 1: $9 Player 2: $5
	B	Player 1: $5 Player 2: $9	Player 1: $6 Player 2: $6

Figure 7. The impact of outcome in k-1 on cooperation rate in k

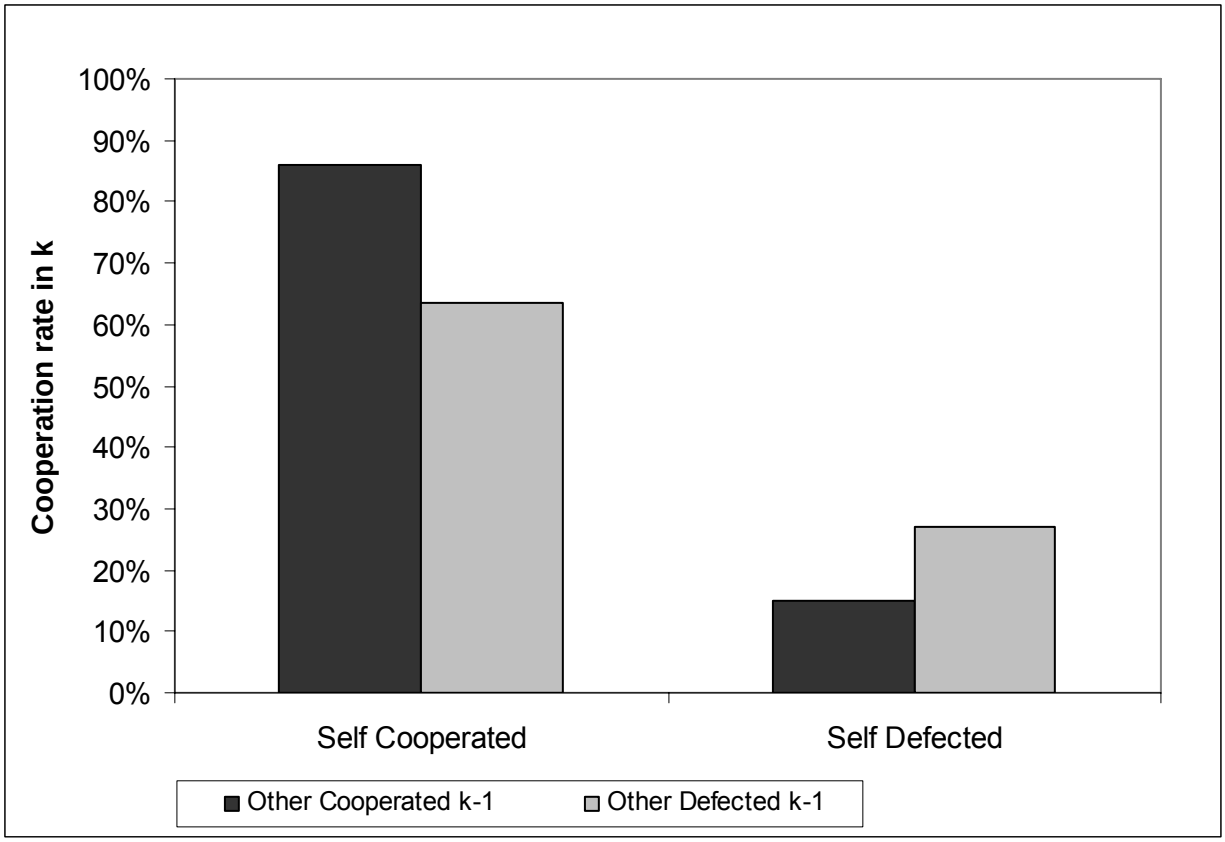

References

Ballinger, T. P., et al., 2003. Precautionary saving and social learning across generations: An experiment. Economic Journal. 113**,** 920-947.

Barron, G., Erev, I., 2003. Small Feedback-based Decisions and Their Limited Correspondence to Description-based Decisions. Journal of Behavioral Decision Making. 16**,** 215-233.

Braver, S. L., Rohrer, V., 1978. Superiority of Vicarious over Direct Experience in Interpersonal Conflict Resolution. Journal of Conflict Resolution. 22**,** 143-155.

Camerer, C. F., 2003. Behavioral Game Theory: Experiments in Strategic Interaction. Princeton University Press, Princeton, New Jersey.

Camerer, C. F., Cachon, G. P., 1996. Loss-Avoidance and Forward Induction in Experimental Coordination Games. Quarterly Journal of Economics. 111**,** 165-94.

Camerer, C. F., Ho, T. H., 1999. Experience-weighted attraction learning in normal form games. Econometrica. 67**,** 827-874.

Camerer, C. F., et al., 2002. Sophisticated experience-weighted attraction learning and strategic teaching in repeated games. Journal of Economic Theory. 104**,** 137-188.

Camerer, C. F., Knez, M., 2000. Increasing cooperation in social dilemmas by establishing a precedent of efficiency in coordination games. Organizational Behavior and Human Decision Processes. 82**,** 194-216.

Camerer, C. F., et al., 2005. Neuroeconomics: How neuroscience can inform economics. Journal of Economic Literature. 43**,** 9-64.

Choi, J. J., et al., 2005. Are Empowerment and Education Enough? Underdiversification in 401(k) Plans. Brookings Papers on Economic Activity. 151-198.

DuBow, F., et al., 1979. Reactions to crime: A critical review of the literature. US Government Printing Office, Washington, DC.

Duffy, J., Feltovich, N. J., 1999. Does observation of others affect learning in strategic environments? An experimental study. International Journal of Game Theory. 28**,** 131-152.

Fudenberg, D., Kreps, D. M., 1993. Learning Mixed Equilibria. Games and Economic Behavior. 5**,** 320-367.

Halpern-Felsher, B. L., et al., 2001. The role of behavioral experience in judging risks. Health Psychology. 20**,** 120-126.

Hertwig, R., et al., 2004. Decisions from experience and the effect of rare events in risky choice. Psychological Science. 15**,** 534-539.

Hough, M., 1985. The Impact of Victimization: Findings from the British Crime Survey. Victimology. 10**,** 488-497.

Kunreuther, H., 1978. Disaster Insurance Protection: Public Policy Lessons. Wiley, New York.

Lind, E. A., et al., 1998. The social construction of injustice: Fairness judgments in response to own and others unfair treatment by authorities. Organizational Behavior and Human Decision Processes. 75**,** 1-22.

Lindell, M. K., Perry, R. W., 2000. Household adjustment to earthquake hazard - A review of research. Environment and Behavior. 32**,** 461-501.

Lucas, R. E., 1973. Some International Evidence on Output-Inflation Tradeoffs. American Economic Review. 63**,** 326-334.

Merlo, A., Schotter, A., 1999. A surprise-quiz view of learning in economic experiments. Games and Economic Behavior. 28**,** 25-54.

Nisbett, R. E., Ross, L., 1980. Human Inference. Prentice-Hall, Inc., Englewood Cliffs, NJ.

Olsson, A., Phelps, E. A., 2004. Learned fear of "unseen" faces after Pavlovian, observational, and instructed fear. Psychological Science. 15**,** 822-828.

Parker, S. D., et al., 1980. Natural Disaster, Perceived Control, and Attributions to Fate. Personality and Social Psychology Bulletin. 6**,** 454-459.

Schiff, M., 1977. Hazard Adjustment, Locus of Control, and Sensation Seeking - Some Null Findings. Environment and Behavior. 9**,** 233-254.

Schotter, A., Merlo, A., 2003. Learning by Not Doing. Games and Economic Behavior. 42**,** 116-136.

Simonsohn, U., 2006. New Yorkers Commute More Everywhere: Contrast Effects in the Field. The Review of Economics and Statistics. LXXXVIII**,** 1-9.

Simonsohn, U., Loewenstein, G., 2006. Mistake #37: The impact of previously observed prices on housing demand. The Economic Journal. 116**,** 175-199.

Skogan, W. G., 1987. The Impact of Victimization on Fear. Crime & Delinquency. 33**,** 135-154.

Small, D. A., Simonsohn, U., Friends of Victims: The Impact of Personal Relationships with Victims on Social Preferences for Other Victims. SSRN Working Paper, 2005.

Svenson, O., et al., 1985. Perceived Driving Safety and Seatbelt Usage. Accident Analysis and Prevention. 17**,** 119-133.

Tyler, T. R., 1980. Impact of Directly and Indirectly Experienced Events - the Origin of Crime-Related Judgments and Behaviors. Journal of Personality and Social Psychology. 39**,** 13-28.

Van Huyck, J., et al., 1990. Tacit Coordination Games, Strategic Uncertainty and Coordination Failure. American Economic Review. 80**,** 234-248.

Weber, R., 2006. Managing growth to achieve efficient coordination in large groups. American Economic Review. 96**,** 114-126.

Weinstein, N. D., 1987. Unrealistic Optimism about illness susceptibility: conclusions from a community-wide sample. Journal of Behavioral Medicine. 10**,** 481-500.

Weinstein, N. D., 1989. Effects of Personal-Experience on Self-Protective Behavior. Psychological Bulletin. 105**,** 31-50.

Yechiam, E., et al., 2005. The role of personal experience in contributing to different patterns of response to rare terrorist attacks. Journal of Conflict Resolution. 49**,** 430-439.

Appendix 1 - Instructions for Weak-link experiment.

Note to the reader:
The experiments and all instructions were originally presented to subjects on a webpage.
The Weak-link experiment was run in Sweden. The instructions that follow are the original English instructions which were translated into Swedish before the experiment took place.
Comments which were not present in the original instructions but which will help the reader better understand them are included inside <parenthesis>

Page 1

This is an experiment about decision making. The decisions that you and other people you interact with make will determine how much money you (and they) make in this experiment. This money will be paid at the end of the experiment in private and in cash.

It is important that during the experiment you remain SILENT. If you have any questions, or need assistance of any kind, RAISE YOUR HAND. DO NOT SPEAK. One of the experiment administrators will come to you and you may whisper your question to him. If you talk, laugh, exclaim out loud, etc., you will be asked to leave and you will not be paid. We expect and appreciate your cooperation.

Continue <this was a link to the next page>

Page 2

You are about to participate in an experiment with 11 other people. The experiment will consist of several rounds. In each round, every one of the participants is randomly assigned to one of groups.

In each round, each group of players will play the game described below:

Each player will pick a number between 1 and 7 by clicking on a button on the table shown below. Each player will receive an amount of money determined by their number and the lowest *other* number chosen, as described by the table below. You will not be able to see anyone else's choices until everybody has made their choice.

<here subjects saw the payoff table which is table 1 in this paper>

Continue *<this was a link to the next page>*

Page 3

Before you participate in the actual experiment, please solve the following two hypothetical examples by indicating what the payoff would be for each of the participants. Once you type in the correct answers to both, you will be able to begin.

Note that all participants will need to solve these (same) examples before the session begins, so you may have to wait a few minutes.

You should have a sheet of paper with the payoffs table that you saw in the previous page, you will need it to figure out the appropriate payments. If you encounter difficulties please raise your hand and the experimenter will assist you.

Example 1

	Player 1	**Player 2**	**Player 3**	**Player 4**
Chosen number	**6**	**6**	**6**	**6**
Payoff				

Example 2

	Player 1	**Player 2**	**Player 3**	**Player 4**
Chosen number	3	5	4	6
Payoff				

ANSWER

<when subjects clicked on the "answer" button, their answers were checked automatically. If they were wrong, they were asked to try again until they answered them correctly. When all subjects succeeded in answering these questions, the experiment begun. Figures 1 and 3 depict the pages they saw as they participated in the experiment.>

Appendix 2. - Instructions for Prisoner's Dilemma Experiment (originally presented to subjects on a webpage).

Page 1

This is an experiment about decision making. The decisions that you and other people you interact with make will determine how much money you (and they) make in this experiment. This money will be paid at the end of the experiment in private and in cash.

It is important that during the experiment you remain SILENT. If you have any questions, or need assistance of any kind, RAISE YOUR HAND. DO NOT SPEAK. One of the experiment administrators will come to you and you may whisper your question to him. If you talk, laugh, exclaim out loud, etc., you will be asked to leave and you will not be paid. We expect and appreciate your cooperation.

Continue <*this was a link to the next page*>

Page 2

You are about to participate in an experiment with <K> other people. The experiment will consist of several rounds. In each round every one of the <K+1> participants is randomly assigned to one of <(K+1)/2> groups.

Each group of 2 players will play the game that's shown below. If both players choose A then they both get $10, if one of them chooses A and the other B, the one choosing A gets $3 and the one choosing B gets $12 and if they both choose B they both get $5.

Continue <*this was a link to the next page*>

Page 3

Once everybody has played, everybody will see what everybody else did. And then a new round will begin.

In every round participant will be re-located to new groups.

After all rounds have been played, one of them will be chosen at random to count for real money.

If you have any questions please raise your hand and wait quietly for an administrator to assist you. Otherwise click on the button below and wait for everybody else to be ready.

Continue <*this was a link to the next page. When all subjects had clicked it the experiment begun. Figures 4 and 5 replicate the pages they saw as they participated in the experiment.*>