ISO Probability Matching

Charles A. Holt[*]

University of Virginia
March 1992

I. Introduction

This paper reports an experiment in which subjects were asked to predict which of two random events would occur in a sequence of Bernoulli trials. In some treatments, financial rewards and/or penalties were used to motivate subject's choices, and in another treatment, each subject was simply told to "do your best". The probability of the more likely event was .75, and therefore, the optimal decision is to predict the more likely event every time, at least in the presence of financial rewards. In the absence of financial rewards, psychologists have observed a tendency for subjects to engage in "probability matching", i.e. predicting the more likely event about as often as it occurs on average, e.g., 3/4 of the time. About 30 years ago, Sidney Siegel and his coauthors replicated the probability matching results that had been reported in the psychology literature in the 1940's and 1950's. Then he showed the tendency for the forecast proportion to match the actual proportion was greatly diminished in the presence of financial rewards. Of the two reward structures used, the payment of 5 cents for correct predictions was less effective than a 5-cent reward for correct predictions together with a 5-cent penalty for incorrect predictions.

One obvious question raised by Siegel's results is whether it is the presence of the penalty per se or the fact that the financial incentive is doubled when the penalty is added. One way to address this issue would be to compare the effects of using a 20 cent reward with that of using a 10-cent reward together with a 10-cent penalty. A second issue arises from the Siegel setup in which the experimenter was seated across from the subject, behind a screen. This setup would seem to suggest that the subject was somehow playing a game with the experimenter. As noted by Simon (1957, p.278):

[*] Department of Economics, Rouss Hall, University of Virginia, Charlottesville, VA 22903.

> To the experimenter who *knows* that the rewards attached to the two behaviors A1 and A2 are random, with constant probabilities, it appears unreasonable that the subject should not learn to behave in such a way as to maximize this expected gain --always to choose A1. To the subject, who *perceives* the situation as one in which the probabilities may change, and who is more intent in outwitting the experimenter (or "nature") than in maximizing expected gain, rationality is something quite different.

This comment suggests using more neutral procedures, at least removing the experimenter from the scene. As Siegel did not tell subjects the probability of the more likely event, Simon's comment also suggests that providing more information about the random process may increase the rationality of forecasts, especially in the no-payoff treatments.

This paper reports a new forecast experiment in which all procedures are controlled by a computer program. Probability matching is not observed in any of the treatments, with or without financial incentives, even when the instructions are adapted directly from those used by Siegel. The care with which the random elements are explained does seem to affect the rationality of forecasts in the first 40-50 trials, however.

The paper is organized as follows. Section II contains a brief discussion of rational forecasting of a Bernoulli event. Section III reviews some early evidence on probability matching and the effects of financial incentives. Procedures for the new experiment are described in section IV, and the data are summarized in the final section.

II. Optimal Forecasting[1]

Consider a choice between 2 alternatives, under the condition that the “payoffs” for these choices are determined by an unknown event, E_1 or E_2. Payoffs for each event are summarized in table 1. Suppose further that the probabilities of the two events are fixed, sum to one, and that the outcome at each stage is independent of previous outcomes. (Such events are called independent Bernoulli trials.) The specific problem for the subject, then, is to predict which event will occur; the predictions of events E_1 and E_2 are denoted by D_1 and D_2 respectively. A

[1] The discussion in this section and the next is based to some extent on the presentation in Holt and Davis (1992), chapter 2.

correct prediction yields a "reward" of \$*R*, and an incorrect predictions yields a "penalty" payoff of -\$*L*. Payoffs are set so that \$*R* exceeds -\$*L*.

Table 1. A Prediction Problem

	Event E_1	Event E_2
prediction D_1	\$*R*	-\$*L*
prediction D_2	-\$*L*	\$*R*

For purposes of experimentation, this prediction setup has the unusual and very desirable property that the optimal decision is independent of a subject's aversion or preference for risk. Let p denote the subject's subjective probability that the event will be E_1, and let $U(.)$ denote the subject's von Neumann-Morgenstern utility function. Then the expected utility for predicting E_1 is $pU(R) + (1-p)U(-L)$, and the expected utility for predicting E_2 is $pU(-L) + (1-p)U(R)$. It follows from simple algebra that the expected utility for predicting event 1 is greater if and only if

(1) $$(2p - 1)[U(R) - U(-L)] > 0.$$

Since the reward is greater than the penalty, the term in square brackets on the left side of (1) is positive, and it follows that the optimal forecast is event 1 if and only if $p > 1/2$. The intuition for why the best decision is independent of risk preferences is that there are only two possible earnings levels that are determined by the symmetric payoff table 1. With only two earnings levels, the curvature of the utility function that determines risk preference is irrelevant. Therefore, this decision problem is one way of using financial incentives to elicit information about which state is the one that a subject thinks is more likely.

III. What Siegel and Coauthors Found

A classic experiment by Siegel and Goldstein (1959) was patterned along the lines of table 1. During the 20 years prior to this paper, psychologists had been doing experiments in which subjects were asked to predict which of two events had occurred, without financial incentives of the type present in this table. Subjects, instead, were told to "do your best to predict correctly". In a typical setup, one of the events would be more likely, e.g., the probability p of event E_1 could be .75. Although subjects were not given the probabilities, the more likely even would quickly become apparent. Despite the fact that it is optimal to predict the more likely event D_1 at every stage, psychologists had observed that the proportion of times that the more likely event is predicted approximately *matched* the probability of this event. Like other psychologists, Siegel and Goldstein also observed probability matching under a no-incentive treatment. Twelve subjects each made a series of 100 predictions without financial rewards for correct guesses or penalties for errors (e.g., \$$R$ = \$$L$ = 0). The proportion of times that the more likely outcome was predicted by individual subjects ranged from .60 to .80, with an average across subjects of .70 (in trials 80-100). When 4 of the 12 subjects were brought back a week later and allowed to continue for 200 additional trials, the average across subjects ended up being exactly equal to .75, the probability of the more likely event.

Siegel was an experimental psychologist who had seen the literature on probability matching. He suspected that the curious matching behavior was caused by boredom; participants might try to overcome the tedium of making repeated binary decisions by trying to out-guess the randomizing device. Absent financial incentives, guessing of this type is costless. Siegel reasoned that the effects of boredom could be decreased by providing subjects with financial rewards and/or penalties for their decisions. This hypothesis was evaluated by comparing results of the above "no-reward" treatment with results in "reward-only" and "reward/penalty" treatments. In the reward-only treatment, participants were paid 5 cents for each correct prediction, but they paid no penalty for an incorrect prediction (e.g., \$$R$ = 5 and \$$L$ = 0). In the reward/penalty treatment, participants both received 5-cent rewards for correct guesses and paid 5-cent penalties for errors (e.g., \$$R$ = 5 and \$$L$ = 5).[2] Procedures for these latter treatments paralleled those used in the no-reward treatment: In each case, 12 subjects made predictions in a series of 100 trials,

[2] Subjects in all treatments were given an initial cash balance in order to reduce the chances of bankruptcy in the reward/loss treatment.

then a subsample of 4 subjects were brought back for an additional 200 trials in the subsequent week.

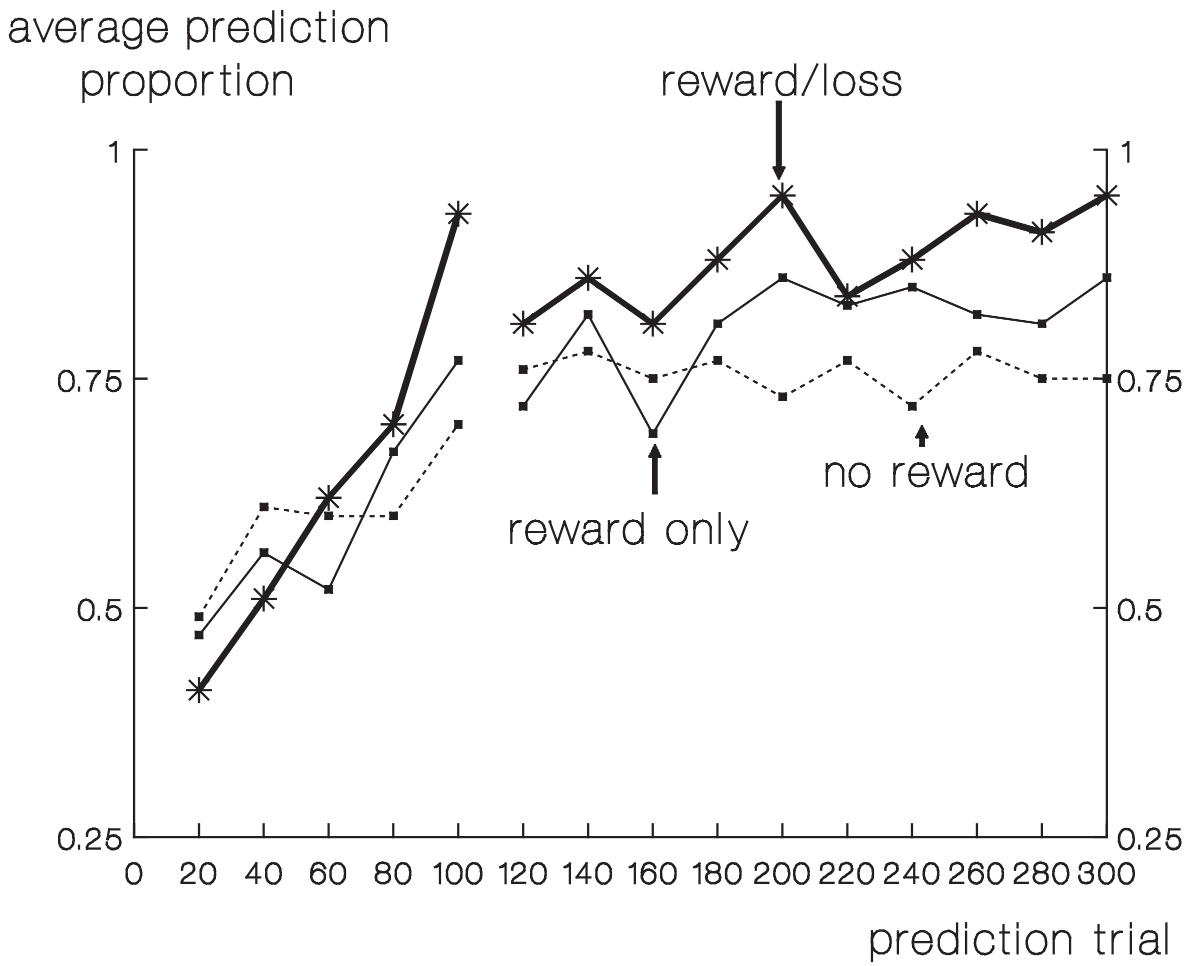

Figure 1. Mean Prediction Rates for an Event that Occurs with Probability .75, Averaged across Subjects and over 20-Trial Blocks (Source: data from Siegel, Siegel, and Andrews, 1964, figure 5, figure from Holt and Davis, 1992, figure 2.4)

Figure 1 presents the temporal pattern of data for each of the 3 treatments. Each line illustrates the average proportion of D_1 responses made by participants in a given treatment. The lines are broken by a space after trial 100, to distinguish the average-response data generated by the original 12 participants from the average-response data generated by the 4 participant subsample in the last 200 trials. The prediction proportions start at about .5 for all three treatments, which is not surprising since subjects were not given the probabilities of the two events. By the end of the first 100 trials, however, there is a rather pronounced separation in responses: In trials 80-100, the boldfaced line for the reward/penalty data reaches .93, the dashed line for the no-reward data reaches .7, and the solid line for the reward-only data is intermediate. These response patterns stabilized in the last 200 sessions. By the final block (trials 280-300), the average prediction proportion for the reward/loss treatment reached .95, the proportion for the reward-only treatment reached .85, and the proportion for the no-reward treatment stabilized at the exact level of probability matching, as noted previously.

IV. Procedures

Subjects were undergraduate students in economics classes at the University of Virginia who volunteered to participate in an experiment on the "economics of decision making". Subjects came to a computer laboratory at a specified time, and were seated at visually isolated personal computers. In a typical session, 5-8 subjects would be present, all using the same instructions and treatment condition, except as noted below. The instructions were read aloud as the subjects followed along on their screens. Subjects were not told the probabilities of the two events. The instructions gave the subjects practice in using the keyboard, but there were no "practice" periods in which predictions were made. After reading the instructions, subjects had a chance to ask questions. If a subject had a question, one of the experimenters would go to their position to answer it. Then subjects were prompted by the software to predict which of two colored boxes would appear in their computer screen, in a series of 150 trials. The two colors, Blue and Magenta, were predicted by pressing the "B" or "M" keys. These colors were used because the B and M keys are located on the bottom row of keyboard. The prompt could appear as: "Press B to choose Blue or M to choose Magenta." The order of the listing of the colors in this prompt was alternated; every other subject had Magenta listed first, and every other subject had Blue listed first. In addition, Blue turned out to be the more likely color for half of the subjects, and less likely color for the other half.[3]

[3] The software automatically alternated the labeling and order of presentation of the more likely event on the basis of a subject's ID number. The four combinations are: Blue is more likely and is presented first, Blue is more likely and is presented last, Magenta is more likely and is presented first, and Magenta is more likely and is presented last. These

Two sets of instructions were used. Before discussing these, it is useful to describe the procedures used by Siegel in his non-computerized sessions. Siegel had each subject sit at a table with a plyboard partition across the middle of the table that separated the subject from the experimenter. There were two light bulbs on the partition. The subject made a prediction as to which bulb would light up by pressing one of two response buttons. After a 3 second delay, the experimenter would switch on one of the bulbs. Two aspects of this procedure are noteworthy. First, the experimenter learned the subject's choice immediately, since the subject's response buttons were connected to small lights on the experimenter's side of the partition. The experimenter recorded this choice manually. This at least raises the possibility that the subject would view the process as a game in which the events were not really random. Second, the event sequence was not really random; it was predetermined in a random manner, subject to two constraints: the more frequent light was never allowed to illuminate more than 6 times in a row, and the 75:25 relative frequency was maintained in each block of 20 trials.

The Siegel instructions were adapted for computer use (see the attached appendix). This required using illuminated windows on the computer monitor instead of light bulbs, and it required a little training on the use of the computer keyboard, with some additional prompts such as "Press ENTER to continue." But the key statements concerning the nature of payoffs and objectives were identical.

A second set of instructions, the "Holt instructions" was also used. These differed in that they provided much more information about the nature of the random process. Subjects were told that a random number between 0.00 to 1.00 would be generated, and that any number in between would be "equally likely". If this random number was less than a fixed number x, the outcome would be one color, and if the number were greater than x, the outcome would be the other color. Subjects were not told the value of x, but they were told (and reminded after each 50-trial round) that it would be unchanged and independent of their own previous predictions. One additional advantage of this description of the random process is that it will be possible to provide probability information in a future treatment by simply using the current instructions and

---

treatment conditions will be blocked by using numbers of subjects in each treatment cell (no-reward, reward, etc.) that are multiples of 4.

replacing *x* with a specific number. These instructions also presented answers to common questions, such as what is meant by "equally likely".

V. Data and Discussion

Figures 2 and 3 plot the average proportion of times that the more likely event was predicted for all subjects in a treatment group. Each data point is an average over 10 trials, so the first point on the left for a treatment is for trials 1-10, etc. There were 150 trials in all treatments. The individual and treatment averages, by 10-trial blocks, are listed in the Data Appendix.

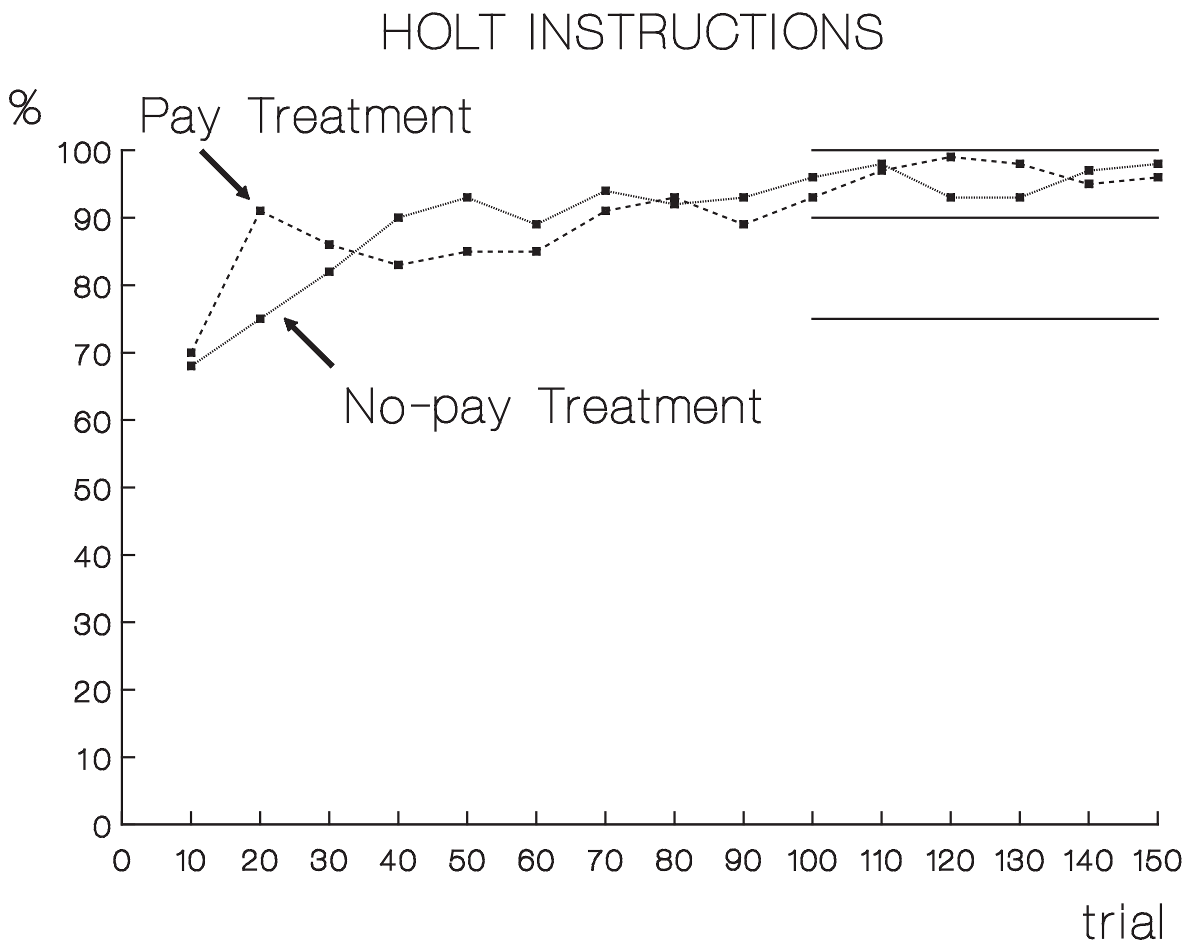

Figure 2. A Comparison of No-Pay and Pay Treatments with Holt Instructions

First consider figure 2, which shows the results for the Holt instructions. There were 12 subjects in the Holt no-pay treatment, and 11 subjects (to date) in the Pay treatment, which used a 20-cent reward for correct predictions. There is no clear effect of providing financial rewards, and both treatments yield a very high proportion of rational predictions in the final 50 trials (the horizontal lines are drawn at .75, .9, and 1.0). There is no tendency for subjects to engage in probability matching.

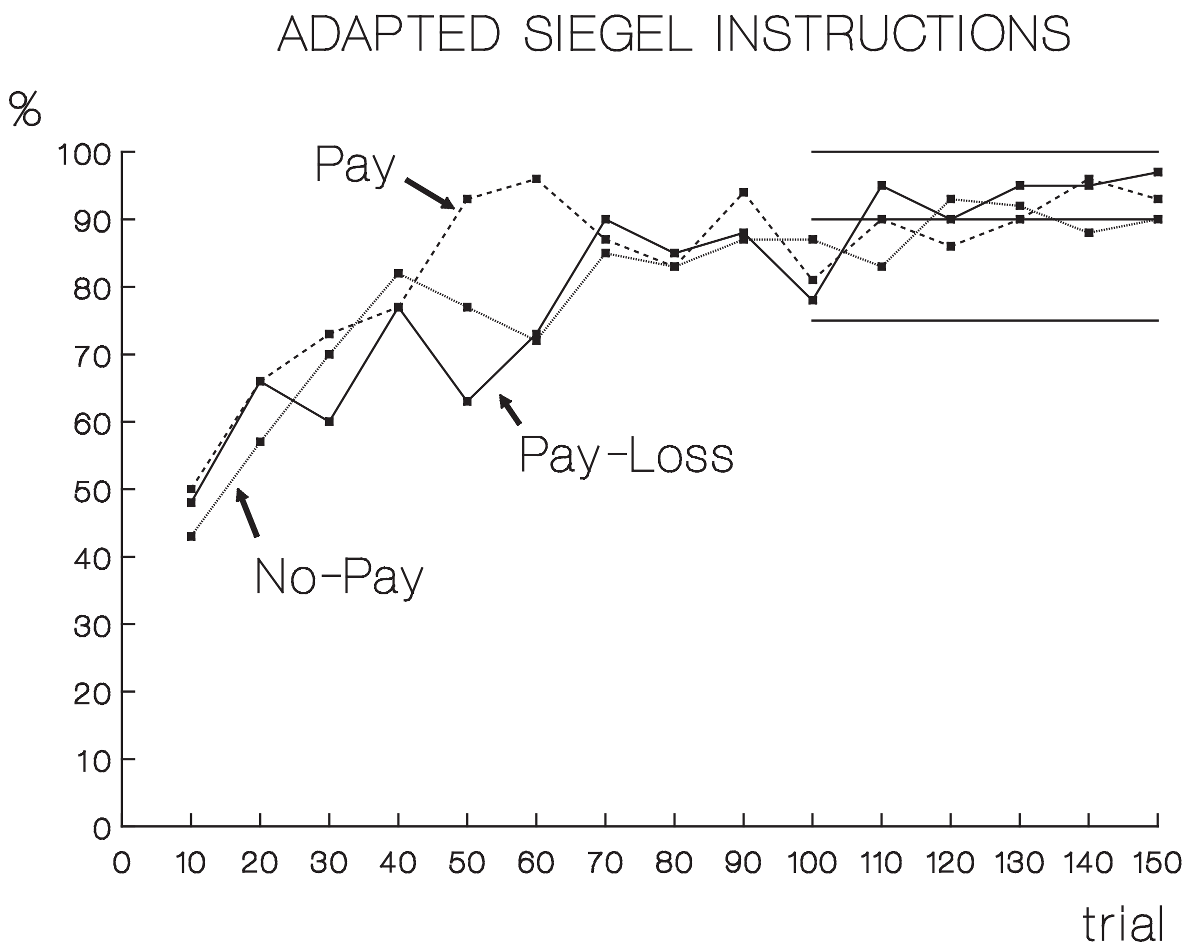

Figure 3. A Comparison of No-Pay, Pay (20 cents), and Pay-Loss (+-10 cents), with the Adapted Siegel Instructions

Next consider figure 3, which shows comparable results from using the Siegel instructions. There were 6 subjects (to date) in the No-Pay treatment, 7 subjects (to date) in the Pay treatment, which used a 20-cent reward for correct predictions, and 6 subjects in the Pay/Loss treatment, with a 10-cent reward for a correct prediction and a 10-cent loss for an incorrect prediction. Again, the prediction proportions end up in the range from .9 to 1.00, which is inconsistent with probability matching. There is little apparent effect of providing financial incentives, which is inconsistent with the earlier Siegel results. In addition, note that the prediction proportions start out near .5 on average, which is consistent with the original Siegel data in figure 1, but which is much lower than with the Holt instructions in figure 2. The tendency for predictions to be more rational with the Holt instructions is highlighted in figure 4, which compares the data for the two instructions for each treatment.

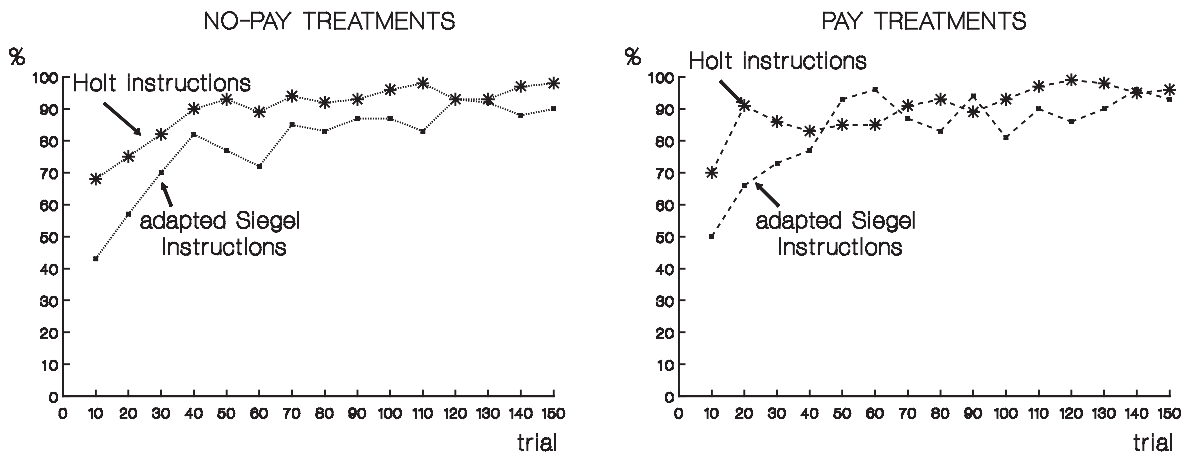

Figure 4. A Comparison of the Holt and the Adapted Siegel Instructions

One may not conclude from these results that participants always behave differently when provided with financial incentives. The matching behavior observed by Siegel and Goldstein (1959) without incentives may have been prompted by a variety of elements in the design and administration of the experiment. As shown in the present paper, it is possible to set up a laboratory forecasting experiment in which such matching behavior is not observed after sufficient experience, with or without incentives. As a rule, however, economists are rather uninterested in sorting out the potential biases of these effects. What we may conclude from the Siegel and Goldstein (1959) experiment is that financial incentives can sometimes eliminate subtle and unintended biases. For this reason, the payment of financial incentives is a critical element in the administration of economics experiments.

References

Davis, Douglas D. and Charles A. Holt (1992) *Experimental Economics*, Princeton, N.J.: Princeton University Press.

Siegel, Sidney and D. A. Goldstein (1959) "Decision-making Behavior in a Two-Choice Uncertain Outcome Situation," *Journal of Experimental Psychology*, *57*, 37-42.

________________, Alberta Siegel, and Julia Andrews (1964) *Choice, Strategy, and Utility*, New York: McGraw-Hill.

Simon, Herbert A. (1957) *Models of Man*, New York: John Wiley and Sons.

Instructions Appendix[4]

The experiment is about to begin. Please wait for further instructions before pressing any keys. Thank you.

Your task in this experiment is to try to predict correctly which colored box will show on the computer screen in each of a large number of trials. Press ENTER to continue.

The colored box that shows will either be blue or magenta. On the keyboard before you, notice the "b" and "m" keys at the bottom of the keyboard. If your prediction is that the square will be blue, press the "b" key. If your prediction is that the square will be magenta, press the "m" key. Press ENTER to continue.

Lets run through an example of each event. If the color turns out to be blue, the blue box will appear above. Press ENTER to see the blue box.
Press the letter "g" to go to the next screen. [If the subject presses any other key, the message box shows: "Press g to continue."]

If the color turns out to be magenta, the magenta box will appear above. Press ENTER to see the magenta box. Press the letter "g" to go to the next screen. [Press g to continue.]

[No-Payoff Treatment]

> Now do your best to predict correctly which color box will show on the screen for every trial. Press ENTER to continue.
>
> At the beginning of this session, we will credit you with a cumulative earnings level of $20.00 This is a payment to compensate you for spending time in this study.

[Reward Treatment:]

> Now, every time you predict correctly you will win $0.20, every time you predict incorrectly you will not win anything. At the end of the session all the money you have won will be yours to keep. That is, for every time you predict correctly you will win $0.20 and all of your cumulative earnings at the end of the session will belong to you. Now, if you make no correct predictions you will won have no money by the end of the game, but if you play carefully you should come out winning at the end of the session. Press ENTER to continue.
>
> Let me make it quite clear: all the money you win you can keep. There are no strings attached, and please understand you will be paid the exact amount

4 These are the adapted Siegel instructions, for the computer implementation in which the blue color is referred to first. Every other subject, however, received the colors listed in the opposite order.

you win. This money does not belong to me; it comes from a research grant, so I am not concerned with how much or how little you win. The amount you win will depend entirely on how you play. Is this clear? All right. Now do your best to predict correctly which color box will show for every trial. Press ENTER to continue.

At the beginning of this session, we will credit you with a cumulative earnings level of $ 1.00 This is your starting stake in the experiment.

[Reward/Penalty Treatment]

Now, every time you predict correctly you will win $0.10; every time you predict incorrectly you will lose $0.10 We will start you off with $16.00, and you will keep all of your winnings at the end of the session. That is, at the end of the session all of your cumulative earnings will belong to you. Now, if you do not play carefully, you are likely to lose your starting stake and therefore end the session with no money. However, if you play carefully you could end the session with a considerable amount of money which would be yours to keep. Press ENTER to continue.

Let me make it quite clear: all the money you win you can keep. There are no strings attached, and please understand you will be paid the exact amount you win. This money does not belong to me; it comes from a research grant, so I am not concerned with how much or how little you win. The amount you win will depend entirely on how you play. Is this clear? All right. Now do your best to predict correctly which color box will show for every trial. Press ENTER to continue.

At the beginning of this session, we will credit you with a cumulative earnings level of $16.00 This is your starting stake in the experiment.

[all treatments:]

PLEASE DO NOT press enter until everyone has had a chance to ask questions.

------- ARE THERE QUESTIONS? ---------

If you have a question, raise your hand and I will come around to answer it. Now, press ENTER to continue.

We now begin the trials. Please do not talk to each other during the experiment. Press ENTER.

If you have a question during the experiment, please raise your hand, and I will come to your desk to answer it. If the keyboard seems to be stuck, check to be sure that you have pressed the appropriate response, which may be "ENTER" or the first letter of a color name to predict that color. Now press ENTER to begin the experiment.

Guide to the Data Appendix

Table 2. Number of Subjects in Each Treatment Cell (with subject numbers in parentheses)

	No-Incentive Treatment	Reward Treatment	Reward/Penalty Treatment
Holt Instructions	12 (m101-m112)	11 (m201-211)	-
Adapted Siegel Instructions	6 (m113-m115, m117-119)	7 (m213-m219)	6 (m313-m318)

In the data appendix that follows (**** available on request *****) , the average proportions for the more likely event are calculated for 10-trial groupings, so each subject has 15 such averages, i.e. for trials 1-10, 11-20, ... 141-150. The 15 average prediction proportions are listed below the person's subject number. Subjects are grouped by treatment, and the far right column shows the average across subjects in a treatment for each 10-trial grouping. For example, 12 subjects, with numbers m101-m112, used the Holt instructions without a "no-pay" condition. The far right column shows that the more likely event was predicted 67.5% of the time in the first 10 trials for these subjects, but the more likely event was selected 98% of the time in the last 10 trials.