

A Laboratory Investigation of Deferral Options ^{*}

Ryan Oprea[†] Daniel Friedman[‡] Steven T. Anderson[§]

March 21, 2007

Abstract

An irreversible investment opportunity has value V governed by Brownian motion with upward drift and random expiration. Human subjects choose in continuous time when to invest. If she invests before expiration, the subject receives $V - C$: the final value V less a given avoidable cost C . The optimal policy is to invest when V first crosses a threshold $V^* = (1 + w^*)C$, where the option premium w^* is a specific function of the Brownian parameters representing drift, volatility and discount (or expiration hazard) rate.

We ran 80 periods each for 69 subjects. Subjects in the Low w^* treatment on average invested at values quite close to optimum. Subjects in the two Medium treatments and the High w^* invested at values below optimum, but with the predicted ordering, and values approached the optimum by the last block of 20 periods. Behavior was most heterogeneous in the High treatment. Subjects underrespond to differences in both the volatility and expiration hazard parameters. A directional learning model suggests that subjects react reliably to ex-post losses due to early investment, and react much more heterogeneously (and on average more strongly) to missed investment opportunities. Simulations show that this learning process converges on a nearly optimal steady state.

Keywords: Real options, optimal stopping, laboratory experiment, fractional factorial design.

JEL codes: G13, D83, C91

^{*}We are grateful for the research assistance of Jerome Bragdon and Todd Feldman, for the programming support of Nitai Farmer and Adam Freidin, and for funding from the National Science Foundation under grant IIS-0527770.

We received helpful comments from participants at the Fall 2006 ESA meetings and a UCSC brown bag seminar.

[†]Economics Department, University of California, Santa Cruz, CA, 95064. roprea@ucsc.edu

[‡]Economics Department, University of California, Santa Cruz, CA, 95064. dan@ucsc.edu

[§]Minerals Information Team, U.S. Geological Survey, Reston, VA 20192. steve@email

1 Introduction

Real options theory came of age in the 1990s. The theory shows why the classic net present value rule for investment can seriously mislead in the presence of irreversibility, new information, and the opportunity to defer. There are now numerous books and articles aimed at practitioners that explain the option to defer and other real options such as those to abandon, to expand, to contract, and to switch, and assorted compound options. These option values turn out to be quite important in all sorts of investment decisions. Yet recent surveys indicate that fewer than 30% of large corporations actually use the theory (Copeland and Antikarov, 2003, p. viii).

What accounts for the disconnect between widely known theory and recent actual practice? One possibility is that experienced practitioners can closely approximate the optimal exercise of available options, and the formal analysis (with its inherent costs and measurement errors) would seldom add much value.

In this paper we consider the most basic, and most intensively analyzed, real option: the option to defer.¹ The underlying investment problem is easy to grasp, but its mathematical solution is subtle and was not known until the late 20th century.

Our conjecture is that ordinary people can closely approximate optimal exercise of deferral options if they are given a decent chance to learn from personal experience. The conjecture might help explain the disconnect between theory and practice just noted, and it is of interest in its own right. There are some mathematically trivial tasks that seem to be quite difficult to learn (e.g., optical illusions, or the three door problem, e.g., Kluger and Wyatt (2004)) and conversely there are some mathematically intractable tasks (such as optimal or equilibrium bidding in a continuous double auction market, e.g., Friedman and Rust (1993)) that even children learn quite quickly. Investigating the conjecture may help sort out learnability issues.

As noted, there is an extensive literature on the theory and application of real options. We are not aware, however, of any studies systematically comparing actual behavior in the field to theoretically optimal exercise of real options. Nor are we aware of any previous laboratory studies. Perhaps the closest existing laboratory studies investigate sequential job search, beginning with Braunstein and Schotter (1981); more recent work includes Cox and Oaxaca (2000) and Gong and Ramachandran (2007). Most of these studies find that actual search behavior approximates optimal behavior fairly well, but the task has only a slight mathematical resemblance to ours and it seems

¹Various authors refer to it as a timing option, a wait option, an initiation or delay option, or a real call option; earlier literature refers to the optimal stopping problem.

completely different to human subjects.

The next section lays the theoretical foundations. After describing the basic investment timing task, it presents the known formula for the deferral (or “wait”) option premium w^* as a function of exogenous Brownian parameters representing the volatility (σ) and trend (α) of the investment value, and the discount rate (ρ). Given a small time step Δt , we approximate the continuous stochastic process with binomial parameters: the proportional tick size h , the uptick probability p and the expiration probability q . The section concludes with comparative static results on how w^* and thus optimal behavior will vary with h, p and q , the parameters controlled in the lab.

Section 3 obtains a series of testable predictions based on the conjecture that actual behavior will approximate optimal behavior. The following section presents the experimental design, which involves 69 human subjects, each assigned to one of four distinct treatments. The Low treatment uses parameter values that lead to $w^* \approx 0.2$, two Medium treatments use rather different combinations of parameter values that lead to $w^* \approx 0.5$, and the High treatment leads to $w^* \approx 0.8$. The treatments form a fractional factorial design that allows us to infer the separate effects h, p and q .

Section 4 collects results. Broadly speaking, they support the conjecture, but there are interesting exceptions. Behavior converges quickly to the optimum in the Low treatment, but in other treatments subjects tend towards impatience, especially in early periods. An underlying reason is under-response to changes in volatility and hazard parameters. With experience subjects attach greater weight to the option value and, over time, behavior approaches optimality in all treatments. A directional learning model suggests that subjects react reliably to ex-post losses due to early investment, and react much more heterogeneously (and on average more strongly) to missed investment opportunities. Simulations show that this learning process converges on a nearly optimal steady state.

Appendix A includes a self-contained derivation of the main mathematical results, including the comparative statics. Appendix B contains technical material on the econometric procedures and includes robustness checks of our main findings. Appendix C reproduces the instructions to subjects.

2 Theoretical Background

An investor with discount rate $\rho > 0$ can launch a project whenever she chooses by sinking a given fixed cost $C > 0$. The present value V of the project follows geometric Brownian motion with

drift parameter $\alpha < \rho$ and volatility parameter $\sigma > 0$. In other words, the value of the project is uncertain and follows a continuous time random walk in which the appreciation rate has mean α and standard deviation σ per unit time. A concise mathematical description is the stochastic differential equation

$$dV = \alpha V dt + \sigma V dz, \tag{1}$$

where z is the standard Wiener process. At times $t \geq 0$ prior to launching the project, the investor observes $V(t)$ (and previous values $V(s)$ for $0 \leq s \leq t$). If she invests at time t , then she obtains payoff $(V(t) - C)e^{-\rho t}$. The project is irreversible, and generates no other payoffs. Thus the task is to choose the investment time so as to maximize the expected payoff.

Chapter 5 of Dixit and Pindyck (1994) centers on this classic problem, giving credit to McDonald and Siegel (1986).² It turns out that the optimal policy takes the form: wait until $V(t)$ hits a specific threshold V^* , then launch immediately. The threshold is proportional to cost, i.e., $V^* = (1 + w^*)C$, where the option premium w^* is an algebraic function of the volatility, drift and discount parameters σ, α and ρ . Specifically,

$$w^* = \frac{1}{B - 1}, \text{ where } B = \frac{1}{2} - \frac{\alpha}{\sigma^2} + \sqrt{\left[\frac{\alpha}{\sigma^2} - \frac{1}{2}\right]^2 + \frac{2\rho}{\sigma^2}} > 1. \tag{2}$$

See Appendix A for a self-contained derivation.

2.1 Discrete Time Representation

Continuous time is useful analytically, but a computer implementation necessarily is in discrete time. Binomial approximations of the Brownian process (1) involve a fixed time interval Δt for each discrete step, and three parameters corresponding to those of the Brownian process (e.g., Dixit and Pindyck, p. 69):

- the step size $h > 0$ of the proportional change in value, i.e., the current value V becomes either $(1 + h)V$ or $(1 - h)V$ at the next step;
- the uptick probability $p \in (0, 1)$, i.e., the probability that the next step is to $(1 + h)V$ rather than to $(1 - h)V$; and
- the expiration probability $q \in (0, 1)$, i.e., the probability that the next step is the last, and the opportunity disappears.

²They in turn note earlier authors going back to Henry (1974), and describe their contribution as allowing for risk averse investors.

The deviation of the uptick probability p from 0.5, times the distance between an uptick and downtick, corresponds to the drift rate α :

$$\alpha = \lim_{\Delta t \rightarrow 0} \frac{(2p - 1)h}{\Delta t}. \quad (3)$$

The volatility σ comes mainly from the stepsize h but when p differs appreciably from 0.5 we must also account for binomial variance $p(1 - p)$. The exact expression is

$$\sigma^2 = \lim_{\Delta t \rightarrow 0} \frac{4p(1 - p)h^2}{\Delta t}. \quad (4)$$

The relation between expiration probability $q \in [0, 1]$ and the discount parameter $\rho > 0$ may require a little more explanation. By definition of discount, one unit value received one time unit from now is equivalent to $e^{-\rho}$ units received immediately. The usual interpretation of such impatience is the foregone interest payments available in a financial market. For practical reasons, we turn to an alternative interpretation (e.g., Kreps, 1990, p.505-6): the discount rate arises from the possibility that the opportunity will expire.³ If that event has probability Q per unit time, then the expected value of 1 unit due one time unit in the future is $1 - Q$. With $T = 1/\Delta t$ time steps per unit time, and expiration probability q per time step, the discount factor is $e^{-\rho} = 1 - Q = (1 - q)^T = (1 - q)^{1/\Delta t}$. Solving for ρ we obtain

$$\rho = \frac{-\ln(1 - q)}{\Delta t}. \quad (5)$$

2.2 Comparative Statics

How do changes in the binomial parameters (h, p, q) affect the option premium w^* ? Equations (3-5) show the effects on the Brownian parameters (α, σ, ρ) , and equation (2) allows us to follow the effects back to B and ultimately to w^* . For example, differentiate the expression (5) to obtain $\frac{\partial \rho}{\partial q} = \frac{1}{\Delta t(1 - q)}$, and differentiate the expression for B in (2) to obtain $\partial B / \partial \rho = 1/D$, where $D = \alpha + \frac{1}{2}\sigma^2(2B - 1) > 0$; see the Appendix for some shortcuts and minor caveats. Then differentiate the expression for w^* in (2) with respect to B to obtain $dw^*/dB = -1(B - 1)^{-2} < 0$. Since $\frac{\partial w^*}{\partial q} = \frac{dw^*}{dB} \frac{\partial B}{\partial \rho} \frac{\partial \rho}{\partial q}$ we obtain

$$\frac{\partial w^*}{\partial q} = \frac{-1}{(B - 1)^2 D (1 - q) \Delta t} < 0. \quad (6)$$

³In laboratory settings like ours, the subject receives all payments at the same time, at the end of the session. Another practical issue is that the time horizon in the lab is not infinite. It might also be worth noting that hazard rates are in theory the ultimate reason for impatience and discounting. Neoclassical theory obtains the real risk-free interest rate from financial market participants' marginal returns and marginal rates of time preference (e.g., Hirshleifer et al, 2005, Chapter 15), while evolutionary theory traces time preference back to the risk of death and other biological hazards (e.g., Robson, 2002).

Treatment	Step Size	Uptick Prob	Expiration Prob	Option Prem
τ	h	p	q	w_τ^*
Low	0.0155	0.513	0.007	0.179
Medium A	0.0155	0.524	0.003	0.490
Medium B	0.03	0.524	0.007	0.499
High	0.03	0.513	0.003	0.804
Effect on Premium	0.317	0.003	-0.308	

Table 1: Binomial parameter values and option premium by treatment. The last row indicates the average change in the option premium when the parameter moves from its lower to its higher value. In all treatments, the time step in minutes is $\Delta t = 0.00\bar{3}$.

The Appendix applies similar reasoning to obtain

$$\frac{\partial w^*}{\partial p} \approx \frac{2hB}{(B-1)^2 D \Delta t} > 0, \quad \text{and} \quad (7)$$

$$\frac{\partial w^*}{\partial h} \approx \frac{[2p-1 + (B-1)\sigma\sqrt{\Delta t}]B}{(B-1)^2 D \Delta t} > 0. \quad (8)$$

For example, fix the time step at 200 milliseconds, so measuring time in minutes a “tick” is $\Delta t = 200/(1000 * 60) \approx 0.0033$. Consider the binomial parameter values $h = .03, p = 0.524$ and $q = 0.007$. Then by straightforward calculation one verifies $B \approx 3.0$ and the optimal option premium is $w^* \approx 0.5$. Other things equal, lowering the per-tick expiration rate from $q = .007$ to $q = .003$ would increase w^* by about $\frac{.007-.003}{2^2(1.11)(0.993)(.0033)} \approx 0.275$. The approximation is far from exact (we will see later that the appropriate number for our experiment is 0.308) because (6) is not constant and indeed is highly non-linear. Moreover, expressions (6-8) have strong interactions among the binomial parameters. Consequently a careful design is needed to untangle the actual impact of the parameters.

3 Treatments and Hypotheses

The experiment has four treatments, as shown in Table 1. Several considerations affected the choice of parameter values. First, the task must have a reasonable “look and feel” for subjects. The task seems continuous with a time step of $\Delta t = 0.00\bar{3}$ minutes, i.e., 5 ticks per second. In this case an expiration probability above $q = 0.010$ ends too many periods before the subject can get focused, and below $q = 0.002$ it takes too long to run a reasonable number of periods. Step sizes above

$h = 0.05$ seem very jagged and hard to display on a consistent V -scale, and the upward trend seems to dominate volatility with uptick probabilities much above $p = 0.53$.

A second consideration, also crucial, is that the treatments should generate at least two and preferably three widely separated option premiums. The Low, Medium and High premiums are distinct and roughly equidistant. A third consideration is to have different parameter configurations produce similar premiums in different ways, hence the two Medium treatments. A fourth consideration is to employ a half factorial design to facilitate comparative static evaluations, described below.

Finally, we chose our parameters in order to avoid flat payoffs around the optimal policy. Figure 1 shows payoffs in simulations where the agent invests at a fixed $w = 0.1, 0.2, \dots, 1.5$. Note that the simulated payoffs are indeed maximized very near the optimum $w = w^*$, but that expected payoff becomes rather flat for $w > w^*$ in treatments with larger w^* . We chose parameter configurations to mitigate the problem, but it seems infeasible to avoid flat payoff landscapes for w^* above 1.0.

The option premiums displayed in Table 1 are the basis of the first two testable hypotheses. As explained in more detail below, the observed premium w_τ in each treatment τ is obtained from the expression $V/C - 1$ where V is the actual investment value in a period for which the realized cost is C .

Hypothesis 1 *Observed option premiums w_τ have the same ordinal rank across treatments as do the theoretical premiums: $w_{Low} < w_{MedA} \approx w_{MedB} < w_{High}$.*

Hypothesis 2 *In each treatment τ , the observed option premium is on average equal to the theoretical premium w_τ^* .*

The half factorial design permits a clean estimate how each binomial parameter affects observed premiums, despite the fact that functions (6-8) are not constant or even linear. In Table 1, each parameter assumes two values, one high and one low, and these values are balanced across treatments. The average effect Δj of a binomial parameter $j = h, p, q$ therefore can be obtained from average observed premium in the two treatments using the high value of j minus the corresponding average in the two treatments using the low value of j . Thus we have

Hypothesis 3 *The average effect of parameter changes on premiums is given by*

$$\Delta h = \frac{w_{High}^* + w_{MedB}^*}{2} - \frac{w_{Low}^* + w_{MedA}^*}{2} = 0.317 \quad (9)$$

$$\Delta p = \frac{w_{MedA}^* + w_{MedB}^*}{2} - \frac{w_{Low}^* + w_{High}^*}{2} = 0.003 \quad (10)$$

$$\Delta q = \frac{w_{Low}^* + w_{MedB}^*}{2} - \frac{w_{High}^* + w_{MedA}^*}{2} = -0.308 \quad (11)$$

The numerical values are repeated at the bottom row of Table 1. Notice that Hypothesis 3 predicts only a small impact for the uptick probability parameter p ,⁴ and predicts that the other two binomial parameters will have effects of roughly equal magnitudes but in opposite directions.

4 Implementation

Figure 2 shows the screen observed by each subject. At the beginning of each of 80 periods, the subject is given a cost C drawn independently and uniformly from integers between 50 and 110. The period starts with the value V at C , and evolves according to the discrete binomial approximation of Brownian motion described in the previous section. When a subject clicks the button to invest, she earns $V - C$ points. Even after clicking, the subject can see the green V -line evolve until the random expiration time, controlled by the parameter q . Costs, value paths and ending times were randomly generated in real time. They were identical across subjects in a given session and were not repeated in subsequent sessions.

In addition to the graphical display, the experimental software shows the numerical values of V and C , earnings in the current period, and cumulative earnings so far (“Total Score”). The user interface allows each subject to view her previous decisions and earnings at any time.

Subjects were 69 undergraduate students at the University of California, Santa Cruz. At the beginning of each of ten sessions, each subject was seated at a visually isolated computer terminal and assigned to a treatment (e.g., Medium B). Instructions were read aloud and the software was displayed on a screen. The binomial parameters for the chosen treatment were explained and written on a white board. Subjects participated in 6 practice periods. Each subject then participated in 80

⁴This is a feature of the particular design we chose in light of the second, third and fourth considerations. Due to intrinsic interactions and other non-linearities, the same parameter p has a large predicted impact in some other designs. For example, in the complementary half factorial design, the predicted impact is 0.456.

periods for pay with no change in treatment. Sessions lasted between 80 and 120 minutes depending on the treatment and random draws made over the course of the session.

A total of 17 subjects participated in Low, Medium A and Medium B treatments and 18 participated in the High treatment. No subject was allowed to participate in more than one session. A subject with cumulative payoff π over all periods received $a(\pi - b)$ cents in cash at the end of the session. To reduce the wide variation in expected earnings across treatments, we chose $b = 200$ in Low and Medium treatments and $b = 250$ in the High treatment, with $a = 4$ in the Medium and High treatments and $a = 5$ in the Low treatment. We also paid subjects a \$5 bonus show-up fee. On average subjects received \$9.18 in the Low treatment, \$18.71 in Medium A, \$19.78 in Medium B and \$27.40 in High.

5 Results

Figure 3 shows raw data, plotting observed investment values and corresponding costs in each treatment. The solid lines are linear fits to the raw data while dotted lines indicate optimal investment decisions according to equation (2). Note that in the Low treatment, observed investment values vary with cost roughly as the theory predicts. In all other treatments, most observations fall well below predicted values.

Figure 3 can be misleading because it doesn't show all the data: it omits the 56% of periods that expired before the subject invested. The omitted periods are more likely to come from higher choices of the premium w . Hence estimates of w based directly on raw data are likely to be biased downward.

To deal with this sampling bias, we use the product-limit estimator (PLE), a maximum-likelihood non-parametric estimator for cumulative distribution functions designed to account for random right censoring (Kaplan and Meier, 1958). The PLE is often used in medical studies to estimate distribution functions of subjects' time-until-death when many subjects exit the study before dying. We adapt the technique to study investment premiums when the opportunity expires before some subjects invest. Appendix B describes the estimator more formally, and also describes the log-rank test, the standard technique for comparing PLEs across treatments.

We apply the PLE estimator to data pooled across subjects, and also apply it separately to data from each individual subject. The pooled estimates assume i.i.d. observations; to the extent that different subjects behave differently, the assumption is violated and the standard errors are

Estimator	Low	Med A	Med B	High
Model Predictions				
	0.179	0.490	0.499	0.804
Means				
Pooled PL	0.182 <small>± 0.004</small>	0.315 <small>± 0.008</small>	0.338 <small>± 0.011</small>	0.605 <small>± 0.037</small>
By-Sub PL Mean	0.182 <small>± 0.011</small>	0.310 <small>± 0.020</small>	0.328 <small>± 0.030</small>	0.575 <small>± 0.065</small>
Raw	0.131 <small>± 0.004</small>	0.241 <small>± 0.006</small>	0.213 <small>± 0.006</small>	0.362 <small>± 0.012</small>
Observations	1360	1360	1360	1440
Uncensored	590	593	622	548

Table 2: Mean observed deferral option premiums. Model Predictions are the theoretical option premiums from Table 1. The Pooled PL row reports the means of the estimated distributions in Figure 4. The next row reports the mean of the separate PL estimates for each subject from Figure 5. The Raw row reports results from the subsample of uncensored observations. Standard errors are reported in smaller font following the \pm symbol.

likely to be understated. On the other hand, individual subject estimates use smaller samples and are noisier. For robustness, we also checked our main results using Tobit models. The next two sections focus on the PLE results; the Tobit results turn out to be similar and are reported in Appendix B.

5.1 Observed Premiums

Figure 4 shows the pooled PL estimated distributions of option premia, while Figure 5 indicates the heterogeneity across individual subjects in terms of the PL estimated means and medians. Table 2 summarizes the central tendencies gleaned from both figures. The sampling bias due to censored observations can be seen by comparing the Raw rows to the PL rows.

The figures and tables seem quite consistent with Hypothesis 1, which states that $w_{Low} < w_{MediumA} \approx w_{MediumB} < w_{High}$. Figure 4 shows a clear separation between the Low, Medium and High cumulative distributions. Recalling that a higher cumulative distribution function indicates a larger fraction of smaller observations, one sees that the figure reflects the predicted ordinal rankings. Log-rank tests comparing each pair of treatments allow us to reject at the one percent level the null hypotheses that Low, Medium and High CDFs are equal, in favor of alternative Hypothesis 1. Only the null hypothesis that the two Medium CDFs are equal survives this log-rank

test ($p = 0.739$).

Figure 5 also suggests that option premia in the Low treatment tend to be smaller than those in Medium treatments which are in turn smaller than option premia in the High treatment. Applying Mann-Whitney tests to mean samples, we reject at the one percent level the null hypothesis in favor of the predicted alternative. The same tests again fail to reject only the null hypothesis that the two Medium treatments produce the same premia ($p = 0.8497$).

Finding 1 *The data support Hypothesis 1. Observed option premiums are significantly different across Low, Medium and High treatments and have the predicted ordinal rank. Also as predicted, we cannot reject the hypothesis that behavior is the same in the two Medium treatments.*

The point predictions of Hypothesis 2 do not fare as well as the ordinal predictions of Hypothesis 1. In Figure 4, the Low PLE distribution curve crosses the dashed line near its median: about 55% of actual option premiums are estimated to be below the predicted value of 0.18 and about 45% above. So far, so good. However, in the three other treatments, median option premiums lie well below predicted levels. Indeed, the figure indicates that more than 80% of actual option premiums are less than predicted in both Medium treatments as well as the High treatment. The by-subject estimates in Table 2 tell a similar story, with Low option premiums coming close to predictions and Medium and High premiums falling well below them.

We formally test this observation using by-subject product-limit estimates. We cannot reject the hypothesis that $w_{Low} = 0.179$ using Wilcoxon signed-rank tests for either mean or median distributions ($p = 0.758$ and 0.286 respectively). Using the same test on mean samples we can reject the hypothesis that $w_{High} = 0.804$ ($p = 0.012$), $w_{MedA} = 0.490$ ($p = 0.000$), $w_{MedB} = 0.499$ ($p = 0.049$).

Finding 2 *The data support only part of Hypothesis 2. In the Low treatment, option premiums center on the predicted level. In other treatments, premiums are significantly lower than predicted.*

5.2 Comparative Statics

Is the premature investment due mainly to under-reaction to one of the binomial parameters? Hypothesis 3 gives the optimal reaction and Table 3 shows that, as predicted, the uptick parameter Δp has small impact, and the prediction lies within the estimated confidence interval. The other

Parameter	Predicted	Mean	Median
Δp	0.003*	-0.067 * (0.040)	-0.03 (0.009)
Δh	0.317	0.223 (0.040)	0.152 (0.009)
Δq	-0.308	-0.200 (0.040)	-0.143 (0.009)

Table 3: Binomial parameter impact estimates. Predicted values are from the bottom row of Table 1. Mean (respectively Median) estimates are obtained from equations (9-11) applied to PL estimates reported in the middle (resp. lower) portion of Table 2. Standard errors are in parentheses. Asterisks indicate estimates whose 95% confidence intervals include the prediction.

parameters have much larger impacts and in the predicted directions. However, neither impact is quite as large as predicted, and the predictions lie outside the 95 % confidence intervals. We collect these observations as a third finding.

Finding 3 *The data qualitatively support Hypothesis 3. Subjects do not respond to small changes in the uptick probability, as predicted. Subjects do increase option premia in response to either an increase in the tick size or a decrease in the expiration probability, but not as much as predicted.*

5.3 Learning

Could the underresponse to the uptick and expiration parameters be due simply to slow learning? Table 4 provides some favorable evidence. The table shows that product-limit estimates usually increase from one time block (e.g., periods 1-20) to the next (e.g., periods 21-40). The estimates in the final block (periods 61-80) are all closer to the prediction than in the initial block, and in every case the difference is significant at the 1% level according to log-rank tests. The mean estimates in the final block are remarkably close to prediction in the High and Low treatments, and within 90% of prediction in the Medium treatments.

To control for intrasubject correlations, we compare the means of individual subject PL estimates across blocks. Mann-Whitney tests reject at the 1% level the null hypothesis that these estimates in the first time block are identical to those in the last block for the Low and Medium A treatment in favor of the learning hypothesis that they are larger in the last block. The same test is significant at the ten percent level in the Medium B treatment. However, we cannot reject

Block of Periods					
	1-20	21-40	41-60	61-80	Overall
Means					
Low	0.156 ± 0.008	0.172 ± 0.009	0.199 ± 0.011	0.198 ± 0.007	0.182 ± 0.004
Medium A	0.263 ± 0.011	0.323 ± 0.017	0.288 ± 0.014	0.370 ± 0.017	0.315 ± 0.008
Medium B	0.309 ± 0.022	0.288 ± 0.012	0.337 ± 0.024	0.395 ± 0.022	0.338 ± 0.011
High	0.480 ± 0.030	0.546 ± 0.029	0.555 ± 0.053	0.797 ± 0.110	0.605 ± 0.037

Table 4: Product-Limit estimates of option premia by time block and treatment. Standard errors are reported in smaller font following the \pm symbol.

the hypothesis that early and late option premium estimates are identical for the High treatment ($p = 0.569$).

Restricting the signed rank tests reported in 2 to the final time block provides another clue. The restricted test results are pretty much the same as the unrestricted tests in the first three treatments: the Low markup option premia are as predicted and both Medium premia are significantly below prediction. However, in the High treatment, restricting to the last time block we can no longer reject the point prediction $w^* = 0.80$.

How can this last finding be reconciled with the absence of a significant learning trend in the High treatment? The culprit is a large increase in the by-subject variance. In the first block the by-subject estimates in the high treatment average 0.48 and have standard deviation 0.21. In the last block the average rises to 0.74 and the standard deviation jumps to 0.82. Evidently in the High treatment, accumulated experience leads different subjects to different strategies.

Finding 4 *In Low and Medium treatments, estimated option premiums tend to increase over time. In the High treatment behavior becomes increasingly heterogeneous over time. We cannot reject the hypothesis that, by the end of a High treatment session, the average option premium reaches the predicted level.*

5.4 A Directional Learning Model

What causes the increasing heterogeneity across subjects in the High treatment? Might it be due to the shape of the payoff function? Recall from Figure 1 that the High treatment has an especially skewed payoff function that is quite flat for option premia above 0.6 or 0.7. On the other hand,

learning seems much more rapid in the Low treatment, with estimated option premia reaching the point prediction halfway through the session. Might this be due to the more symmetric and sharply peaked payoff function in this treatment?

We investigate such questions using a parametric learning model adapted from Cason and Friedman (1999), which in turn is based on Selten's directional learning model (e.g., Selten and Buchta, 1998). The idea is that subjects increase the intended option premium when they invest before the ex-post optimal time, and decrease the intended premium when a profitable opportunity expires before investment.

To specify the model, let w_t be the option premium in period t , let π_t be the earnings in that round, and let π_t^o be the ex-post maximum earnings available in that round. Let D_t^- be the indicator variable that is 1 if a subject invested at a value lower than the ex-post maximum round t and otherwise is zero. Likewise, let D_t^+ be 1 if a subject did not invest in round t and otherwise be zero. The model describes adjustment in the option premia as follows:

$$w_t - w_{t-1} = \nu + [\delta_- D_{t-1}^- + \delta_+ D_{t-1}^+][\pi_{t-1}^o - \pi_{t-1}] \quad (12)$$

Here, ν is a trend variable which should be equal to zero if the model is correctly specified. The coefficient δ_- is the subject's sensitivity to a dollar lost in the last period from investing lower value than the ex-post optimum, and δ_+ is the analogous sensitivity to a dollar lost by waiting so long that a profitable opportunity expires.

At first it seems natural to estimate the parameters $(\nu, \delta_-, \delta_+)$ of (12) directly, by simply appending an additive error term, say

$$w_t - w_{t-1} = \nu + [\delta_- D_t^- + \delta_+ D_t^+][\pi_{t-1}^o - \pi_{t-1}] + \epsilon_{t-1} \quad (13)$$

with ϵ_{t-1} assumed to be distributed $N(0, \sigma^2)$. However, there is a censoring problem. Although we always observe D^+ , D^- , π and π^o , we never observe w when $D^+ = 1$. Therefore, consider the latent model

$$\tilde{w}_t - \tilde{w}_{t-1} = \nu + [\delta_- D_t^- + \delta_+ D_t^+][\pi_{t-1}^o - \pi_{t-1}] + \epsilon_{t-1} \quad (14)$$

where \tilde{w} is an observed option premium $w = V/C - 1$ when $D^+ = 0$ and is latent when $D^+ = 1$. Consider a round t in which $D^+ = 0$, and let N_t be the number of consecutive periods preceding t in which $D^+ = 1$. Then we can rewrite (14) entirely in terms of observables:

$$\tilde{w}_t - \tilde{w}_{t-N_t-1} = \nu[N_t + 1] + \delta_- \sum_{i=t-N_t}^{t-1} [\pi_i^o - \pi_i] + \delta_+ [\pi_{t-N_t-1}^o - \pi_{t-N_t-1}] + \sum_{k=t-N_t-1}^{t-1} \epsilon_k \quad (15)$$

We can estimate (15) conveniently using a standard weighted least squares package on the subset of observations for which $D_t^- = 1$. First, specify the initial value of the option premium, w_0 , say by setting it equal to the first observed $V/C - 1$ and dropping previous data. Then form the dependent variable $w_t - w_{t-N_t-1}$ and estimate

$$w_t - w_{t-N_t-1} = \nu[N_t + 1] + \delta_- \sum_{i=t-N_t}^{t-1} [\pi_i^o - \pi_i] + \delta_+ [\pi_{t-N_t-1}^o - \pi_{t-N_t-1}] + \zeta_t \quad (16)$$

where ζ_t is distributed $N(0, \sigma^2 N_t)$.

Estimates with clustering at the subject level are reported in Table 5 under column (1). Note that, for ease of reading, Table 5 reports the effects on option premia of having foregone 1000 points of potential earnings in the previous round. In this specification there is no evidence of autonomous trend ($p = 0.382$), indicating that our specification is reasonable. Our results indicate that ex post underinvestment has a significant positive effect on option premia while ex post overinvestment has a marginally significant ($p=0.051$) negative effect. The point estimates indicate a stronger effect per dollar foregone by investing at too high a premium than a dollar foregone by investing at too low a premium. Wald tests, however, do not allow us to reject the hypothesis that these two effects are equal. In a second specification (2), we drop the insignificant autonomous trend variable and get even more disparate per-dollar effects (still not significantly different in absolute value).

Finding 5 *Subjects adjust their option premium adaptively in response to foregone earnings. When subjects forego earnings by investing too early relative to the ex post optimum, they increase their premium in the following period. When subjects forego earnings by losing the opportunity to invest, they decrease option premia in the following period. We find no evidence of an autonomous trend (e.g. non-adaptive learning) in our data.*

5.5 Implications of the Learning model

To what degree can the learning model reported above explain patterns observed in the data? We conduct simulations using fits from the learning model and initial observed option premia (w_0 as above) as seeds. Each subjects' initial option premium decision was used 20 times, yielding a total of 1380 simulations. In each simulation, 80 unique periods were generated and option premia were adjusted according to (12) using the coefficients in Table 5, specification (2). Means of PL estimates for each treatment and block of the simulated data are presented in Table 6.

Simulated data match experimental data in several respects. First, option premia, which are initially low, increase over time towards optimality. Moreover, as in the experiment, the rate

Coefficient	(1)	(2)
$\nu \times 1000$	-1.838 (2.089)	
$\delta_- \times 1000$	-0.6179 (0.310)*	-0.7502 (0.223)***
$\delta_+ \times 1000$	0.4402 (0.0711)***	0.428 (0.0642)***

Table 5: Estimated effects of losses on subsequent option premium choices. The coefficients $(\nu, \delta_-, \delta_+)$ of (12) respectively represent an autonomous trend and reactions to ex-post losses from low investment and from not investing. One, two and three asterisks denote coefficients significantly different from zero at the ten percent, five percent and one percent levels respectively.

and overall magnitude adjustment depends strongly on the treatment. Option premia in the Low treatment change only a very small amount over 80 periods, while premia in the High treatment increase a great deal. Indeed, the difference between option premia in the first and last time block is many times greater in the High treatment than in the Low treatment in both the experimental and simulated data. Likewise, changes in behavior in the two Medium treatments are similar to one another and distinct from the rates of change observed in Low and High treatments. Simulations also replicate the increasing heterogeneity in the High treatment: The standard errors in the last block are twice those in the previous block and both are far higher than those in other treatments. Finally, simulated data generate overall means that match quite well with those observed in the experiment.

Simulations also suggest that subjects approach optimal behavior after 60 to 80 periods. Is the steady state of the learning process nearly optimal as well? We ran a second set of simulations to find out.⁵ For each treatment we conducted 100 unique simulations, each lasting 300 periods. Agents began the simulation utilizing an option premium of 0.1 and adjusted the premium according to (12) using the coefficients in Table 5, specification (2). The resulting data is binned at the end of each set of 20 periods and the means over simulations are plotted (as Actual Parameters) for each treatment in Figure 6. Horizontal lines denote optimal option premia. In Low and Medium treatments the process converges to virtually optimal decision rules. In the High treatment convergence is to a level slightly above the optimum but still close.

The parameter values in Table 5 are biased in that the absolute value of δ_+ is greater than

⁵We also include some technical notes on the steady state in the Appendix A.

Block of Periods					
	1-20	21-40	41-60	61-80	Overall
Means					
Low	0.154 ± 0.001	0.161 ± 0.001	0.163 ± 0.001	0.164 ± 0.001	0.160 ± 0.000
Medium A	0.231 ± 0.002	0.355 ± 0.002	0.426 ± 0.002	0.449 ± 0.003	0.365 ± 0.001
Medium B	0.223 ± 0.002	0.348 ± 0.002	0.423 ± 0.003	0.457 ± 0.003	0.362 ± 0.002
High	0.410 ± 0.004	0.606 ± 0.005	0.698 ± 0.005	0.775 ± 0.017	0.625 ± 0.006

Table 6: Product-Limit estimates of simulation option premia from by time block and treatment. Standard errors are reported in smaller font following the \pm symbol.

δ_- . Subjects adjust their behavior more when they lose a dollar by waiting too long (in an *ex post* sense) than they do when they lose a dollar by not waiting long enough. To what degree does near-optimal convergence depend on this sort of bias in the learning process? One way to answer this question is to examine the steady state under *non-biased* learning. We repeat the simulations for each treatment with agents who use even (in absolute value) learning parameters. Specifically, we set $\delta_- = -0.0005941$ and $\delta_+ = 0.0005941$, which are based on the average of the absolute value of the estimated learning parameters. Results are binned at the end of each block of 20 periods and plotted in Figure 6 as the series “Even Parameters”. In all treatments other than Low, non-biased learning converges to levels which are much further away from the optimum than the observed biased learning. The bias subjects incorporate in their learning is therefore remarkably well suited to finding optimal decision rules under wide the range of parameters examined in our experiment.

Finding 6 *Estimated learning parameters imply convergence to nearly optimal investment. Biased learning of the sort observed converges to much more optimal option premia than a non-biased learning process.*

6 Discussion

Ordinary people can indeed learn to closely approximate optimal exercise of deferral options, despite their mathematical subtlety. In our laboratory experiment, 69 undergraduates each faced 80 stochastic investment opportunities governed by Brownian motion (or a close binomial approximation thereof). At first they tended to exercise the deferral option too soon, but over time their average behavior converged close to the optimum. In our Low treatment (using binomial parameters that imply an optimal premium of about 20% of investment cost) convergence was virtually

complete by the 40th opportunity. Convergence (albeit with greater heterogeneity) was almost complete by the last block of 20 opportunities in the High treatment (optimal premium about 80% of cost). In the Medium treatments (optimal premium about 50% of cost), simulations indicate that another block or two would allow subjects to achieve the optimum.

It bears emphasizing that the task was quite demanding. Several colleagues (and at least one of the coauthors!) expected that our subjects would persistently undervalue the deferral option and invest too early. It is, after all, a major disappointment when a valuable investment opportunity disappears before it is exploited, and such disappointments could easily bias people towards early investment. On the other hand, the expected profit is rather flat above the optimum in some treatments, and this could lead to the opposite bias in those treatments. The optimum differs considerably across treatments, so getting things right for the wrong reason in one treatment is not a harbinger of success in the other treatments. Actual computation of the optimum is surely beyond the powers of any of our subjects.

Nonetheless, the subjects performed quite well overall, and not just in the last few periods. Our experimental design allows us to infer the separate effects of the underlying parameters. The optimum predicts correctly the actual responses of our subjects to changes in the discount rate (or expiration probability) and correctly predicts the direction (but somewhat overpredicts the average magnitude) the overall average response to the binomial parameters controlling the volatility and the drift of the investment opportunity.

Our study provides several clues on how people learn in this demanding task. A model of directional learning suggests that they respond to ex-post errors. They tend to increase the option premium following an opportunity where that would have been more profitable, and tend to decrease the option premium following a missed opportunity. There is some evidence that they tend to respond more strongly to the second sort of ex-post error, but a remarkable finding is that, without such bias, directional learning would lead to convergence much further away from the optimum.

Are there other explanations for our findings? Risk aversion is perhaps the first standard explanation to come to mind, but the analysis in Appendix A.3 suggests that it plays at best a minor role in our task. We have not yet systematically considered quasihyperbolic discounting or other known behavioral explanations.

Our findings may shed some light on previous results from the related experimental search literature. As Gong and Ramachandran (2007) note, payoff functions in most search experiments are relatively flat on the left hand side of the optimum (corresponding to undersearch) and steep

on the right hand side (corresponding to oversearch). Gong and Ramachandran conjecture, ala Harrison (1989), that the under-searching observed in much of this literature may be related to these flat payoffs. Intuitively, subjects lose little by under-searching in typical designs and therefore are more likely to under-search than over-search. As is evident from Figure 1, our design does not suffer from such flat payoffs to the left of the optimum and yet in most treatments we observe something roughly equivalent to under-searching, (low option premia) especially in early sessions. In fact, in our High treatment (and to a lesser extent in our Medium treatments) payoffs are relatively flat *to the right* of the optimum so that payoff dominance issues problems would lead to something like over-searching (high option premia). However, we observe just the opposite in these cases. Although our experiment is only loosely related to this literature, our results suggest that search behavior in general and under-searching in particular may be driven by factors other than payoff dominance.

References

- [1] Braunstein, Y. and Schotter, A., 1981, "Economic Search: An Experimental Study." *Economic Inquiry*. 19: 1-25.
- [2] Cason, T. and Friedman, D., 1999, "Learning in a Laboratory Market with Random Supply and Demand." *Experimental Economics*. 2:277-298.
- [3] Copeland, T. and Antikarov, V., 2003, "Real Options: A Practitioner's Guide." Norton.
- [4] Cox, James C. and Ronald L. Oaxaca., 2000, "Good News and Bad News: Search from Unknown Wage Offer Distributions," *Experimental Economics*. 2: 197-225.
- [5] Dixit, A. and Pindyck, A., 1994, "Investment Under Uncertainty." Princeton: Princeton University Press.
- [6] Friedman, D. and Rust, J., 1993, "The Double Auction Market: Institutions, Theories and Evidence", Santa Fe Institute Proceedings Volume XIV, Addison-Wesley, 1993.
- [7] Gong, B. and Ramachandran, V. "Payment Dominance in Sequential Search.", 2006, manuscript.
- [8] Harrison, G., 1989, "Theory and Misbehavior of First-Price Auctions," *American Economic Review*, 79: 749-762.

- [9] Henry C., “Investment Decisions under Uncertainty: the Irreversibility Effect.”, 1974, *American Economic Review*, 64(6), 1006-1012.
- [10] Hirshleifer, J., Glazer A., and Hirshleifer D., 2005, “Price Theory and Applications”, 7th ed. Cambridge.
- [11] Kaplan, E., and Meier, P., 1958. “Nonparametric estimation from incomplete observations.” *Journal of the American Statistical Association*. 53:457-481
- [12] Kluger, B. and Wyatt, S., 2004, “Are Judgement Errors Reflected in Market Prices and Allocations? Experimental Evidence Based on the Monty Hall Problem” *Journal of Finance* 59: 969-997
- [13] Kreps, D., 1990. “A Course in Microeconomic Theory.” Princeton: Princeton University Press.
- [14] MacDonald, R. and Siegel, D., 1986, “The Value of Waiting to Invest,” *Quarterly Journal of Economics*. 101: 707-728.
- [15] Robson, A., 2002, “Evolution and Human Nature.” *Journal of Economic Perspectives*. 16: 89-106.
- [16] Selten, R. and Buchta, J., 1998, “Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions.” In D. Budescu, I. Erev, and R. Zwick (eds.), *Games and Human Behavior: Essays in Honor of Amnon Rapoport*. Mahwak, NJ: Erlaauum Ass.

A Mathematical Details

A.1 Deriving the Option Premium Formulas

Consider again the investment timing problem when the gross value of investment V is governed by the stochastic differential equation

$$dV = \alpha V dt + \sigma V dz, \tag{17}$$

where z is the standard Wiener process. Following Chapter 5 of Dixit and Pindyck (1994), we will show that expected discounted profit $E[(V - C)e^{-\rho t}]$ is maximized when the given cost C is sunk as soon as V exceeds $V^* = (1 + w^*)C$, for

$$w^* = \frac{1}{B - 1}, \text{ where } B = \frac{1}{2} - \frac{\alpha}{\sigma^2} + \sqrt{\left[\frac{\alpha}{\sigma^2} - \frac{1}{2}\right]^2 + \frac{2\rho}{\sigma^2}} > 1. \tag{18}$$

Let $F(V)$ be the value of the deferral option, i.e., the maximized value of $E[(V - C)e^{-\rho t}]$, given initial project value $V \geq 0$. Prior to investment, the Bellman equation equates the expected return on the opportunity to the expected rate of appreciation,

$$\rho F(V) dt = E_t[dF]. \quad (19)$$

The second order Taylor expansion of dF yields

$$dF = F'(V) dV + \frac{1}{2} F''(V) (dV)^2. \quad (20)$$

Expand dV using Ito's Lemma and equation (17), recalling that $E_t dz = 0$ and $E_t [dz]^2 = dt$, to obtain

$$E_t [dF] = \alpha V F'(V) dt + \frac{1}{2} \sigma^2 V^2 F''(V) dt. \quad (21)$$

Insert equation (21) into (19) and divide by dt to obtain

$$\frac{1}{2} \sigma^2 V^2 F''(V) + \alpha V F'(V) - \rho F(V) = 0. \quad (22)$$

We now derive the value function $F(V)$ and the optimal threshold value V^* by solving the second order ordinary differential equation (22) subject to the boundary conditions

1. $F(0) = 0$,
2. $F(V^*) = V^* - C$, and
3. $F'(V^*) = 1$.

The first boundary condition is implied by geometric Brownian motion; by (17), if $V(0) = 0$ then $V(t) = 0$ for all $t > 0$. The second condition is called value matching: if the initial value of the project makes it worthwhile to launch immediately, then the realized value is simply that value less the cost.⁶ The third condition is called smooth pasting. It rules out a kink in the Bellman value function F at the threshold, which (it can be shown) would permit profitable arbitrage at points arbitrarily close to V^* .

To solve the problem, suppose it has a solution of the general form $F(V) = AV^B$. Insert this into (22) and cancel the common factor AV^B to obtain the quadratic equation

$$\frac{1}{2} \sigma^2 B(B - 1) + \alpha B - \rho = 0. \quad (23)$$

⁶Actually, the condition says more. When she invests, the investor gains V but loses the deferral option $F(V)$. Thus she should set the threshold V^* so that the gain net of opportunity cost, $V - F(V)$, is just equal to the out-of-pocket cost C .

A straightforward calculation shows that the larger root of (23) is

$$B = \frac{1}{2} - \frac{\alpha}{\sigma^2} + \sqrt{\left[\frac{\alpha}{\sigma^2} - \frac{1}{2}\right]^2 + \frac{2\rho}{\sigma^2}} > 1. \quad (24)$$

To complete the derivation, substitute the general solution $F(V) = AV^B$ into boundary condition 2 and re-arrange slightly to obtain

$$C = (1 - AV^{B-1})V. \quad (25)$$

Inserting the general solution in boundary condition 3 and rearranging yields $AV^{B-1} = 1/B$. Substituting this last expression into (25) yields $C = (1 - 1/B)V$ which is easily rearranged to obtain the desired formula

$$V^* = \left(1 + \frac{1}{B-1}\right)C. \quad (26)$$

To clean up loose ends, note that one can obtain an explicit expression for A and thus the deferral option value $F(V) = AV^B$ by rewriting the second boundary condition as $A = (V^* - C)/(V^*)^B$ and inserting the expression (26) for V^* . Boundary condition 1 is used to show that the smaller (negative) root of the quadratic equation (23) is irrelevant and that $F(V) = AV^B$ is indeed the general solution of (22). That second order differential equation would have a unique solution given only one other fixed boundary condition, but since V^* is endogenous (i.e., is a free boundary) we need a third condition to determine a specific solution.

A.2 Comparative Statics

To obtain comparative statics for the option premium factor $w^* = \frac{1}{B-1}$, first note that $dw^*/dB = -1(B-1)^{-2} < 0$. Then implicitly differentiate the quadratic equation (23) with respect to each Brownian parameter and solve to obtain the derivatives of B . For example, $\partial B/\partial \rho = 1/D$, where $D = \alpha + \frac{1}{2}\sigma^2(2B-1) > 0$.⁷ Similarly $\partial B/\partial \alpha = -B/D$ and $\partial B/\partial \sigma = -\sigma B(B-1)/D$.

Modify equations (3-5) describing the relation between the binomial and Brownian parameters by fixing Δt at a small positive value, and differentiate to obtain

$$\partial \alpha / \partial h \approx (2p - 1) / \Delta t; \quad \partial \alpha / \partial p \approx 2h / \Delta t; \quad \partial \alpha / \partial q = 0 \quad (27)$$

$$\partial \sigma / \partial h \approx 2\sqrt{p(1-p)} / \Delta t; \quad \partial \sigma / \partial p \approx (1-2p)h / \sqrt{p(1-p)\Delta t}; \quad \partial \sigma / \partial q = 0 \quad (28)$$

$$\partial \rho / \partial h = 0; \quad \partial \rho / \partial p = 0; \quad \partial \rho / \partial q = 1 / [(1-q)\Delta t]. \quad (29)$$

⁷Of course, $D < 0$ when the drift parameter α is too negative, i.e., when $\alpha < -\frac{1}{2}\sigma^2(2B-1)$. However, in that case it can be shown that the classic present value rule holds: launch immediately whenever $V \geq C$. Hence $D > 0$ whenever the option premium formula (18) is valid.

Thus

$$\frac{\partial w^*}{\partial h} = \frac{dw^*}{dB} \left[\frac{\partial B}{\partial \alpha} \frac{\partial \alpha}{\partial h} + \frac{\partial B}{\partial \sigma} \frac{\partial \sigma}{\partial h} \right] \approx \frac{[2p - 1 + (B - 1)\sigma\sqrt{\Delta t}]B}{(B - 1)^2 D \Delta t} > 0, \quad (30)$$

$$\frac{\partial w^*}{\partial p} = \frac{dw^*}{dB} \left[\frac{\partial B}{\partial \alpha} \frac{\partial \alpha}{\partial p} + \frac{\partial B}{\partial \sigma} \frac{\partial \sigma}{\partial p} \right] \approx \frac{2hB}{(B - 1)^2 D \Delta t} > 0, \text{ and} \quad (31)$$

$$\frac{\partial w^*}{\partial q} = \frac{dw^*}{dB} \left[\frac{\partial B}{\partial \rho} \frac{\partial \rho}{\partial q} \right] = \frac{-1}{(B - 1)^2 D (1 - q) \Delta t} < 0. \quad (32)$$

A.3 Risk Aversion

MacDonald and Siegel (1986) argue that the impact of risk aversion is that the discount rate ρ in the investment timing problem will increase by a risk premium $\lambda > 0$. What are the observable implications of a shift from ρ to $\rho + \lambda$?

From equation (5) we have $q = 1 - \exp(-\rho\Delta t)$. Define $\hat{q} = 1 - \exp(-(\rho + \lambda)\Delta t)$, and Taylor expand the difference to obtain $\hat{q} - q = [(\rho + \lambda)\Delta t - \rho\Delta t] + [(\rho + \lambda)^2\Delta t^2 - \rho^2\Delta t^2]/2 + \dots = \lambda\Delta t - (\rho\lambda + 0.5\lambda^2)\Delta t^2 + \dots \approx \lambda\Delta t$.

The conclusion is that shifting ρ by $\lambda > 0$ has the same observable effect as replacing q by $q + \lambda\Delta t$. Up to first order, the theoretical premiums w_τ^* in each treatment τ would decrease by $|\frac{\partial w^*}{\partial q}| \lambda \Delta t$.

Non-trivial values of $\lambda > 0$ will drive apart the two Medium treatment values of w^* , weakening Finding 1. One can find values for λ that improve Finding 2 for one treatment, but they degrade it for other treatments. Finding 3 is not appreciably changed by assuming $\lambda > 0$, and such an assumption can not explain the dynamics summarized in Findings 4-6. Thus allowing for an additional free parameter $\lambda > 0$ does not help account for the observed data.

A.4 Optimal Premium versus Steady State Premium

Let $H(v)$ with density $h(v)$ be the distribution function for the maximum value of the investment $v = \sup_{t>0} V(t)$. The dependence of H on treatment τ is suppressed to streamline notation. Assume that the investor uses a threshold of the optimal form $V(w) = (1 + w)C$ but that $w > 0$ is not necessarily the optimal premium w^* given by equation (18). Then the actual (undiscounted) profit will be wC if $v \geq (1 + w)C$, an event with probability $1 - H(V(w))$, and otherwise profit will be 0. Hence expected profit (as simulated in Figure 1) will be $G(w) = wC(1 - H(V(w)))$.

The optimal premium then can be written as $w^* = \operatorname{argmax} G(w)$. The associated first order

condition is $0 = 1 - H(V(w^*)) - w^*Ch(V(w^*))$, using $V' = C$. An alternative expression is $w^* = \text{argmin}EL(w)$, where the expected loss $EL(w)$ sets $\delta_+ = \delta_- = 1$ in the expression $\delta_+ \int_C^{V(w)} (\pi^o - 0)h(v)dv + \delta_- \int_{V(w)}^\infty (\pi^o - wC)h(v)dv$. Noting that the ex-post optimal profit is always $\pi^o = v - C$, the expression becomes

$$\delta_+ \int_C^{V(w)} (v - C)h(v)dv + \delta_- \int_{V(w)}^\infty (v - V(w))h(v)dv. \quad (33)$$

It is straightforward to verify for $\delta_+ = \delta_- > 0$ that the expression (33) has the same first order condition as before, $0 = 1 - H(V(w^*)) - w^*Ch(V(w^*))$.

The steady state of the learning model in section 5.4, by contrast, does not involve taking the derivative of (33). Instead, for given $\delta_+ > 0 > \delta_-$, the steady state premium \hat{w} equates (33) to zero. One can see that if the density h is unimodal and approximately symmetric around w^* , and if the δ coefficients have approximately equal magnitude with opposite signs, then $\hat{w} \approx w^*$. Figure 1, however, shows that h is positively skewed in the Medium and High treatments, so here the steady state premium \hat{w} exceeds the optimal premium w^* .

B Econometric Details

B.1 The Product-Limit Estimator

The product-limit estimator produces an estimate of the distribution function, $F(w_i) = \text{prob}(x \leq w_i)$ while taking account of random right censoring. Consider a sample consisting of n observations. In uncensored observations denote w as the observed option premium decision ($w = V/C$) and in censored cases denote w as the highest premium available before censoring ($w = v/C$). We can then construct an n -vector of these option premia $w = (w_0, w_1, \dots, w_n)$ ordered so that $i < j$ if and only if $w_i < w_j$. Include in this vector $w_0 = 1$, the lowest premium possible at investment.

The product-limit estimate of $F(w_i)$ exploits the fact that the complement of the distribution function can be written as a product of conditional probabilities. Note that $1 - \text{prob}(x \geq w_i) = 1 - \text{prob}(x \geq w_i | x \geq w_{i-1}) \times \text{prob}(x \geq w_{i-1})$. Recursively then,

$$F(w_i) = 1 - \prod_{j=1}^i \text{prob}(x \geq w_j | x \geq w_{j-1}) \quad (34)$$

Let c_i denote the number of censored observations smaller than w_i and let u_i denote the number of uncensored observations smaller than w_i . Finally, define $n_i = n - c_i - u_i$, the number of

observations equal to or greater than w_i . The product-limit estimate of $p(x \geq w_i | x \geq w_{i-1})$ is the proportion of investments greater than w_{i-1} which are also greater than w_i :

$$\hat{p}(x \geq w_i | x \geq w_{i-1}) = \frac{n_i - u_i}{n_i} \quad (35)$$

The product-limit estimator is then, following (34), the cumulative product of these individual conditional probabilities at each w_i

$$\hat{F}(x) = 1 - \prod_{w_i < x} \frac{n_i - u_i}{n_i} \quad (36)$$

Without censoring (that is when all w_i denote option premia at investment) it can be shown that $\hat{F}(w_i)$ is simply the empirical distribution function – the proportion of investments which are lower than w_i for each w_i . Kaplan and Meier (1958) show that the product-limit estimator is the maximum likelihood *non-parametric* estimator of the distribution function in environments with censoring problems analogous to ours.

Typically, hypothesis tests comparing product-limit distribution functions are conducted using a log-rank test. Consider two samples, labeled $j = 1, 2$. In what follows we will sometimes pool the two samples and construct the statistics described above, in which case we omit the j subscript. In other cases we'll use statistics described above computed only for pool j in which cases we include the subscript. That implies, for instance, that $n_i = n_{i1} + n_{i2}$. We'll also introduce d_i , defined as the total number of investments made at premium w_i where $d_i = d_{i1} + d_{i2}$.

Under the hypothesis that the two samples are the same at w_i , expected investments in group 1 are $n_{i1}d_i/n_i$ while the actual observed investment is simply d_{i1} . As with many non-parametric techniques, the log-rank test relies on a test statistic based on the difference between these observed and expected statistics. To construct the test statistic, the log-rank test computes the hypergeometric variance for the number of investments at premium w_i as

$$V_i = \frac{n_{i1}n_{i2}(n_i - d_i)d_i}{n_i^2(n_i - 1)} \quad (37)$$

The test statistic for the log rank test is then

$$z = \frac{\sum_i^n (d_{i1} - \frac{n_{i1}d_i}{n_i})}{\sqrt{\sum_i^n V_i}} \quad (38)$$

which is approximately distributed standard normal under the hypothesis that the hazard rates for the two samples are equal .

B.2 Parametric Robustness Check

As a robustness check, we conduct a parametric analysis of the data using a tobit models with censoring at the highest value drawn before the end of the round. This analysis is then used to test our three hypotheses and test the robustness of our first three findings. A methodological problem with Tobit analysis in our setting is that, in our data, the value of the censoring point, \bar{V} , varies across observations making direct Tobit estimation across rounds and sessions impossible. Our workaround is to estimate the relationship $V/\bar{V} = (1+w)C/\bar{V}$ which has a common censoring point across sessions and rounds of 1. If V is linearly related to C by a slope term β , $\frac{V}{\bar{V}}$ is related to $\frac{C}{\bar{V}}$ by the same slope term. After coding censored observations as $V = 1$ we can estimate the following Tobit model:

$$\frac{\hat{V}_{ijk}}{\bar{V}_{ij}} = \beta \frac{C_{ij}}{\bar{V}_{ij}} + \epsilon_{ijk} \quad (39)$$

where the value of V we actually observe in the data is

$$V_{ijk} = \begin{cases} \frac{\hat{V}_{ijk}}{\bar{V}_{ij}} & \text{if } \frac{\hat{V}_{ijk}}{\bar{V}_{ij}} < 1 \\ 1 & \text{if } \frac{\hat{V}_{ijk}}{\bar{V}_{ij}} > 1 \end{cases} \quad (40)$$

Subscript i , j and k index session, period and subject respectively. The disturbance term ϵ_{ijk} is assumed to be normally distributed with a mean μ_ϵ and variance σ_ϵ . From the estimates, the options premium is $w = \beta - 1$.

We estimate this model separately for each subject and study the distribution of premium estimates across subjects. Although it would be more efficient to pool subjects and use a random effects estimator, such a model will necessarily be misspecified. Because project values begin at cost, and the brownian motion in our environment is geometric, a non-zero value of investment at a zero cost is not possible. It is therefore important that the intercept be constrained at zero. Random effects estimators, however, can only constrain the mean of the intercept to zero.⁸ An

⁸Perhaps more importantly, our main interest is the variation in the slope term across subjects, and random coefficient Tobit models are not available.

Treatment	Tobit		RE Tobit	
	Pooled	By Treatment	Pooled	By Treatment
High	0.411 (0.007)	0.452 (0.010)	0.353 (0.010)	0.407 (0.013)
Low	0.199 (0.005)	0.170 (0.003)	0.226 (0.008)	0.168 (0.007)
Medium A	0.293 (0.006)	0.284 (0.005)	0.274 (0.015)	0.266 (0.009)
Medium B	0.289 (0.006)	0.289 (0.006)	0.274 (0.011)	0.264 (0.010)

Table 7: Estimates from the random effects tobit model described in (39) and (40). Standard errors are in brackets.

alternative is to estimate a standard tobit on pooled data with an omitted intercept, though this introduces repeated measure problems. We will include estimates of both tobit and random effect tobits pooled across subjects in our analysis as robustness checks, with the caveat that both may be problematic.

The by-subject Tobit estimates resemble PL estimates and lend further support to Finding 1. Figure 7 shows histograms of individual subject premium estimates for each treatment. Note the strong similarities to the by-subject PL estimates presented in Figure 5. The modal premium is similar in the two Medium treatments, lower in the Low treatment and higher in the High treatment. A Jonckheere-Terpstra test allows us to reject the hypothesis that estimated premiums fail to ordinally rank as predicted across treatments ($p = 0.000$). We further use Mann-Whitney tests to test for treatment differences in pairwise comparisons. We are unable to reject this hypothesis when comparing the two Medium treatments ($p = .90$). In all other pairwise comparisons, we can reject this hypothesis at the one percent level.

We find further support for Finding 1 using both standard and random-effect tobit models with data pooled across subjects. Because there is strong by-treatment heteroskedasticity (evident in Figure 3), we include specifications in which models are estimated separately by treatment (called “By Treatment” in Table 7). In the case of the standard tobits, the specification is as described in (39) and (40). In RE specifications, to (39) we add a random effect, u_k , on subjects which is assumed to be normally distributed with a mean μ_u and variance σ_u . Results are displayed in Table 7. In all specifications coefficient estimates rank as predicted. Likelihood ratio tests (in By Treatment

specifications) and Wald tests (in Pooled specifications) allow us to reject the hypotheses (at the one percent level) that Low, Medium and High coefficients are identical in pairwise comparisons. Tests in the By Treatment specifications also indicate a statistical difference between the two Medium treatments at the five percent level, though these differences are, economically, very small. In Pooled specifications we cannot reject the hypothesis that premiums in the Medium treatments are different ($p = 0.587$ for the standard Tobit and $p = 0.999$ in the random effects Tobit).

Results of by-subject tobit models also support Finding 2. In Figure 7, virtually all subject estimates are close to (in the same bin as) predictions in the Low treatment while most lie well below predictions in other treatments. Median premium estimates are 0.17 in the Low treatment (compared to $m_{Low}^* = 0.18$), 1.28 in Medium A ($m_{MedA}^* = 0.49$), 0.30 in Medium B ($m_{MedB}^* = 0.5$) and 0.47 in the High treatment ($m_{High}^* = 0.80$). Sign tests⁹ allow us to reject the hypothesis that $m_t = m_t^*$ at the one percent level in High, MediumA and MediumB treatments. We cannot reject the same hypothesis using such a test in the Low treatment ($p=.629$).

Finding 2 is also supported by estimates on data pooled over subjects. As is clear from Table 7, premium estimates for Medium and High treatments are dramatically lower than m_t^* under all pooled specifications. Wald tests for all specifications show that these differences from predictions are significant at the one percent level. Estimates for the Low treatment are economically similar though statistically different from the predicted premium of 0.18 in all specifications (by Wald tests, all at the one percent level). These statistical differences stand as poor evidence against Finding 2, however, since in half of the specifications estimates are greater than predicted and in half they are lower. Overall pooled tobit estimates correspond well with PL estimates supporting our second finding.

Finally, using the pooled Tobit results, we construct estimates of the impact of binomial parameter changes analogous to those constructed in support of Finding 3. These estimates are presented in Table 8. In all specifications, results support Finding 3. Moreover, in all but the pooled random effects specification, the magnitude of these effects match those from PL estimates, which are reported in Table 3. In the pooled random effects specification, the magnitude of the effects of the h and q parameters are somewhat smaller than those generated using either PL estimates or estimates from other Tobit specifications.

⁹We use the sign test instead of the more powerful Wilcoxon sign rank test because observations appear to be from non-symmetric parent distributions. All results are based on two sided tests.

Parameter	Predicted	Tobit		RE Tobit	
		Pooled	By Treatment	Pooled	By Treatment
Δp	0.003	-0.014 (0.012)	-0.024 (0.013)*	-0.015 (0.023)	-0.023 (0.020)
Δh	0.317	0.104 (0.012)*	0.143 (0.013)*	0.064 (0.023)*	0.118 (0.020)*
Δq	-0.308	-0.108 (0.012)*	-0.139 (0.013)*	-0.064 (0.023)*	-0.121 (0.020)*

Table 8: Binomial parameter impact estimates using pooled Tobit models. Predicted values are from the bottom row of Table 1. Estimates are obtained from equations (9-11) applied to Tobit estimates reported in t Table 7. Standard errors are in parentheses.

C Instructions

The following instructions were handed out to subjects and read aloud. A demonstration of the experimental software was projected on a screen during the instructions to help with comprehension. Actual parameters for the particular treatment were written on a white board and referred to during the instructions. Figure 1 in the instructions was similar to Figure 2.

Instructions for Investment Timing Game You are about to participate in an experiment in the economics of decision-making. The National Science Foundation and other agencies have provided the funding for this project. If you follow these instructions carefully and make good decisions, you can earn a CONSIDERABLE AMOUNT OF MONEY, which will be PAID TO YOU IN CASH at the end of the experiment. Your computer screen will display useful information. Remember that the information on your computer screen is PRIVATE. To insure best results for yourself and accurate data for the experimenters, please DO NOT COMMUNICATE with the other participants at any point during the experiment. If you have any questions, or need assistance of any kind, raise your hand and one of the experimenters will come. In the experiment you will make investment decisions over several rounds. At the end of the last period, you will be paid \$5.00, plus the sum of your investment earnings over all rounds. **The Basic Idea.** Each round you will decide when (if ever) to seize an investment opportunity. At the beginning of the round you will be assigned a cost, C of investing. The value V of investing will change randomly over time. You earn V-C points if you seize the opportunity before it disappears. If you wait longer, V might go higher,

earning you more points. Or V might go lower. The opportunity to invest might evaporate before you seize it, in which case you earn 0 points that round. Investor Screen Information. Your cost C is shown on your screen as a horizontal red line, as in Figure 1. The value V of the investment is shown as a jagged green line that scrolls from left to right, with the rightmost tip (the leading edge) representing the current value V . Previous values move left, as on a ticker tape. At the start of each round the value line V starts at your cost line C and randomly evolves from there. There is an Invest! button located at the bottom of the screen. Point the cursor arrow at the Invest! button and click when you want to invest. When the investment opportunity evaporates, that button will disappear, replaced by a gray message. Other useful messages appear in the window to the right labeled Your Performance. For example, in Figure 1, that window tells you the round number, your cost, and that you have no competitors for the investment opportunity. You can see messages and results from previous rounds by clicking the Previous button at the top of the window. The round will continue until the investment opportunity disappears even if you have already invested. Payment. Points translate into dollars according to a formula written on the board. You will be paid in cash at the end of the experiment for the points earned in all rounds plus the \$5 show-up fee. For example, if the formula is \$0.02 per points in excess of 1000, and if you earn 1682 points, then your cash payment is $\$5.00 + (1682 - 1000) \times 0.02 = \$5.00 + \$13.64 = \18.64 . Details. In case you want to know, here are a few details of how V unfolds. You can skip these if you prefer to learn just from experience.

- The round is a series of many ticks (e.g., 5 ticks per second).
- Each tick the value V moves randomly up or down by a fixed percentage, e.g., 3%.
- Upticks are slightly more likely than downticks, e.g., each tick is up with probability 51% or down with probability 49%.
- The round ends (the investment evaporates) with a small probability each tick, e.g., ? of 1%.
- The actual values (for ticks per second, tick size, uptick probability, and evaporation probability) will be written on the board before the experiment begins.
- The value always starts at the cost line, e.g. at $C = V = 20$.
- The computer will not allow you to seize the investment opportunity when V is less than C , because that would give you a negative number of points.

Frequently Asked Questions Q1. Is this some kind of psychology experiment with an agenda you haven't told us? Answer: No. It is an economics experiment. If we do anything deceptive, or don't pay you cash as described, then you can complain to the campus Human Subjects Committee and we will be in serious trouble. These instructions are meant to clarify how you earn money, and our interest is in seeing how people make investment timing decisions.Q2. How long does a round last? Is there a minimum or maximum? Answer: The length of time is random. In the example, the probability is 0.005 that any tick is the last, and there are 5 ticks per second. In this case, the average length of a round is 200 ticks or 40 seconds. Many rounds will last less than the average, and a few will last much longer. Rounds longer than 7 minutes are so unlikely that you probably will never see one. The minimum length is one tick, but it is unlikely you will ever see a round quite that short!Q3. How many rounds will there be? Answer: Lots. We aren't supposed to say the exact number, but there will be dozens and dozens of rounds.Q4. Are there patterns in upticks and downticks? Answer: No. We've tried very hard to make it random. No matter what the recent history of upticks and downticks, the probability that the next tick is up is always the same (and is written on the board).




Figure 1: Cumulative payoffs by premium based on 1000 simulations of 80 periods. Curves are fits to means and error bars are set at one standard error. Vertical lines mark the theoretical optimum.




Figure 2: Computerized display used in the experiment. The green jagged line shows present (Time =0) and previous values of V in the current period. The subject invests by clicking the button near the bottom of the screen. The window on the right displays current period information, and toggles to summarize results from previous periods.




Figure 3: Observed investment values with linear fits to the data.




Figure 4: Product-Limit estimated distributions of option premium by treatment. Vertical dashed lines show predicted factors in each treatment.




Figure 5: Histograms of means of product-limit estimates of option premia estimated separately by individual subject. Stars indicate the histogram bins containing predicted option premia.




Figure 6: Results from simulations of directional learning using observed learning parameters (Actual parameters) and using a non-biased set of parameters (Even parameters). Each series is composed of means over 100 simulations and each observation represents the average over the preceding 20 periods.




Figure 7: Distribution of option premia estimates from the Tobit model described in (39) and (40) estimated on each individual subject. Stars indicate the histogram bins containing predicted option premia.