THE REVIEW OF ECONOMIC STUDIES

Review of Economic Studies (2009) **76**, 1103–1124 0034-6527/09/00000000$02.00

Learning to Wait: A Laboratory Investigation

RYAN OPREA
University of California

DANIEL FRIEDMAN
University of California

and

STEVEN T. ANDERSON
U.S. Geological Survey

First version received August 2007; *final version accepted October* 2008 (*Eds.*)

Human subjects decide when to sink a fixed cost C to seize an irreversible investment opportunity whose value V is governed by Brownian motion. The optimal policy is to invest when V first crosses a threshold $V^* = (1 + w^*)C$, where the wait option premium w^* depends on drift, volatility, and expiration hazard parameters. Subjects in the Low w^* treatment on average invest at values quite close to optimum. Subjects in the two Medium and the High w^* treatments invested at values below optimum, but with the predicted ordering, and values approached the optimum by the last block of 20 periods.

1. INTRODUCTION

The proper timing of investment is tricky. Only with the work of Henry (1974), and better known expositions such as MacDonald and Siegel (1986) and Dixit and Pindyck (1994), did economists come to realize that the classic net present value rule needs major modifications in the presence of irreversibility, new information, and the opportunity to wait. A highly mathematical and computational theory of real options blossomed in the 1990's, featuring the wait option,[1] and also covering the options to abandon, to expand, to contract, and to switch, as well as assorted compound options (*e.g.*, Copeland and Antikarov, 2003).

Does real options theory describe actual investment behaviour? Evidence so far is surprisingly slender. Moel and Tufano (2002) remark "empirical research on real options has lagged considerably behind [the] conceptual and theoretical contributions". The handful of empirical articles to date face formidable obstacles: key variables are unavailable (often unobservable) and the data permit only very indirect tests of the theory.[2] Laboratory studies are a natural way to help fill the empirical gap. All relevant variables are not only observable but also controllable, permitting small scale but direct tests of the theory.

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

2. The main empirical studies in this area include Paddock, Siegel and Smith (1988), Quigg (1993), Moel and Tufano (2002) and Berger, Ofek and Swary (1996).

In this paper, we study wait options in the laboratory. Although there are distantly related studies on sequential job search,[3] ours seems to be the first laboratory study of wait options or indeed of any real options.

Our conjecture is that ordinary people can closely approximate optimal exercise of wait options if they are given a decent chance to learn from personal experience. The conjecture is as much about learning as it is about real options. There are some mathematically trivial tasks that seem quite difficult to learn (*e.g.*, optical illusions, or the three door problem, cf. Kluger and Wyatt, 2004) and conversely, there are some mathematically intractable tasks (such as equilibrium bidding in a continuous double auction market, *e.g.*, Friedman and Rust, 1993) that even children learn quite quickly. Investigating the conjecture may shed light on learnability issues.

The next section lays the theoretical foundations. After describing the basic investment timing task, it presents the known formula for the optimal wait option premium w^* as a function of exogenous Brownian parameters representing the volatility and trend of the investment value and the discount rate. The section also explains binomial approximations of the continuous stochastic process and their parameterization.

Section 3 obtains three testable predictions based on the conjecture that actual behaviour will approximate optimal behaviour. It also 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 $w^* \approx 0{\cdot}2$, two Medium treatments use rather different combinations of parameter values that lead to $w^* \approx 0{\cdot}5$, and the High treatment leads to $w^* \approx 0{\cdot}8$. The treatments form a fractional factorial design that allows us to infer the separate effects of the binomial parameters.

Section 4 collects results. Broadly speaking, they support the conjecture, but there are interesting qualifications. Behaviour converges quickly to the optimum in the Low treatment, but subjects tend towards impatience in early periods of other treatments. A proximate cause is under-response to changes in two of the binomial parameters. However, over time subjects attach greater weight to the option value and their behaviour approaches optimality in all treatments. A directional learning model suggests that subjects react reliably to *ex post* losses due to early investment and react more strongly (and more heterogeneously) to missed investment opportunities. Simulations show that the asymmetric reactions allow the learning process to converge on a nearly optimal steady state.

Appendix A contains technical material on the econometric procedures. Appendix B reproduces the instructions to subjects. A self-contained theoretical derivation and supplementary econometric results can be found in our working paper, Oprea, Friedman and Anderson (2007).

2. THEORETICAL BACKGROUND

A risk neutral 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 $\sigma > 0$:

$$dV = \alpha V dt + \sigma V dz, \qquad (1)$$

where z is the standard Wiener process. That is, the value follows a continuous time random walk in which the appreciation rate has mean α and S.D. σ per unit time. 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

3. This literature is quite extensive, beginning with Braunstein and Schotter (1981). Recent work includes Cox and Oaxaca (2000) and Gong and Ramachandran (2007).

no other payoffs. Thus, the task is to choose the investment time so as to maximize the expected payoff.

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 $w^* \geq 0$ 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 Dixit and Pindyck (1994, especially chapter 5) for an accessible derivation.

There are two barriers to laboratory implementation of the model. First, while continuous time is useful analytically, 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 two parameters corresponding to those of the Brownian process (*e.g.*, Dixit and Pindyck, 1994, 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; and
- the uptick probability $p \in (0, 1)$, *i.e.*, the probability that the next step is to $(1+h)V$ rather than to $(1-h)V$.

Second, while the model features an infinite time horizon, laboratory sessions are necessarily finite in length. We solve this problem by replacing the discount parameter, ρ with a third parameter:

- the expiration probability $q \in (0, 1)$, *i.e.*, the probability that the next step is the last and the opportunity disappears.

The deviation of the uptick probability p from 0·5, times the distance between an uptick and downtick, corresponds to the drift rate α:

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

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

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

The relation between the expiration probability $q \in [0, 1]$ and the discount parameter $\rho > 0$ is obtained as follows. By definition of discount, 1 unit value received a 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, pp. 505–506): the discount rate arises from the possibility that the opportunity will expire. If that event has probability Q per unit time then the expected value of 1 unit due a 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

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

3. LABORATORY TREATMENTS, EXPERIMENTAL DESIGN, AND TESTABLE HYPOTHESES

Taking any theory to the laboratory presents challenges. For us, the first challenge was to map the continuous time, infinite horizon model to a computerized laboratory environment that must be in discrete time (due to computers' digital internal clocks) and have a finite horizon (due to subjects' obligations outside the lab). We used the standard discrete time approximation just described, with binomial parameters $\Delta t, h, p$, and q. The expiration parameter q induces impatience as required by the model.[4] It also makes the horizon effectively finite, in the sense that q can be chosen so that the probability of a period lasting longer than, say, 5 minutes is arbitrarily small. There are alternative methods to render the horizon finite but, as noted in the concluding discussion, they are less common and on the whole more problematic in the laboratory.

A second challenge was to choose appropriate parameter values. Our choices, shown in Table 1, satisfy several different considerations. To begin with, the task must have a reasonable "look and feel" for subjects. The task seems continuous with a time step of $\Delta t = 0{\cdot}00\overline{3}$ minutes, *i.e.*, five ticks per second. In this case, an expiration probability above $q = 0{\cdot}010$ too often ends the period before the subject can get focused and below $q = 0{\cdot}002$ it takes too long to run a reasonable number of periods. Step sizes above $h = 0{\cdot}05$ 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 $p = 0{\cdot}53$.

Another consideration, also crucial, is that the treatment's binomial parameters should generate at least two and preferably three widely separated option premiums. This enables tests of the model's comparative static predictions in addition to its point predictions. Wide separation between predictions affords cleaner tests on possibly noisy data. The premiums generated by Low, Medium, and High parameters are distinct and roughly equidistant. To sharpen tests of the theory, it is also desirable to have different parameter configurations that produce the same premium in different ways, hence the two Medium treatments.

Finally, we chose our parameters to avoid flat payoffs around the optimal policy. Figure 1 shows payoffs from simulations where the agent invests at a fixed $w = 0{\cdot}1, 0{\cdot}2, \ldots, 1{\cdot}5$. Note that the simulated payoffs are indeed maximized very near the optimum $w = w^*$, but that the expected payoff becomes rather flat for $w > w^*$ in treatments with larger w^*. We chose parameter configurations to mitigate the problem; it seems infeasible to avoid flat payoff landscapes for w^* above $1{\cdot}0$.

3.1. *Testable hypotheses*

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 $V/C - 1$, where V is the investment value actually chosen in a period with realized cost C. A first qualitative hypothesis is that binomial parameters generate premiums in each treatment that ordinally match the model's prediction.

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

4. Arguably, expiration hazard is 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, Glazer and Hirshleifer, 2005, chapter 15), while evolutionary theory traces time preference back to the risk of death and other biological hazards (*e.g.*, Robson, 2002).

TABLE 1

Binomial parameter values and option premium by treatment

Treatment (τ)	Step size (h)	Uptick probability (p)	Expiration probability (q)	Option premium (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.	

Notes: 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 $\Delta t = 0{\cdot}00\overline{3}$.

A stronger hypothesis is that parameters not only induce an optimal ordering of premiums across treatments but also induce optimal valuation of the wait option in each treatment.

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

The third testable hypothesis concerns comparative statics. As explained in the next subsection, our experimental design allows estimation of the separate impact of each binomial parameter, and the predictions have wide separation and different signs across parameters. Thus, we have

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

$$\bar{h} = \frac{w^*_{High} + w^*_{MedB}}{2} - \frac{w^*_{Low} + w^*_{MedA}}{2} = 0{\cdot}317 \tag{6}$$

$$\bar{p} = \frac{w^*_{MedA} + w^*_{MedB}}{2} - \frac{w^*_{Low} + w^*_{High}}{2} = 0{\cdot}003 \tag{7}$$

$$\bar{q} = \frac{w^*_{Low} + w^*_{MedB}}{2} - \frac{w^*_{High} + w^*_{MedA}}{2} = -0{\cdot}308. \tag{8}$$

3.2. *Fractional factorial design*

The expressions in the third hypothesis exploit a classic experimental design (*e.g.*, Box, Hunter and Hunter, 1978, p. 394ff) that we employ. In a fractional factorial design, each variable is typically controlled at two levels, high (*e.g.*, $h = 0{\cdot}03$) and low (*e.g.*, $h = 0{\cdot}0155$). Our four treatments are chosen as a balanced subset of the eight possible configurations (hence half factorial) of our three variables.

The advantage of a half factorial design is that it allows us to use the entire data set to estimate the main effects of each of the parameters. Moreover, the design allows us to make these inferences without running a full factorial experiment with eight treatments.

In Table 1, each parameter assumes two values, one high and one low, and these values occur with equal frequency across the design. Parameters are also balanced in the sense that the average effect $\bar{j}$ of a binomial parameter $j = h, p, q$ 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. The corresponding predicted values are shown at the bottom row of Table 1. Notice that hypothesis 3 predicts only a small impact for the uptick probability parameter p and predicts that the other two binomial parameters will have effects of roughly equal magnitude but in opposite directions.

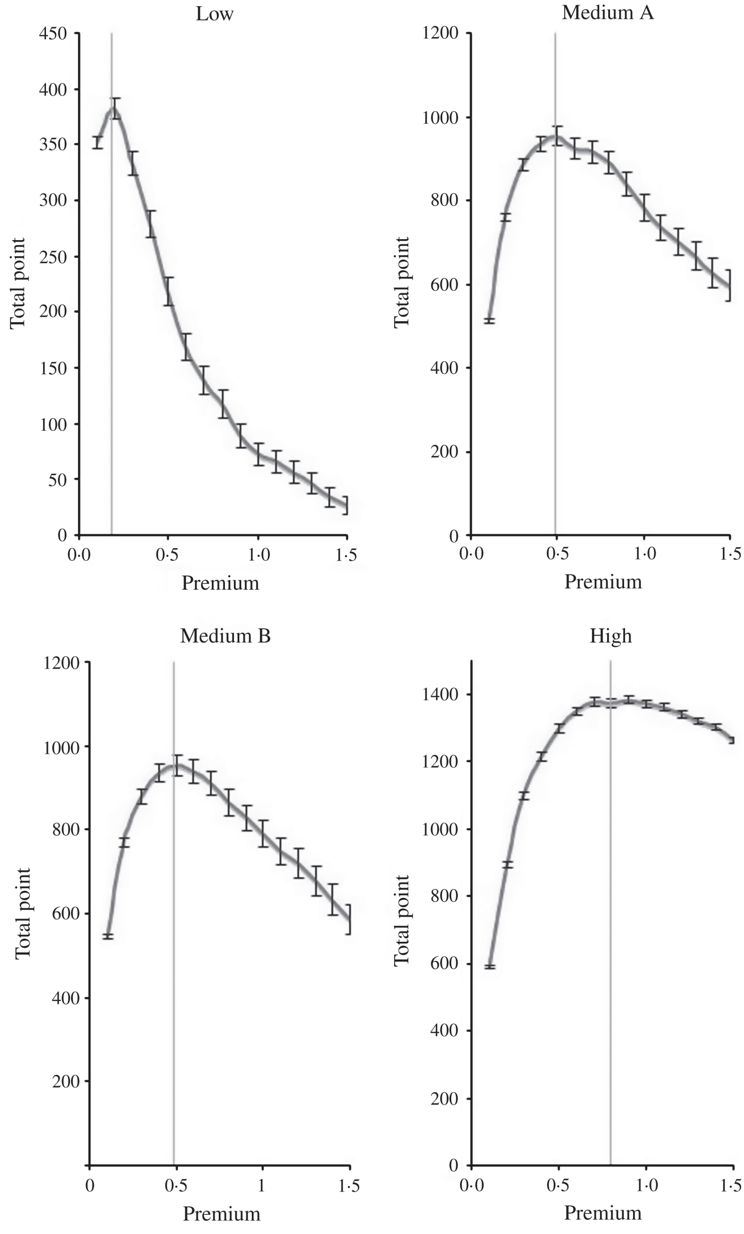

FIGURE 1

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

3.3. *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

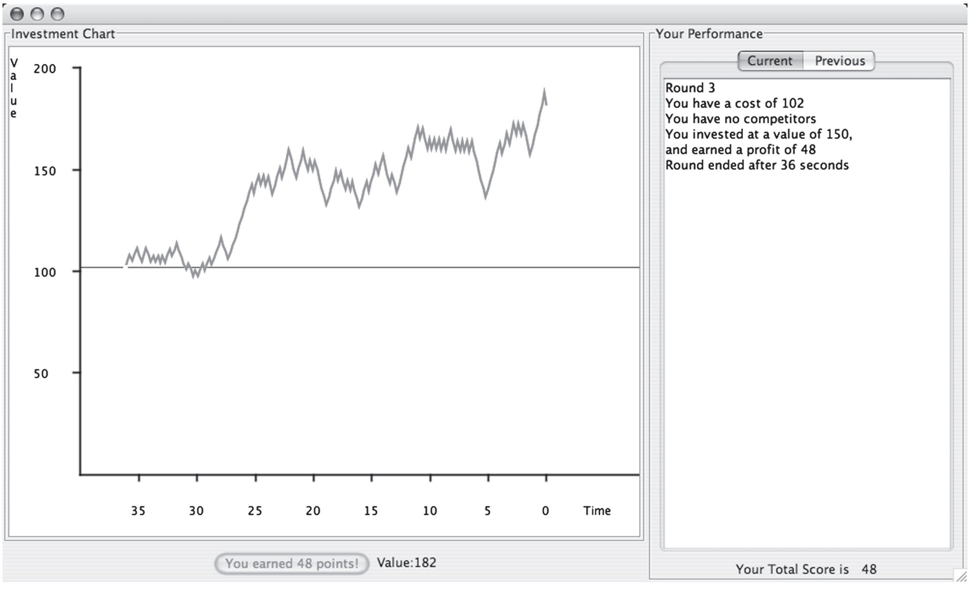

FIGURE 2

Computerized display used in the experiment using Medium B parameters. The 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. The screenshot is taken at the end of a period and shows feedback on the subject's decision and earnings. Because the software has the capacity to run experiments in which subjects compete for the investment opportunity, the information panel on subjects' screens notifies them of how many competitors they have—in this case, none.

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 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.

In addition to the graphical display, the subject's screen shows the numerical values of V and C, earnings in the current period, and cumulative earnings so far ("total score"). It also allows each subject to view her previous decisions and earnings at any time.

Subjects were 69 students (primarily) at the University of California, Santa Cruz, who were randomly assigned to the experiment using an online recruitment software. The subject pool included students with many different majors. Participation was purely voluntary. No subject participated in more than one treatment.

At the beginning of each of 10 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 interface was displayed on a wall screen. The binomial parameters for the chosen treatment were explained and written on a white board. Subjects participated in six practice periods. Each subject then participated in 80 periods for pay with no change in treatment. Sessions lasted between 80 and 120 minutes, depending on the treatment and the random draws made over the course of the session.

Nonoverlapping groups of 17 subjects participated in each of the Low, Medium A, and Medium B treatments and 18 in the High treatment. 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. In the Low treatment, subjects earned between \$5·50 and \$12·50 with median earnings of \$9·75. In the Medium A treatment, subjects earned between \$14·50 and \$25·00 with median earnings of \$18·75. In the Medium B treatment, subjects earned between \$15·50 and \$24·75 with median earnings of \$20·00. In the High treatment, subjects earned between \$17·75 and \$39·00 with median earnings of \$26·75.

4. RESULTS

Figure 3 presents several slices of the raw data that preview our main findings. Panel (a) shows empirical cumulative distribution functions (CDFs) of the option premiums ($w = \frac{V}{C} - 1$) observed in each treatment. Because the premium is observed only when the subject invested prior to expiration, the CDFs in panel (a) represent only 44% of periods. Dashed vertical lines are placed at each treatment's optimal premium. Two regularities stand out. First, the CDFs affirm hypothesis 1: Low premiums are stochastically smaller than Medium premiums, which are smaller than High premiums. Second, in most treatments, subjects undervalue the wait option.

The sample of observed investment decisions, however, suffers from a known source of statistical bias. An investment decision is more likely to be observed when the subject chooses a low value, and therefore, the observed decisions constitute a downward biased sample. One simple way of correcting this bias is to restrict attention to periods in which the maximum realization $\hat{v}$ of the investment opportunity V is particularly large. The bias is attenuated in such periods and, because $\hat{v}$ is unpredictable and exogenous, such sampling creates no new biases.

In panel (b), we provide CDFs of premiums from periods in which, at some time before expiration, the actual investment value exceeds the highest optimal threshold. Thus, for treatment τ, the "feasible draws" sample consists of periods in which $\hat{v} > (1 + w_\tau^*)\bar{C}$, where $\bar{C} = 110$ is the highest possible cost. In this sample, censoring rates drop from 56% to 4%, greatly reducing potential bias. Qualitatively, the results look similar to those derived from the unrestricted sample, once again supporting hypothesis 1. There is clear evidence of censoring bias in the High treatment (CDFs are further rightward in the restricted sample than in the unrestricted sample) and weaker evidence in the other treatments. Investment values once again tend to fall short of the optima, especially in the Medium and High treatments.

Panels (c)–(f) divide the "feasible draws" sample into time blocks. "Early" observations are from the first half of each session and "late" are from the final quarter.[5] Here, vertical lines are placed at the optimum and horizontal lines at the median. In each treatment, investment values move towards optimality over time. In Medium treatments, observed investment values are substantially closer to the optimum at the end than at the beginning, correcting at least half the initial undervaluation. In the Low and High treatments, median values converge to the optimum by the final 20 periods.

To summarize, visual inspection of raw data suggests that (i) observed option values and optimal option values have the same ordering across treatments, (ii) subjects on average undervalue the option in most treatments, but (iii) over time subjects get closer to the optimum. In the next few sections, we subject these impressions to greater econometric scrutiny and find that they hold up well.

5. We first considered defining "early" as the first quarter (first 20 periods) but found that in one treatment, there were too few "feasible draws".

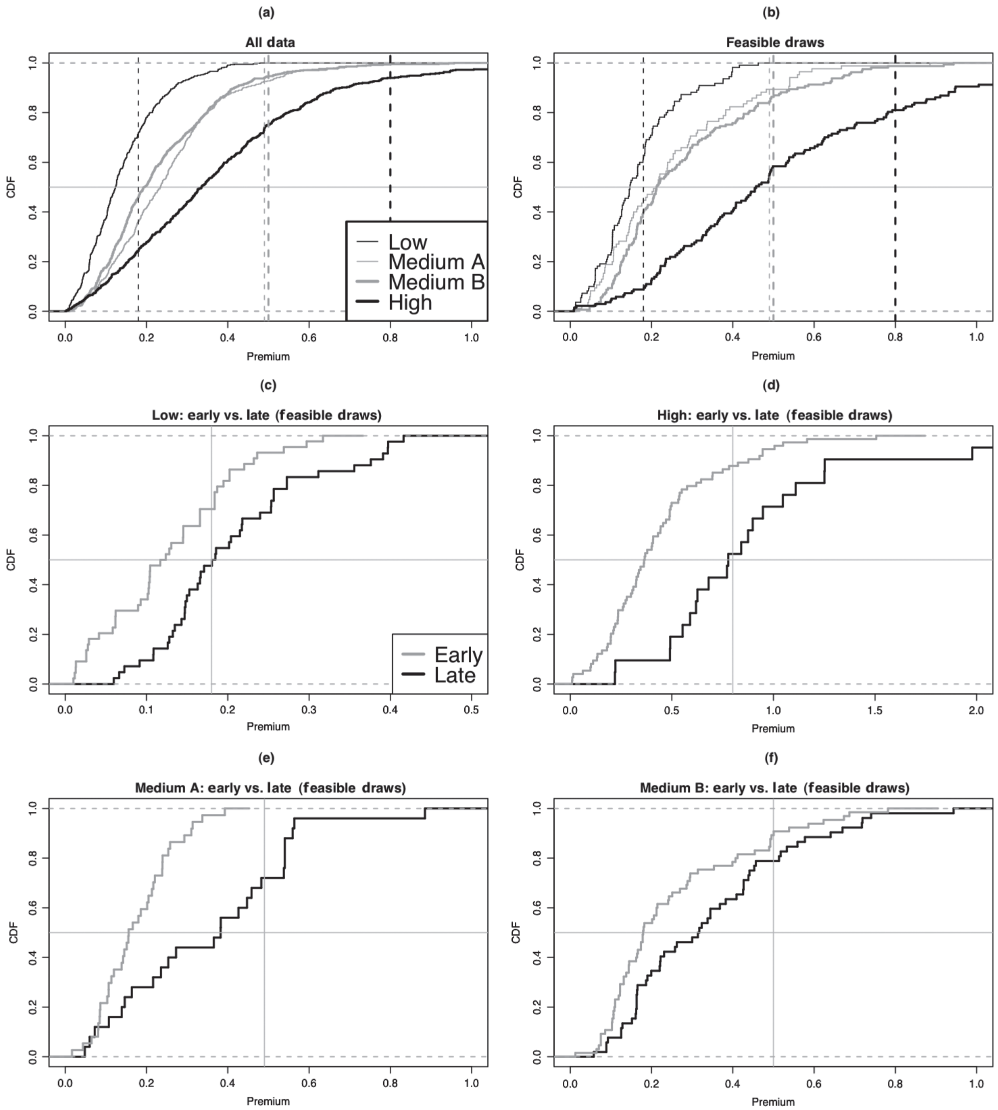

FIGURE 3

Empirical CDFs of observed option premiums. Panel (a) shows the premiums for the sample of observed investment. Panel (b) shows the premiums for the subset of observations in which the value line admits optimal investment for any cost. Vertical lines are optima. Panels (c)–(f) divide this sample into early and late periods for each treatment.

4.1. *Product limit estimation*

While the bias reduction method used in the previous section is transparent, it is inefficient since it throws out a great deal of potentially informative data. An established method for dealing with sampling bias that uses all data is the product-limit (PL) estimator, a maximum likelihood nonparametric estimator for CDFs designed to account for random right censoring (Kaplan and Meier, 1958). The PL estimator 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 option premiums when the opportunity expires before some subjects invest.

Appendix A describes the estimator more formally and also describes the log-rank test, the standard technique for comparing PL estimates across treatments.

We apply the PL estimator in two ways. First, we construct PL estimates (and standard errors) using data pooled across subjects. These estimates assume i.i.d. observations; to the extent that different subjects behave differently, the assumption is violated and the standard errors are likely to be understated. To control for within-subject dependence, we also construct a PL mean estimate for each individual subject and conduct tests on these "by-subject" samples. Such observations are completely independent within each treatment, and therefore, hypothesis tests conducted on them, which depend only on between-subject variation, are rather conservative.

4.2. *Main hypothesis tests*

Figure 4 shows the pooled PL estimated cumulative distributions of option premiums. It reveals a clear separation between the CDFs for Low, Medium, and High treatments. The CDFs are remarkably similar to those for the high value ("feasible draws") sample in Figure 3 (b). Recalling that a rightward shift in the CDF indicates a larger fraction of larger observations, one sees that the figure reflects the predicted ordinal rankings. A joint log-rank test of the equality of CDFs across treatments allows us to decisively reject the hypothesis that CDFs are equal across all treatments ($p = 0{\cdot}000$). Pairwise log-rank tests allow us to reject at the 1% level the null hypotheses that Low, Medium, and High CDFs are equal, in favour of alternative hypothesis 1. The only null hypothesis that survives the pairwise log-rank tests is that the two Medium CDFs are equal ($p = 0{\cdot}739$).

Figure 5 shows histograms of the population of by-subject means for each treatment and indicates that option premiums in the Low treatment tend to be smaller than those in the Medium

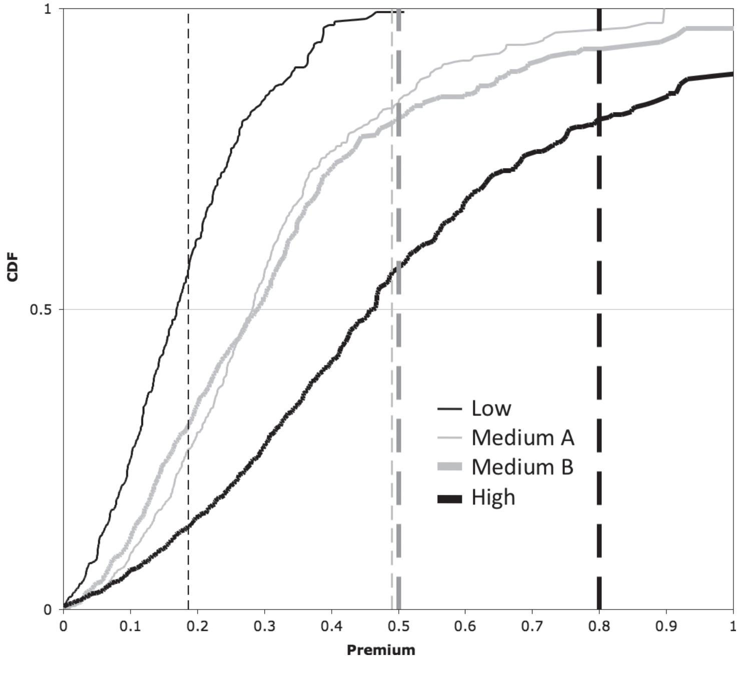

FIGURE 4

PL estimated distributions of option premium by treatment. Vertical dashed lines show predicted factors in each treatment.

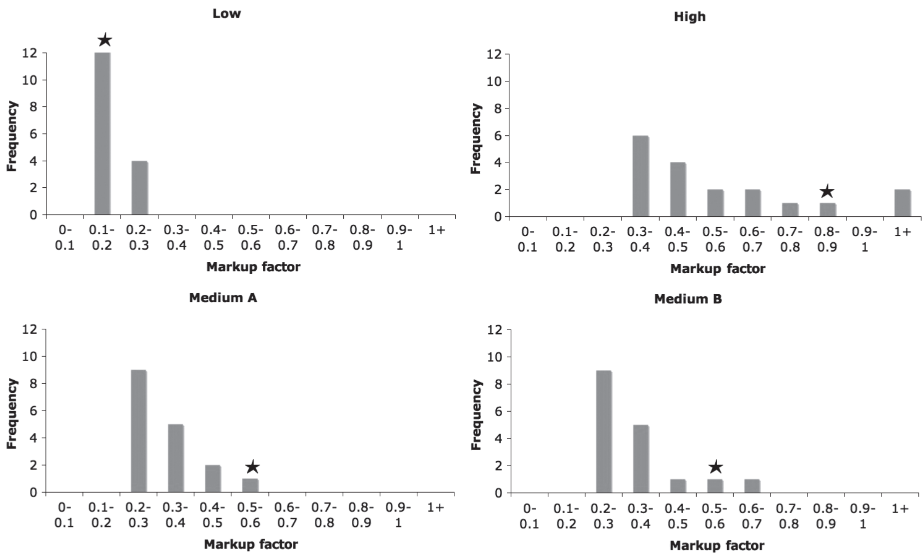

FIGURE 5

Histograms of means of product-limit estimates of option premiums estimated separately by individual subject. Asterisks indicate the histogram bins containing predicted option premiums.

TABLE 2

Mean observed option premiums

Estimator	Low	Medium A	Medium B	High
		Model predictions		
	0·179	0·490	0·499	0·804
		Means		
Pooled PL	$0{\cdot}182 \pm 0{\cdot}004$	$0{\cdot}315 \pm 0{\cdot}008$	$0{\cdot}338 \pm 0{\cdot}011$	$0{\cdot}605 \pm 0{\cdot}037$
By-subject PL	$0{\cdot}182 \pm 0{\cdot}011$	$0{\cdot}310 \pm 0{\cdot}020$	$0{\cdot}328 \pm 0{\cdot}030$	$0{\cdot}575 \pm 0{\cdot}065$
Raw	$0{\cdot}131 \pm 0{\cdot}004$	$0{\cdot}241 \pm 0{\cdot}006$	$0{\cdot}213 \pm 0{\cdot}006$	$0{\cdot}362 \pm 0{\cdot}012$
Observations	1360	1360	1360	1440
Uncensored	590	593	622	548

Notes: Model predictions are the theoretical option premiums from Table 1. The pooled PL row reports the means of the product limit estimates in Figure 4. The next row reports the mean across individual subject PL estimates in Figure 5. The raw row reports results from the subsample of uncensored observations. The values following the ± symbol are standard errors. Standard errors in the second row are for the mean of individual subject means.

treatments, which are in turn smaller than option premiums in the High treatment. Applying Mann–Whitney tests to compare these samples across each pair of treatments, we reject at the 1% level the null hypothesis in favour of the predicted alternative. The same tests fail to reject only the null hypothesis that the two Medium treatments produce the same premiums ($p = 0{\cdot}9588$).

Table 2 summarizes the central tendencies gleaned from Figures 4 and 5. The sampling bias due to censored observations can be seen by comparing the raw rows to the PL rows. Once again, the table confirms the hypothesis that $w_{Low} < w_{MediumA} \approx w_{MediumB} < w_{High}$.

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 behaviour 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 PL estimate 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 and the High treatment. The by-subject estimates in Figure 5 and 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 PL estimates. We cannot reject the hypothesis that $w_{Low} = 0{\cdot}179$ using Wilcoxon signed-rank tests ($p = 0{\cdot}9811$). Using the same test on mean samples, we can reject the hypotheses that $w_{High} = 0{\cdot}804$ ($p = 0{\cdot}0057$), $w_{MedA} = 0{\cdot}490$ ($p = 0{\cdot}0004$), $w_{MedB} = 0{\cdot}499$ ($p = 0{\cdot}0008$).

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

Is the premature investment due mainly to underreaction to one of the binomial parameters? Hypothesis 3 gives the optimal reaction and Table 3 shows that, as predicted, the uptick parameter $\bar{p}$ has small impact and the prediction lies within the estimated confidence interval. The other 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 premiums in response to either an increase in the tick size or a decrease in the expiration probability, but not as much as predicted.

4.3. *Learning*

Could the underresponse to the uptick and expiration parameters be due simply to slow learning? Table 4 suggests an affirmative answer. It shows that PL estimates usually increase from one

TABLE 3

Binomial parameter impact estimates

Parameter	Predicted	Mean
$\bar{p}$	0·003	−0·067† (0·040)
$\bar{h}$	0·317	0·223 (0·040)
$\bar{q}$	−0·308	−0·200 (0·040)

*N*otes: Predicted values are from the bottom row of Table 1. Mean estimates are obtained from equations (6–8) applied to by-subject PL estimates reported in the middle portion of Table 2. Standard errors are given in parentheses.
†Estimates whose 95% confidence intervals include the prediction.

TABLE 4

PL estimate means of option premiums by time block and treatment

	Block of periods				
	1–20	21–40	41–60	61–80	Overall
	Means of pooled data				
Low	$0{\cdot}156 \pm 0{\cdot}008$	$0{\cdot}172 \pm 0{\cdot}009$	$0{\cdot}199 \pm 0{\cdot}011$	$0{\cdot}198 \pm 0{\cdot}007$	$0{\cdot}182 \pm 0{\cdot}004$
Medium A	$0{\cdot}263 \pm 0{\cdot}011$	$0{\cdot}323 \pm 0{\cdot}017$	$0{\cdot}288 \pm 0{\cdot}014$	$0{\cdot}370 \pm 0{\cdot}017$	$0{\cdot}315 \pm 0{\cdot}008$
Medium B	$0{\cdot}309 \pm 0{\cdot}022$	$0{\cdot}288 \pm 0.012$	$0{\cdot}337 \pm 0{\cdot}024$	$0{\cdot}395 \pm 0{\cdot}022$	$0{\cdot}338 \pm 0{\cdot}011$
High	$0{\cdot}480 \pm 0{\cdot}030$	$0{\cdot}546 \pm 0{\cdot}029$	$0{\cdot}555 \pm 0{\cdot}053$	$0{\cdot}797 \pm 0{\cdot}110$	$0{\cdot}605 \pm 0{\cdot}037$

Notes: The values following the ± symbol are standard errors.

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 0·11 of prediction in the Medium treatments.

To control for intrasubject correlations, we compare the means of individual subject PL estimates across blocks. Paired Wilcoxon tests reject at the 5% 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 favour of the learning hypothesis that they are larger in the last block. The same test is significant at the 10% level in the Medium B treatment. However, we cannot reject the hypothesis that early and late option premium estimates are identical for the High treatment ($p = 0{\cdot}569$).

Restricting the signed rank tests reported in Finding 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 premiums are as predicted and both Medium premiums are significantly below predictions. However, in the High treatment, restricting to the last time block, we can no longer reject the point prediction $w^* = 0{\cdot}80$.

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 S.D. 0·21. In the last block, the average rises to 0·74 and the S.D. 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, behaviour 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.

4.4. *Evidence of adaptive learning*

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 rather flat for option premiums above 0·6 or 0·7. On the other hand, learning seems much more rapid in the Low treatment, with estimated option premiums 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 these questions using a 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 after they see that they invested at a value below the *ex post* optimum and they decrease the intended premium after they miss a profitable investment opportunity.

To measure these effects, let w_t be the intended option premium in period t, let π_t be the earnings in that period, and let π_t^o be the *ex post* maximum earnings available in that period. Let D_t^E be the ("early") indicator variable that is 1 if a subject invested at a value lower than the *ex post* maximum value round t and otherwise is 0. Likewise, let the ("late") indicator D_t^L be 1 if a subject did not invest in round t and otherwise be 0. The model describes adjustment in the intended option premiums as follows:

$$w_t - w_{t-1} = \left[\delta_E D_{t-1}^E + \delta_L D_{t-1}^L\right]\left[\pi_{t-1}^o - \pi_{t-1}\right]. \tag{9}$$

The parameter δ_E is the subject's sensitivity to a dollar lost in the last period from investing too early relative to the *ex post* optimum; it corresponds to "winner's regret" in Filiz-Ozbay and Ozbay (2007). Likewise, δ_L is the sensitivity to a dollar lost by waiting so long that a profitable opportunity expires or "loser's regret".

Since D^E and D^L are mutually exclusive, the model is especially simple and easy to estimate. Indeed, we can obtain approximate parameter measurements without imposing a parametric structure. For each pair of consecutive periods of investment t and $t-1$ with some *ex post* loss in $t-1$, we can impute the marginal impact by:

$$\hat{\delta}_{Et} = \frac{w_t - w_{t-1}}{\pi_{t-1}^o - \pi_{t-1}}. \tag{10}$$

We report the median value (since the mean value gives too much weight to outliers, arising for instance when *ex post* errors are small) in Table 5 (in the second column). The table also reports a similar nonparametric estimate of δ_L though, as explained in Appendix A.2, the procedure for δ_L is a bit more complicated because we do not actually observe target premiums during late investment. As hypothesized, intended option premiums tend to decrease after *ex post* "late" investment and to increase after *ex post* "early" investment. Both estimates are significantly different from 0 at the 1% level, according to a Wilcoxon test applied to by-subject median responses. As a robustness check, Table 5 also presents structural parametric estimates of the model (in the third column) using clustering at the subject level; see Appendix A.3 for procedural details. These estimates also indicate significant negative adaptation to too high investment and significant positive adaptation to too low investment.

TABLE 5

Estimated effects of losses on subsequent option premium choices

Coefficient	Nonparametric	Parametric
$\delta_L \times 1000$	$-0{\cdot}9185^\dagger$	$-0{\cdot}7502^\dagger$ (0·223)
$\delta_E \times 1000$	$0{\cdot}5486^\dagger$	$0{\cdot}428^\dagger$ (0·0642)

Notes: The coefficients (δ_E, δ_L) of (9), respectively, represent reactions to *ex post* losses from early investment and from not investing. Standard errors of the parametric estimates are shown in parentheses.

†Coefficients significantly different from 0 at the 1% level.

Consistent with Filiz-Ozbay and Ozbay's (2007) finding that "loser's regret" is stronger than "winner's regret" in first price auctions, we find that δ_L is larger in absolute value than δ_E in both nonparametric and parametric estimates. Wald tests from the parametric specification indicate that the difference (in absolute values) is marginally significant ($p = 0{\cdot}100$).

Finding 5. Subjects adjust their option premiums adaptively in response to foregone earnings, both from too "early" and too "late" exercise. They respond especially strongly to the latter, *i.e.*, to potential earnings missed by not exercising prior to expiration.

4.5. *Simulations of adaptive learning*

To what degree can the adaptive learning parameters reported above explain patterns observed in the data? We conducted 20 Monte Carlo simulations using the the nonparametric adjustment parameters (simulations using the parametric estimates are virtually identical) and the initial observed option premiums for each subject, a total of 1380 simulations of 80 periods each. Table 6 reports the results.

The simulated data match the pooled laboratory data reported in Table 4 in several respects. First, option premiums, which are initially low, increase over time towards optimality. Moreover, as in the experiment, the adjustment rate varies with the treatment. Option premiums in the Low treatment change only by a very small amount more than 80 periods, while premiums in the High treatment increase a great deal. Indeed, the difference between option premiums 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 the simulated data. Likewise, changes in behaviour 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 closely match those observed in the experiment.

The simulations also suggest that subjects approach optimal behaviour after 60–80 periods. Is the steady state of the learning process nearly optimal as well? To find out, for each treatment, we ran a second set of 100 simulations, each lasting 300 periods. From an initial option premium of 0·1, the simulated agents again adjusted the premium according to estimates in Table 5. The resulting data are binned at the end of each set of 20 periods and the means over simulations are plotted in Figure 6. Horizontal lines denote optimal premiums. In Low and Medium treatments,

TABLE 6

PL estimates of simulation option premiums by time block and treatment

	Block of periods				
	1–20	21–40	41–60	61–80	Overall
	Means				
Low	0·156 ± 0·001	0·165 ± 0·001	0·164 ± 0·001	0·165 ± 0·001	0·162 ± 0·000
Medium A	0·247 ± 0·002	0·391 ± 0·003	0·453 ± 0·003	0·477 ± 0·003	0·391 ± 0·002
Medium B	0·245 ± 0·002	0·382 ± 0·003	0·445 ± 0·003	0·470 ± 0·003	0·386 ± 0·002
High	0·421 ± 0·004	0·648 ± 0·008	0·760 ± 0·007	0·817 ± 0·008	0·662 ± 0·004

Notes: The values following the ± symbol are standard errors.

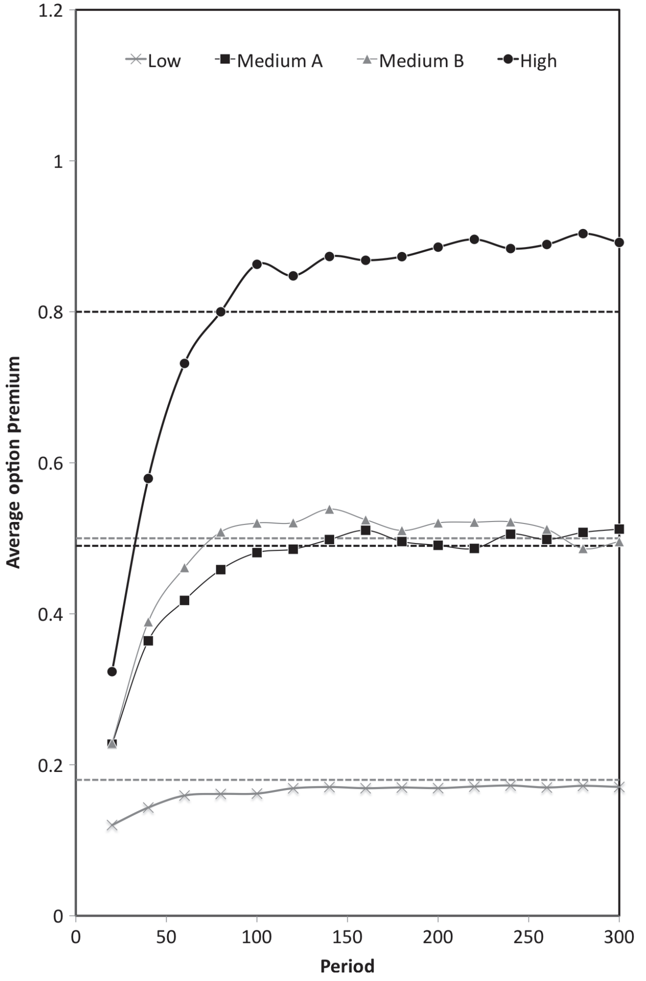

FIGURE 6

Results from simulations of adaptive learning. Each plotted point is the average over the preceding 20 periods of 100 simulations. Horizontal dashed lines show, for each treatment, the optimum.

the process converges to a tight neighbourhood of the optimum. In the High treatment, convergence is to a level slightly above the optimum.

Finding 6. Estimated adjustment parameters imply convergence to nearly optimal investment in each treatment.

It may seem surprising that unbalanced adjustment parameters (roughly −0·9 for "late" and 0·5 for "early") can lead towards optimality. Apparently, the imbalance compensates for the asymmetry of the loss functions shown in Figure 1: "early" losses tend to be larger than "late" losses. Indeed, overall our subjects had median (mean) *ex post* losses of 8·5 (12) when $D_t^L = 1$ and 17·4 (41·8) when $D_t^E = 1$. Similar simulations with balanced adjustment parameters show wide divergence from optima. Therefore, subjects' stronger response to the "late" losses helped them move towards optimality.

5. DISCUSSION

Ordinary people can indeed learn to closely approximate optimal exercise of wait options, despite options' 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 wait option prematurely, but over time their average behaviour 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 suggest 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 option and invest prematurely. 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. *Ex post* feedback is often misleading, *e.g.*, an investor who is already too impatient may see a profitable opportunity expire before he seizes it, and an overly patient investor might see a trial where it would have paid to wait even longer. 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 parameter controlling the drift of the investment opportunity. The optimum also correctly predicts the direction (but somewhat overpredicts the average magnitude) of responses to parameters controlling the discount rate (or expiration hazard) and volatility.

Our study provides several clues on how people learn in this demanding task. We report evidence that subjects respond to *ex post* errors. They tend to increase the option premium following an opportunity where a higher premium would have been more profitable and tend to decrease the option premium following a missed opportunity. They tend to respond more strongly to the second sort of *ex post* error. Remarkably, the imbalance in response seems to compensate for an asymmetric loss function and leads to convergence towards the optimum.

Are there other explanations for our findings? Risk aversion is a perennial explanation for departures from expected wealth maximization, and in our setting, it does suggest premature exercise of the wait option. Likewise, if subjects view missed investment (but not premature investment) as a loss then loss aversion could lead to premature investment. Such preference-based explanations alone, however, cannot account for the fact that, although subjects invest prematurely in early periods, they approach optimal risk neutral behaviour over time. Cognitive explanations seem a better fit to the data. A plausible explanation, suggested by a reviewer, is that subjects initially choose strategies meant to optimize single-period earnings (and so exhibit risk averse behaviour) but, after observing accumulated multiperiod earnings, switch to strategies that optimize their earnings over a broader portfolio of periods (and so exhibit more nearly risk neutral behaviour). Another possibility is that subjects initially overweight the small probability of expiration but with experience develop more realistic assessments of the expiration hazard. A third possibility—one that we favour—is that subjects initially choose conservative strategies but adjust them in response to *ex post* losses.

Several avenues of future research are now open. Longer runs, with 100 or even 200 periods (perhaps split over 2 days to avoid fatiguing subjects), would permit a more detailed analysis of heterogeneity across individual subjects and would permit more direct tests of convergence. Another avenue departs from the feature of our experiment that subjects get a lot of feedback on missed opportunities. While such feedback is present in many real-world settings (*e.g.*, re-opening a mine), there might be other settings in which participants never observe the potential value of an investment opportunity after it is seized. Learning may be different in such settings, and it is easy to implement them in the laboratory.

As noted earlier, we use the most prevalent method—an expiration probability—to induce impatience. A few experiments use an alternative method such as a visually shrinking pie (*e.g.*, Shin and Ariely, 2004). That method creates some technical problems such as a new form of heteroskedasticity (the error amplitude increases over time because the incentives to optimize shrink but the optimization costs do not) and does not eliminate the censoring problem (since eventually, the pie shrinks to invisibility). Nevertheless, it might affect the learning process and therefore seems worth investigating.

Finally, it would also be quite interesting to investigate in the laboratory other sorts of real options. We conjecture that ordinary people, even when given a decent chance to learn from personal experience, will not approximate optimality so closely for abandonment and switch options, nor perhaps for expansion options.

APPENDIX A. ECONOMETRIC DETAILS

A.1. *PL estimation*

The PL estimator produces an estimate of the distribution function, $F(w_i) = prob(x \leq w_i)$ that takes account of information contained in censored observations to correct censoring bias. Intuitively, a discrete distribution function can be expressed as the product of interval conditional probabilities. By calculating each of these conditional probabilities using only observations that have not been censored at smaller values, unbiased conditional probabilities can be formed. The product of these unbiased interval probabilities for premiums less than or equal to x is the PL estimate of the distribution function at premium x.

Consider a sample consisting of n observations. In uncensored observations denote w as the observed option premium decision $\left(w = \frac{V}{C} - 1\right)$ and in censored cases denote w as the highest premium available before censoring $\left(w = \frac{\hat{v}}{C} - 1\right)$. We can then construct a vector of these option premiums $w = (w_0, w_1, \ldots w_n)$ ordered so that $i < j$ if and only if $w_i < w_j$. Include in this vector $w_0 = 0$, the lowest premium possible at investment.

The PL 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 $1 - prob(x \geq w_i) = 1 - prob(x \geq w_i | x \geq w_{i-1}) \times prob(x \geq w_{i-1})$. Recursively then,

$$F(w_i) = 1 - \prod_{j=1}^{i} prob(x \geq w_j | x \geq w_{j-1}). \tag{11}$$

Let c_i denote the number of observations that are censored in $(w_{i-1}, w_i]$ and let d_i denote the number of investments at premium w_i. Finally, define $n_i = n_{i-1} - c_{i-1} - d_{i-1}$, the number of subjects available to invest (who have not yet invested nor been censored) at w_i. The PL estimate of $p(x \geq w_i | x \geq w_{i-1})$ is the proportion of still-in-sample investors at w_i who do not invest at w_i:

$$\hat{p}(x \geq w_i | x \geq w_{i-1}) = \frac{n_i - d_i}{n_i}. \tag{12}$$

The PL estimator is then, following (11), the cumulative product of these individual conditional probabilities at each w_i

$$\hat{F}(x) = 1 - \prod_{w_i < x} \frac{n_i - d_i}{n_i}. \tag{13}$$

Without censoring (*i.e.*, when all w_i denote option premiums 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 PL estimator is the maximum likelihood nonparametric estimator of the distribution function in environments with censoring problems analogous to ours.

Typically, hypothesis tests comparing PL distribution functions are conducted using a log-rank test. Consider two samples, labelled $j = 1, 2$. 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 will use statistics described above computed only for pool j in which cases we include the subscript. That implies, for instance, that $n_i = n_{i1} + n_{i2}$. Let d_i now represent the total number of investments made at premium ω_i so that $d_i = d_{i1} + d_{i2}$.

Under the hypothesis that the two samples are the same at ω_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 nonparametric 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 ω_i as

$$s_i^2 = \frac{n_{i1} n_{i2} (n_i - d_i) d_i}{n_i^2 (n_i - 1)}. \tag{14}$$

The test statistic for the log-rank test is then

$$z = \frac{\sum_i^n \left(d_{i1} - \frac{n_{i1} d_i}{n_i} \right)}{\sqrt{\sum_i^n s_i^2}}, \tag{15}$$

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

A.2. *Calculating learning parameters*

We can easily measure δ_E by applying (10) to all periods in which subjects invest two periods in succession. Each pair contains all the information necessary to compute (10). Although we would like to estimate δ_L in a similar way, this is complicated by the fact that we do not observe all relevant data. In particular, we never observe ω_{t-1} because late investment involves choosing a target premium that is never reached before expiration. This problem seems to only compound when there are several successive periods of noninvestment between periods of investment.

To derive an estimate of δ_L, we look at blocks in which there is at least one period of noninvestment. These blocks of N noninvestments lie between observed investments ω_{t-N-1} and ω_t. The change observed between these two periods, $\omega_t - \omega_{t-N-1}$ can be attributed to (i) any losses suffered due to early investment in period $t - N - 1$ and (ii) losses suffered due to late investment in the N intervening periods.

We therefore use our estimate of δ_E and the observed loss from early investment in $t - N - 1$, $\pi^o_{t-N-1} - \pi_{t-N-1}$, to estimate the portion of $\omega_t - \omega_{t-N-1}$ due to the initial early investment. The change in premiums $\omega_t - \omega_{t-N-1} - \delta_E \left[\pi^o_{t-N-1} - \pi_{t-N-1} \right]$ remaining after accounting for adjustment due to the early initial investment is attributed to the intervening N periods of late noninvestment. We can therefore calculate δ_L using the median estimate of δ_E reported in Table 5 as

$$\delta_L = \frac{\omega_t - \omega_{t-N-1} - \delta_E \left[\pi^o_{t-N-1} - \pi_{t-N-1} \right]}{\sum_{i=t-N}^{t-1} \left[\pi^o_i - \pi_i \right]}. \tag{16}$$

One potential worry about these nonparametric estimates is that δ_E is only measured in periods in which subjects followed early investment by another nonlate investment. This raises the possibility that the sample on which δ_E is estimated is not random and it is conceivable that estimates may be downward biased. Were this bias to exist, correcting it would only strengthen our result, increasing the magnitude of both δ_E and δ_L and to a similar degree. Our parametric analysis, which uses the complete sample to measure δ_E actually produces (insignificantly) *smaller* estimates, suggesting that our method does not in fact introduce such bias.

A.3. *Parametric learning estimation*

At first, it seems natural to estimate the parameters (δ_E, δ_L) of (9) directly, by simply appending an additive error term, say

$$\omega_t - \omega_{t-1} = \left[\delta_E D_t^E + \delta_L D_t^L \right] \left[\pi^o_{t-1} - \pi_{t-1} \right] + \epsilon_{t-1} \tag{17}$$

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

$$\tilde{\omega}_t - \tilde{\omega}_{t-1} = \left[\delta_E D_t^E + \delta_L D_t^L \right] \left[\pi^o_{t-1} - \pi_{t-1} \right] + \epsilon_{t-1}, \tag{18}$$

where $\tilde{w}$ is an observed option premium $w = V/C - 1$ when $D^L = 0$ and is latent when $D^L = 1$. Consider a round t in which $D^L = 0$, and let N_t be the number of consecutive periods preceding t in which $D^L = 1$. Then, we can rewrite (18) entirely in terms of observables:

$$\tilde{w}_t - \tilde{w}_{t-N_t-1} = \delta_L \sum_{i=t-N_t}^{t-1} \left[\pi_i^o - \pi_i\right] + \delta_E \left[\pi_{t-N_t-1}^o - \pi_{t-N_t-1}\right] + \sum_{k=t-N_t-1}^{t-1} \epsilon_k. \tag{19}$$

We can estimate (19) conveniently using a standard weighted least squares package on the subset of observations for which $D_t^E = 1$. 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

$$w_t - w_{t-N_t-1} = \delta_E \sum_{i=t-N_t}^{t-1} \left[\pi_i^o - \pi_i\right] + \delta_L \left[\pi_{t-N_t-1}^o - \pi_{t-N_t-1}\right] + \zeta_t, \tag{20}$$

where ζ_t is distributed $N(0, \sigma^2 N_t)$. The parameter N_t shows up in the error term but only affects the amplitude of variance and therefore cannot generate bias due to endogeneity.

APPENDIX B. 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 the Figure 2 seen earlier in this article.

B.1. *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.

B.2. *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.

B.3. *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.

B.4. *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)*0.02 = $5.00 + $13.64 = $18.64.

B.5. *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., $\frac{1}{2}$ 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.

B.6. *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).

Acknowledgements. 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 Juan Carrillo, Tom Copeland, Tek Ho, John Morgan, and Lenos Trigeorgis and from participants at a University of California, Santa Cruz, brown bag seminar, the Fall 2006 ESA meetings in Tucson, the 2007 11th International Real Options Conference in Berkeley, and the 2007 Invitational Choice Symposium at the University of Pennsylvania and also from three anonymous referees of this *Review*.

REFERENCES

BERGER, P., OFEK, E. and SWARY, I. (1996), "Valuation of the Abandonment Option", *Journal of Financial Economics*, **42**, 257–287.

BOX, G., HUNTER, J. and HUNTER, W. (1978), *Statistics for Experimenters: An Introduction to Design, Data Analysis, and Model Building* (New York: John Wiley and Sons).

BRAUNSTEIN, Y. and SCHOTTER, A. (1981), "Economic Search: An Experimental Study", *Economic Inquiry*, **19**, 1–25.

CARILLO, J. and MARIOTTI, T. (2000), "Strategic Ignorance as a Self-Disciplining Device", *Review of Economic Studies*, **67**, 529–544.

CASON, T. and FRIEDMAN, D. (1999), "Learning in a Laboratory Market with Random Supply and Demand", *Experimental Economics*, **2**, 277–298.

COPELAND, T. and ANTIKAROV, V. (2003), *Real Options: A Practitioner's Guide* (New York: Texere).

COX, J. C. and OAXACA, R. L. (2000), "Good News and Bad News: Search from Unknown Wage Offer Distributions", *Experimental Economics*, **2**, 197–225.

DIXIT, A. K. and PINDYCK, R. S. (1994), *Investment Under Uncertainty* (Princeton: Princeton University Press).

FILIZ-OZBAY, E. and OZBAY, E. (2007), "Auctions with Anticipated Regret: Theory and Experiment", *American Economic Review*, **97**, 1407–1418.

FRIEDMAN, D. and RUST, J. (1993), "The Double Auction Market: Institutions, Theories and Evidence", in *Santa Fe Institute Proceedings*, Vol. XIV (Boston: Addison-Wesley) 3–26.

GONG, B. and RAMACHANDRAN, V. (2007), "Payment Dominance in Sequential Search" (manuscript, University of Maryland).

HENRY, C. (1974), "Investment Decisions under Uncertainty: The Irreversibility Effect", *American Economic Review*, **64**, 1006–1012.

HIRSHLEIFER, J., GLAZER A. and HIRSHLEIFER D. (2005), *Price Theory and Applications*, 7th edn (Cambridge: Cambridge University Press).

KAPLAN, E. and MEIER, P. (1958), "Nonparametric Estimation from Incomplete Observations", *Journal of the American Statistical Association*, **53**, 457–481.

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.

KREPS, D. (1990), *A Course in Microeconomic Theory* (Princeton: Princeton University Press).

MACDONALD, R. and SIEGEL, D. (1986), "The Value of Waiting to Invest", *Quarterly Journal of Economics*, **101**, 707–728.

MOEL, A. and TUFANO, P. (2002), "When are Real Options Exercised? An Empirical Study of Mine Closings", *Review of Financial Studies*, **15**, 35–64.

OPREA, R., FRIEDMAN, D. and ANDERSON, S. (2007), "Learning to Wait: A Laboratory Investigation" (manuscript, University of California, Santa Cruz).

PADDOCK, J., SIEGEL, D. and SMITH, J. (1988), "Option Valuation of Claims on Real Assets: The Case of Offshore Petroleum Leases", *Quarterly Journal of Economics*, **103**, 479–508.

QUIGG, L. (1993), "Empirical Testing of Real Option-Pricing Models", *The Journal of Finance*, **48**, 621–640.

ROBSON, A. (2002), "Evolution and Human Nature", *Journal of Economic Perspectives*, **16**, 89–106.

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: Erlaaum Ass) 79–102.

SHIN, J. and ARIELY, D. (2004), "Keeping Doors Open: The Effect of Unavailability on Incentives to Keep Options Viable", *Management Science*, **50**, 575–586.