Exp Econ (2007) 10:79–104
DOI 10.1007/s10683-006-9134-0

Searching for the sunk cost fallacy

Daniel Friedman · Kai Pommerenke · Rajan Lukose · Garrett Milam · Bernardo A. Huberman

Received: 17 March 2005 / Revised: 15 June 2006 /
Accepted: 15 June 2006 / Published online: 7 February 2007

Abstract We seek to isolate in the laboratory factors that encourage and discourage the sunk cost fallacy. Subjects play a computer game in which they decide whether to keep digging for treasure on an island or to sink a cost (which will turn out to be either high or low) to move to another island. The research hypothesis is that subjects will stay longer on islands that were more costly to find. Eleven treatment variables are considered, e.g. alternative visual displays, whether the treasure value of an island is shown on arrival or discovered by trial and error, and alternative parameters for sunk costs. The data reveal a surprisingly small sunk cost effect that is generally insensitive to the proposed psychological drivers.

Keywords Sunk costs · Sunk cost fallacy · Search · Self-justification · Loss aversion

JEL Classification C91, D11

1 Introduction

A cost is *sunk* when it cannot be recovered. Once a cost is sunk, it has no effect on the incremental payoffs of future decisions, and therefore plays no role in rational choice. Indeed, sunk costs should play no role in any outcome-oriented process, rational or otherwise.

Economists generally assume that people are rational, and pride themselves on internal consistency. It is therefore remarkable that they devote so much class time and so many textbook pages to teaching undergraduate and MBA students to ignore sunk costs. A favorite anecdote is a concertgoer who realizes after the first five minutes

Electronic Supplementary Material Supplementary material is available in the online version of this article at http://dx.doi.org/10.1007/s10683-006-9134-0.

D. Friedman (✉) · K. Pommerenke · R. Lukose · G. Milam · B. A. Huberman
Department of Economics, University of California, Santa Cruz 95064, USA
e-mail: dan@ucsc.edu

that the show is horrible, but sticks around just to “get his money’s worth” from his purchase of an admission ticket. The textbooks explain that expenditures on the ticket are sunk and thus irrelevant. The student is warned never in her career to commit the *sunk cost fallacy* of taking some otherwise undesirable action simply because of a sunk cost, e.g., investing another million dollars on an unprofitable product line just because ten million has already been invested.

So who is right, the economist model builders who assume rationality, or the economist teachers and textbook writers who think that hard work is needed to stamp out the fallacy? How widespread really is the sunk cost fallacy? Our own interest in these questions began with a practical issue in e-commerce: is it true, as claimed by several observers, that people stay at a website longer when it takes longer to download?

Our investigation encountered many surprises. The first is that the published evidence for the fallacy is not as definitive as we had supposed. The next section shows that one can rationalize the choices featured in most studies and anecdotes, for example.

We therefore devised a direct laboratory test inspired by the e-commerce issue. We put subjects in front of a computer screen, present them sequentially with “islands” that contain various amounts of “buried treasure”, and grant them a limited number of mouse clicks for uncovering the treasure. To get to a new island the subject must sink a cost that will turn out to be either high or low. The sunk cost fallacy leads to the following testable hypothesis: subjects click more often on high cost islands than on otherwise identical low cost islands.

The next surprise was the difficulty in demonstrating the fallacy. Our initial treatments produced essentially the same distribution of clicks on low and high cost islands. We created a design capable of detecting very small effects and, following advice of colleagues, tried many new treatments. The most recent data confirm a sunk cost effect, but it is much smaller and less robust than we had originally expected. Variables representing rational choice are much more powerful in explaining the data than any of the psychological variables we have investigated.

The next section discusses existing empirical evidence on sunk costs, emphasizing the logical gaps. Section 3 describes our experiment. A final surprise was that the formal analysis of the apparently simple treasure hunt game is rather complex; Section 4 sketches the results. Section 5 collects testable hypotheses and notes some results from previous laboratory search experiments. Section 6 reports our own results. The concluding section summarizes, offers possible interpretations and suggests directions for new research. Appendix A presents the formal derivation of optimal search for our task.

2 Existing evidence

Recent news stories suggest that the sunk cost fallacy exists on a grand scale, but at the same time they underscore ambiguities. Was the US government’s final decision to invade Iraq in March 2003 a sunk cost fallacy? Iraq’s dictator seemed ready to agree to very intrusive weapons inspections and to the placement of tens of thousands of NATO troops, meeting the stated goals of the US. On the other hand, in readying the attack, the US had already spent tens of billions of dollars and disrupted the lives of more than one hundred thousand soldiers. Some commentators argued that these sunk costs precluded calling off the invasion. Several other interpretations are equally plausible, however. For instance, US policymakers may have believed that a last minute cancellation of the invasion would hurt their credibility, or may have had goals beyond those announced at the time.

The loss of space shuttle Columbia in February 2003 brought to mind numerous previous decisions to continue the NASA's shuttle program. From its inception in the 1970s, the shuttle was criticized as extremely cost ineffective and dangerous. Yet each time its supporters pointed to the lives lost and dollars already spent as a reason for continuing the program, and so far Congress has always agreed (Economist, 2003). Again there are other interpretations. For NASA managers and Congress, for example, the interests of contractors and other clienteles may sometimes seem more urgent than making the space program safer and more cost effective. And admitting a huge mistake might not be good for managers 'future careers or for their mentors' place in history.

Psychologists have studied the fallacy for several decades (e.g., Staw, 1976; Bazerman, 1986, chapter 4; see also Thaler, 1980), often referring to it as "irrational escalation of commitment." The underlying mechanism mentioned in older papers is cognitive dissonance (Festinger, 1957) or self-justification (Aronson, 1968). More recent discussions (e.g., Whyte, 1993) often tie it instead to prospect theory (Kahneman and Tversky, 1979), specifically to a fixed reference point and loss aversion.

Most of the evidence consists of responses to hypothetical survey questions. For example, subjects are asked to imagine that they have spent $50 on a ticket for event A and $100 on a ticket for event B (e.g., A and B are ski weekends in Wisconsin and Michigan). They are also to imagine that they really prefer A to B, that they have just discovered that the events are mutually exclusive, and that the tickets have no salvage value. When asked which event they would then choose, about half the subjects select the more expensive but less preferred alternative B. The psychologist authors urge the interpretation that their subjects respond more strongly to the $50 difference in sunk costs than to the "true" preferences. Since the choices are not salient, a skeptic could offer alternative interpretations, e.g., that subjects attend more to their actual homegrown preferences between A and B, or to the impression they make on the person asking the question.

Psychologists report that the sunk cost effect increases in the size of the hypothetical sunk cost, especially in proportional terms (e.g., Garland and Newport, 1991). The effect is very sensitive to framing, and is reduced by emphasizing the salience of the incremental costs (e.g., Northcraft and Neale, 1986; Tan and Yates, 1995). Indeed, Heath (1995) finds a strong *reverse* sunk cost effect when the sum of incremental and sunk cost would exceed the total benefit.

Posing hypothetical questions on whether to grant an additional bank loan for continuing a project, Garland and Conlan (1998) find that sunk costs are less important than whether the additional loan will allow project completion. Their interpretation is that the goal of project completion psychologically displaces the profit goal. Again, a skeptic could offer an alternative explanation, e.g., that the subjects respond to imbedded real options. Refusing the additional loan to complete a project would extinguish the wait option for the project, and might hurt the bank's reputation.

Eyster (2002) expresses a consensus view that "the most convincing single experiment comes from Arkes and Blumer (1985)," experiment 2. In this field experiment, 20 randomly selected buyers were given a small ($2) discount, 20 others a large discount

($7, almost half the price), and 20 others no discount, on season tickets to the campus theater. After excluding 10% of the subjects who bought tickets as couples, the authors report that the no-discount group used more tickets than either discount group in the first half of the theater season ($p < .05$); no significant differences are detected between the small and large discount groups or in the second half of the season. Thus this field experiment provides some evidence consistent with a sunk cost fallacy, but to the best of our knowledge it still awaits replication with a larger sample.

The animal behavior literature reports a controversy dating back to Trivers (1972) on the "Concorde effect," an allusion to continuing government subsidies of the uneconomic supersonic passenger plane. Arkes and Ayton (1999) conclude that "there are no unambiguous instances" of the sunk cost fallacy among animals, or even human children. They argue that adult humans commit the fallacy by misapplying the "don't waste" rule.

There are only a few relevant laboratory experiments using salient payments. Phillips et al. (1991) found that increases in the sunk price of a lottery ticket led to increased valuations by one quarter of the subjects, but to decreased valuations by another quarter and no response by half. With opportunities to learn in a market setting, very few subjects responded to the sunk price. Meyer (1993) reported that about half of his subjects bid more relative to a benchmark when an auction entry fee became larger. However, questions remain about his theoretical benchmark, the symmetric equilibrium bidding function for a different kind of auction than used in the experiment. We are also aware of an unpublished study by Offerman and Potters (2001) that shows sunk costs can facilitate coordination, and an inconclusive study by Elliott and Curme (1998).

Some non-experimental field evidence suggests the sunk cost fallacy. Camerer and Weber (1999) confirm Staw and Hoang's (1995) observation that first year professional basketball players who are drafted earlier (and thus, by the nature of the draft system, represent larger sunk costs) get more playing time, conditional on measured performance. Again a skeptic could offer alternative explanations. Possibly there are unobserved components of performance, or perhaps the Bayesian priors that led the team to draft a player may also lead to more playing time during the first season. Barron et al. (2001) find that US firms are significantly more likely to terminate projects following the departure of top managers. This might reflect the new managers' insensitivity to costs sunk by their predecessors, or it might simply reflect two aspects of the same broad realignment decision.

Do Internet users respond to sunk time costs? Manley and Seltzer (1997) report that after a particular website imposed an access charge, the remaining users stayed longer. A rival explanation to the sunk cost fallacy is selection bias: the users with shortest stays when the site was free are those who stopped coming when they had to pay. Klein et al. (1999) report that users stick around longer on their site after encountering delays while playing a game, but again selection bias is a possible alternative explanation. The issue is important in e-commerce because "stickier" sites earn more advertising revenue. Schwartz (1999) reports that managers of a Wall Street Journal site deliberately slowed the login process in the belief that users would then stay longer. One of us (Lukose) took a sample of 2000 user logs from a website and found a significant positive correlation between residence time at the site and download latency. One alternative explanation is unobserved congestion on the web, and users may have been responding more to expected future time costs than to time costs already sunk. Also, good sites may be more popular because they are good, leading to (a) congestion and (b) more time spent on the site.

The issues can be summarized as follows. There are at least two distinct psychological mechanisms that might create an irrational regard for sunk costs. Self-justification (or cognitive dissonance) induces people who have sunk resources into an unprofitable activity to irrationally revise their beliefs about the profitability of an additional investment, in order to avoid the unpleasant acknowledgment that they made a mistake and wasted the sunk resources. Loss aversion (with respect to a reference point fixed before the costs were sunk) might induce people to choose an additional investment whose incremental return has negative expected value but still has some chance of allowing a positive return on the overall investment.

There are also several possible rational explanations for an apparent concern with sunk costs. Maintaining a reputation for finishing what you start may have sufficient value to compensate for the expected loss on an additional investment.[1] The "real option" value (e.g., Dixit and Pindyck, 1994) of continuing a project also may offset an expected loss. Agency problems in organizations may make it personally better for a manager to continue an unprofitable project than to cancel it and take the heat from its supporters (e.g., Milgrom and Roberts, 1992; see also Kanodia et al., 1989).

The available evidence is extensive and varied,[2] but remains somewhat ambiguous. Besides confounding the various rational and irrational explanations, the studies often are unable to control for unobserved Bayesian priors, selection biases, etc. Clearly there is room for a new laboratory experiment that eliminates the rational explanations and the unobservable factors, and that allows alternative psychological explanations to demonstrate their explanatory power.

3 Experiment design

Subjects play a computerized game called "Treasure Hunt" in which they visit a sequence of "islands." In the baseline treatment, each island has 20 sites the subject can "dig up" by clicking the mouse; the subject earns 5 points each time she clicks a site with buried treasure. The "voyage" ends when she exhausts a fixed click budget, e.g. 200 clicks. Budget permitting, she can click as many of the 20 sites as she wants before "sailing North or South" to the next island.

Leaving for a new island involves a sunk cost. The subject is told that because of unpredictable weather at sea, "... your cost (in points) of reaching the next island is either high or low. The amount of buried treasure on an island is not affected by the cost of getting there." Figure 1 shows the user interface.

[1] In a variant on this theme, Carmichael and MacLeod (2003) argue that attending to sunk costs can help solve the hold-up problem.

[2] An anonymous referee cogently argues that the principle of parsimony favors a single explanation—the sunk cost fallacy—over the long list of alternative explanations for the varied evidence. If the published evidence included all potential evidence we would have been convinced by this argument.

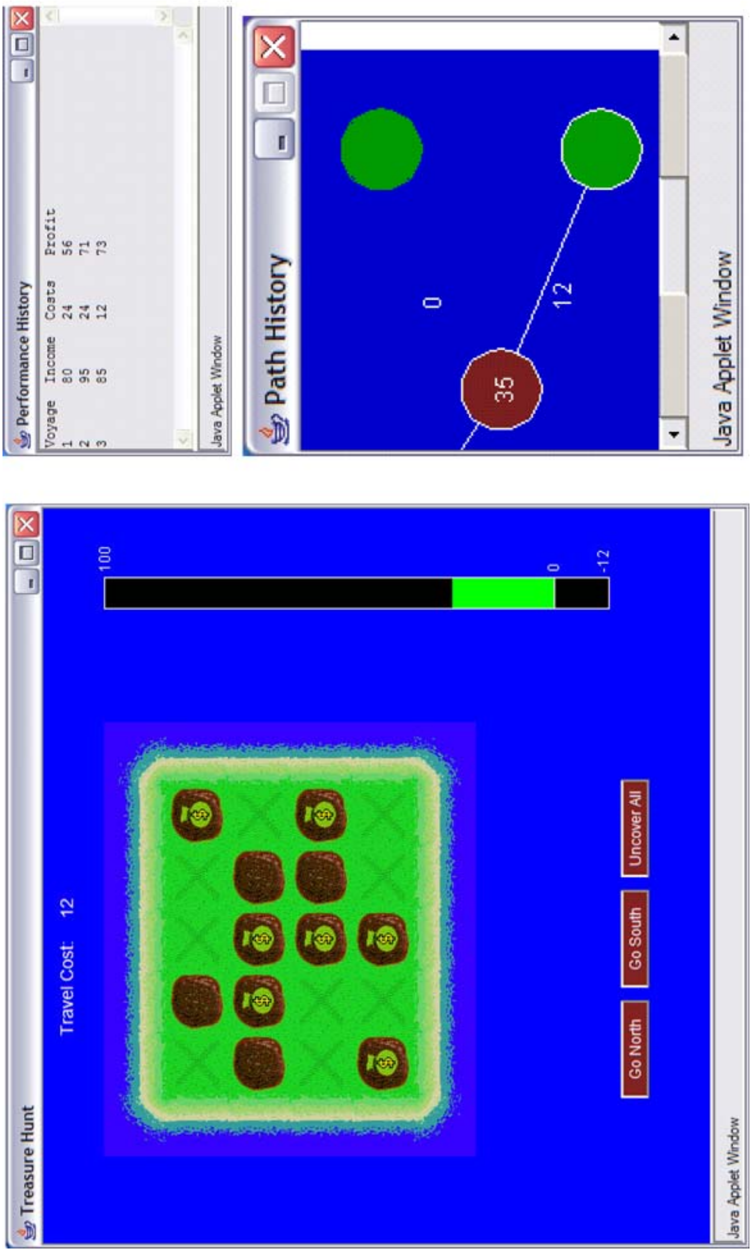

Fig. 1 User interface

The number of treasures buried on each island is an independent uniform random draw from consecutive integers {L, …,U} with $0 \leq L < U \leq 20$, and the travel cost is either $c_L \geq 0$ or $c_H > c_L$. Subjects are told the values of these parameters, e.g., that each island will contain between L = 2 and U = 18 treasures. The pay rate, e.g., two cents per point, is posted on the board. Subjects are paid for one voyage, selected at the end of the session by a public random device.

As the experiment progressed, we introduced new treatments suggested by colleagues and by members of our team. The following lists the reported treatments,[3] with the baseline values in parentheses.

1. Displaying upon arrival the initial number of treasures on the island (Show Island Value, SIV = y). In the alternative treatment, SIV = n, subjects must rely on their own estimates.
2. Requiring subjects to click all sites or none (Require Complete Uncovering, RCU = y), used only in conjunction with SIV = y. In the alternative, subjects can click sites one at a time.
3. The click budget (Nclicks = 200); variants included Nclicks = 100.
4. In the baseline (Choose Next Island, CNI = y), subjects choose whether to sail North or South, but these choices have no effect on the distribution of sunk costs (c_L and c_H are equally likely) and island values. In alternative treatment CNI = a, the distributions differ in ways known to the subjects, e.g. c_H and more buried treasure are more likely when sailing North. At the other extreme, in treatment CNI = n, the subject chooses only when to leave, not whether to go North or South.
5. In the baseline (Show Other Island Cost, SOIC = y), after choosing North or South, the subject sees the cost of her own choice and also the cost of the direction not chosen. In the alternative treatment SOIC = n, she sees only her own cost. Treatments 4 and 5 are intended to manipulate the saliency of self-justification.
6. When traveling to the next island, subjects experience a time delay proportional to the cost in points (Cost Pause, CP = 0.8 seconds per point); alternatives include CP = 0, 0.4 and 1.6.
7. In the baseline (Thermometer Displayed, TD = y), the screen displays a thermometer-like graph of net cumulative points earned on the island. It starts out in the red (as the negative of realized travel cost c_L or c_H) and turns green when it reaches positive territory as treasures are found. In the alternative treatment TD = no, this part of the graphical display is suppressed. Treatments 6 and 7 are intended to manipulate the saliency of loss aversion.
8. Alternative choices to the baseline parameters ($c_L = 2$ points, $c_H = 18$ points, L = 2 sites and U = 18 sites) are often used; some colleagues conjectured that the sunk cost effect will be stronger when $c_L = 0$.
9. The baseline treatment draws the probability of hitting a treasure with no replacement (Replace = n) from the discrete uniform distribution with endpoints $0 \leq L < U \leq 20$. The alternative treatment Replace = y draws a probability p for each island independently from a continuous uniform distribution on [l, u]

[3] Pilot experiments not reported here explored two additional treatments: more formal instructions, and sunk costs incurred as clicks rather than as points. Neither treatment had a discernable effect on responsiveness to sunk costs.

$\subset$ [0, 1], and each click on that island has independent probability p of hitting treasure. Sometimes p is displayed on arrival to the island: Show Hit Probability, SHP = y. The alternative treatments simplify the rational choice computation.

10. Experience = y indicates that the subject participated in an earlier session of Treasure Hunt.
11. Intense incentives were used in the last five sessions. Here, in addition to other salient payments, subjects were given the option to return the $5 show-up fee and take one last 200-click "bonus voyage" that paid at a much higher rate. All 42 of the subjects accepted this option and earned a bonus between $10 and $15.

Another design feature deserves mention. Except in treatment CNI = a, subjects are paired so that for each subject who reaches a given island at high cost, there is another subject who reaches the same island at low cost. This pairing reduces experimental error, and it is feasible (except in CNI = a) because the random sequences of travel costs and of hits and misses on each island are drawn in advance.

A total of 155 UCSC undergraduates played Treasure Hunt as inexperienced subjects and 52 of them returned in a later session as experienced subjects. In a quarter of the sessions, the subjects played Treasure Hunt exclusively and most earned $20–22 (including a $5 show up fee) during 80–90 minutes of play. In other sessions subjects earned about the same amount overall but played an interactive game as well as Treasure Hunt; the latter accounted for a third to a half of salient payments and time.

Table A1 in the Appendix summarizes the design of the 36 sessions analyzed below. A factorial design with so many treatments would require thousands of sessions, so the actual design is sparse. It focuses on treatment combinations that seem most likely to produce the strongest sunk cost effect, according to published articles, our own intuition, and suggestions of experienced colleagues.

4 Search theory

Our subjects report that the game was easy and fun. To most of them it seemed like a rather simple video game that was much easier to play than it was to read the original instructions.[4] However, the formal analysis of the game is quite complicated; indeed it took us several tries over a period of months to get it right. There are several different cases, depending on whether the island value is displayed (SIV), whether subjects may click sites one at a time (RCU), and whether replacement is used in probabilities of hitting treasure (Replace). Appendix A collects the formal results, which can be summarized as follows.

Table 1 below splits the data by case and reports a significant sunk cost effect only in Case 3, again a reverse effect.[5] Panel B of the table separates out the non-obvious islands, defined as those whose displayed value is within 15 of the reservation value R.

[4] Hence the reported sessions used only the one-page informal instructions available in the online version of this article. See also the previous footnote.

[5] This result is inconsistent with Heath (1995), who finds that mentioning benefits more explicitly in questionnaires tends to reverse the sunk cost effect. The reverse effect is strongest in our data when the benefits are least explicit, in Case 3 where SIV = n. In other respects, our work is consistent with Heath, of whom we became aware only after completing our experiment and initial write up. Heath also finds

Table 1 Weighted average click difference

Panel A: By case

Case	W. avg. click difference	t-value	nobs
1: SIV = y, RCU = y	−0.25	−0.75	966
2: SIV = y, RCU = n	0.07	0.47	4449
3: SIV = n, RCU = n	−0.45	−2.56	3609

Panel B: By case and island value

Case	Island value	W. avg. click difference	t-value	nobs
1: SIV = y, RCU = y	Low	−0.01	−0.02	437
	Medium	−0.63	−0.78	317
	High	−0.21	−1.05	212
2: SIV = y, RCU = n	Low	0.16	0.65	2014
	Medium	0	−0.01	1684
	High	−0.03	−0.19	751

Note: Weighted average (W. avg.) click difference is the difference between the mean number of clicks by the high cost group and the low cost group on each island, with islands weighted by the number of subjects in the smaller group. The t-values compare the weighted average click difference to zero. Low (resp. High) Island values in Panel B are those less than R − 15 (resp. greater than R + 15), where R is the reservation value defined in Section 4. Panel B excludes Case 3 since the High/Medium/Low classification is not defined when the island value is not displayed.

Table 2 Choices by case and cost

Case	Cost	Choices (% of total) Optimal	Stubborn	Impatient	Chi-square (nobs)	p-value
Panel A: By case (all islands)						
1: SIV = y, RCU = y	Low	78.5	21.4	0.2	2.35	0.31
	High	77.4	21.9	0.8	(1,193)	
2: SIV=y, RCU = n	Low	64.8	34.6	0.6	0.72	0.70
	High	65.1	34.3	0.6	(81,292)	
3: SIV = n, RCU = n	Low	80.0	19.0	1.0	5.01	0.08
	High	79.4	19.6	1.0	(82,843)	
Panel B: By case for medium island values						
1: SIV = y, RCU = y	Low	63.4	36.1	0.5	2.06	0.36
	High	59.1	39.1	1.9	(409)	
2: SIV = y, RCU = n	Low	68.7	30.3	0.9	6.3	0.04
	High	69.9	29.2	0.9	(39,790)	

Contrary to our conjecture, the sunk cost effect is not stronger for such islands; indeed, Panel B shows a stronger (but still not significant) reverse sunk cost effect for Medium island values in Case 1. Thus direct tests of H1 fail to establish a sunk cost effect.

ambiguity in earlier studies, and his results suggest that the cost-in-points treatment we use will produce a stronger sunk cost effect than the cost-in-clicks treatment mentioned in footnote 3 above.

Case 1. SIV = y, RCU = y, i.e. the island value v (or hit probability p) is shown on arrival and the click choice is all-or-none, as in approximately 8% of the voyages listed in Table A1. This case requires only a minor extension of classic search theory, and the optimal search is characterized by a reservation value R. Budget permitting, the optimal strategy is to click out the island (expend 20 clicks) if $v \geq R$, and immediately to sink the cost and move on if $v \leq R$. For the uniform distributions on $[L, U]$ used in the experiment, the risk neutral reservation value is closely approximated by $R = bU - \sqrt{2cb(U - L)}$. For example, with expected sunk cost $c = (c_L + c_H)/2 = (2 + 18)/2 = 10$, treasure value $b = 5$, and the number of treasures per island between $L = 9$ and $U = 18$, we get $R = 60$. Then a rational player would click out the islands showing values $v = 60$ and above and would pass on the islands showing values 55 and below. For Replace = y, set $v = 20\ bp = 100p$; e.g., in the example, click out islands showing $p = 0.60 = 60\%$ and above.

Case 2. SIV = y, RCU = n, approximately 42% of the voyages. This case is identical to case 1, except that the player has the option to leave the island after clicking some but not all sites. With the number of clicks divisible by 20 and with replacement (Replace = y), the option has no value: if it is optimal to click once on a given island, then it is optimal to click 20 times, so the rule is the same as in Case 1. For Replace = n, however, the option is valuable. For a player who gets sufficiently lucky on the first several clicks, the remaining sites are not worth clicking because his luck must "catch up" to the displayed island value. Hence optimal behavior is more complicated and is computed using techniques similar to those discussed for case 3b below.

Case 3. SIV = n, RCU = n, approximately 50% of the voyages. Even when it does not hit treasure, each click has an information value because the player can update his estimate of the number of remaining treasures. This would seem to give more scope for self-justification and perhaps strengthen the sunk cost effect. There are two subcases.

(a) When Replace = y, we obtain essentially a click-by-click reservation value $h^*(n)$, the minimum number of hits in the first n clicks required to justify staying on an island. This value also depends on the size of the remaining click budget.
(b) When Replace = n, the optimal policy is even more complicated because the catch-up effect (mentioned in Case 2 above) opposes the effect of information updating.[6] Usually the optimal choice (for a given number of clicks on an island and given number of clicks remaining) is to leave if the number of hits falls into a middle range and otherwise to click once more, but occasionally the optimal strategy is still more complicated.

Case 4. SIV = n, RCU = y. The player must decide all-or-none whether to click out an island about which he can obtain no information. This case is uninteresting and is not used in the experiment.

[6] The authors were confused on this point for quite a while. An example may provide intuition. Suppose that you know initially that between 2 and 18 of the 20 sites contain treasure, and the first 4 clicks do not hit anything. The 16 remaining sites must therefore contain between 2 and 16 treasures, increasing the probability of hitting treasure on the next click. This catch-up effect sometimes dominates the Bayesian updating effect that negatively skews the posterior distribution. The treatment Replace = y eliminates the catch-up effect and simplifies the analysis.

5 Testable hypotheses

The Matlab program available in the online version of this article identifies the optimal decision for each click on each voyage, and the expected loss for each suboptimal decision. It allows us to classify each decision as either

- Impatient: the subject left the island when it was optimal to stay, thus incurring a loss $x > 0$ in expected earnings; or
- Stubborn: the subject clicked another site on the island when it was optimal to leave, thus incurring a loss $y > 0$ in expected earnings; or
- Optimal: no alternative action would generate higher expected earnings, and $x = y = 0$.

The classification allows sharper characterizations of subjects' behavior. For example, suppose that on average subjects stay longer on high cost islands. The sunk cost fallacy is confirmed if this arises from stubbornness on high cost islands and approximate optimality on low cost islands. However, if we find optimality on high cost islands but impatience on low cost islands, then the data reflect other departures from rationality.

The following testable hypotheses guide the data analysis.

H1.0. The null hypothesis is that average number of clicks on each island is the same for players who reached it with a low sunk cost as for players who reached it with a high sunk cost.

H1.R. The research hypothesis, a one-sided alternative, is that the average number of clicks is higher for high sunk cost players.

In testing H1 when the island value is displayed on arrival (Cases 1 and 2), we anticipate a stronger sunk cost effect for "close-calls," islands whose value is in the vicinity of the optimal reservation value R. The sunk cost fallacy should have less impact when islands are obviously worth clicking out or obviously better to skip.

The three-way classification of choices permits a more refined test:

H2.0. Impatient choices and stubborn choices have the same distribution on islands reached with high sunk cost as on islands reached with low sunk cost.

H2.R. Impatient choices are less frequent and stubborn choices are more frequent on islands reached with high sunk cost than on islands reached with low sunk cost.

It is reasonable to say that costly mistakes are more meaningful than mistakes that incur negligible losses. Hence we also run payoff domain tests:

H3.0. The loss of expected earnings due to impatient choices ($\text{IL} = \sum x$) and the loss of expected earnings due to stubborn choices ($\text{SL} = \sum y$) have the same distribution on islands reached with high sunk cost as on islands reached with low sunk cost.

H3.R. Average IL is smaller and average SL is larger on islands reached with high sunk cost than on islands reached with low sunk cost.

The experiment design encourages a more detailed dissection of treatment effects and individual choice. Each click by each subject gives us a value of the indicator variable for staying: $Z = 0$ if the player sinks a cost to go to the next island, and $Z = 1$ if the player clicks another site on the current island. A logit regression can explain this dependent variable by variables representing rational and psychological motives, other treatment variables, and their interactions. The null hypothesis is:

H4.0. Estimated coefficients in the logit regressions will be large and significantly positive for variables representing rational motives for staying, and will be insignificant for the dummy (indicator variable) for high sunk cost and for its interactions with other variables.

H4.R. The dummy variable for high sunk cost and its interactions with several of the treatment variables will be significant. In particular, self-justification theory suggests positive interactions with treatments 4 and 5, especially in Case 3, and loss aversion suggests positive interactions with treatments 6 and 7.

Before turning to our own results, we note that several authors have investigated search behavior in the laboratory, e.g., Schotter and Braunstein (1981), Hey (1982, 1987), Cox and Oaxaca (1989, 1992), Kogut (1990), Harrison and Morgan (1990), and more recently Sonnemans (1998, 2000) and Cura-Juri and Galiani (2003). With the exception of the last paper (which investigates dynamics when inflation increases price dispersion), the typical setting is analogous to our Case 1 with a minimal click budget: the subject samples wage offers from a known distribution and the trial ends as soon as one offer is accepted. Most subjects appear to use suboptimal heuristics rather than the optimal reservation price strategy (e.g., they often recall earlier offers when permitted), but they do not usually incur very large payoff losses. Overall there seems to be a slight bias towards what we call impatience.

6 Results

First consider H1, the most direct test of the sunk cost fallacy. For each island we calculate the average number of clicks under high cost c_H minus average clicks under low cost c_L, so positive differences represent a sunk cost effect. Differences are averaged across islands using weights proportional to the number of players reaching the island. The overall click difference is $-.17$, an insignificant reverse sunk cost effect (the paired t-test value is -1.63).

We turn now to a finer grained examination of hypotheses H2 to H4. Table 2 shows that, despite the complexity of the calculation, from 65 percent (in Case 2) to 80 percent (in Case 3) of all choices are optimal. Stubborn choices account for most of the departures. This choice asymmetry is not surprising: on a given island one can be stubborn many times but impatient at most once.

Table 2 also shows that in Cases 1 and 3 stubborn choices are more frequent and optimal choices are less frequent on high cost islands than on low cost islands, consistent with H2.R. However, the shift is small and is reversed in Case 2, and no consistent picture emerges for impatient choices. A chi-square test indicates the shifts are insignificant except perhaps in Case 3 (significant at a marginal 8% level). Panel B restricts the analysis to medium island values (R $-15 \leq$ island value $\leq$ R $+ 15$), and here the choice shift becomes significant for Case 2 at the 4% level, but in the wrong direction.

Now consider the payoff domain. Cases 2 and 3 allow the calculation of value gained or lost on each click. Define total potential value (TPV) as the sum of the absolute difference in expected profit between immediately leaving the island and

Table 3 Value gained or lost by cost

Case	Cost	Value as % of TPV			Nobs
		Gain	Stubborn loss	Impatient loss	
2: SIV = y, RCU = n	Low	93.6	6.0	0.4	38,090
	High	93.4	6.3	0.4	43,292
3: SIV = n, RCU = n	Low	96.3	3.0	0.7	39,889
	High	96.3	3.0	0.7	42,954

staying for another click. Otherwise put, TPV= actual value gained + SL + IL. Table 3 shows that subjects overall lose less than 7% of TPV. There really is an asymmetry in that 3–6% of TPV is lost due to stubbornness but less than 1% due to impatience. The data support H3.R for Case 2. The difference in value lost due to stubbornness between the low and high cost groups is small (0.3%) but highly significant ($t = -4.1$, $p < 0.0001$). In Case 3, however, stubborn losses and impatient losses are the same on high and low cost islands, so here we can't reject the null hypothesis.

Case 2 (SIV = y, RCU = n) without replacement (replace = n) permits sharp tests of two well-known heuristics. According to the "win-stay, lose-shift" heuristic (e.g., Eyster, 2002), subjects who just experienced success are more likely to stay on an island. Optimality predicts exactly the opposite because in Case 2 the catch-up effect is not offset by an update effect. To test these opposing predictions, and to refine the estimates of the sunk cost effect, we ran the logit regression reported in Table 4.

The second line indicates that higher sunk cost (cost = 1 if $c = c_H$ and =0 otherwise) increases the log odds of staying by 0.09, significant at the 5% level, consistent with the sunk cost effect and hypothesis H4.R. The third line shows a strong impact of the difference between the value of clicking and the value of leaving the island (Stay surplus), e.g., increasing stay surplus by 2 treasures or 10 points increases the log odds by 3.6, consistent with noisy rational search. The last line investigates the heuristic. The dummy variable Last-click-successful is set to one if the previous click hit treasure, and to zero otherwise. The line indicates that the rational catch-up effect dominates the win-stay, lose-switch heuristic.

A second heuristic is based on loss aversion. Eyster (private conversation) predicts that subjects whose cumulative earnings on a given island are still negative (i.e. have not yet covered the sunk cost) are more likely to stay. The variable Cumulative loss is the minimum of zero and the cumulative earnings on that island. Hence it is negative until a subject recoups the travel cost by finding treasures on the island. Panel A of

Table 4 Logistic regression for win-stay, lose-switch heuristic

Parameter	Estimate	Standard error	Wald Chi-Square	Pr > ChiSq
Intercept	3.48	0.05	6303	<.0001
Cost (high cost = 1)	0.09	0.04	3.9	0.05
Stay surplus	0.36	0.01	971	<.0001
Last-click-successful dummy	−0.70	0.05	227	<.0001

Notes: Dependent variable $Z = 1$ if subject clicks once more, $Z = 0$ if subject leaves the island. Case 2 data. Number of Observations = 74,952. The interaction term cost* Last-click-successful is insignificant when included in the regression.

Table 5 Logistic regression for loss aversion

Parameter	Estimate	Standard error	Wald chi-square	Pr > ChiSq
Panel A: Case 2. Number of Observations = 81,382				
Intercept	2.85	0.03	9581	<<.0001
Cost	0.34	0.05	56.7	<.0001
Stay surplus	0.35	0.01	1249	<<.0001
Cumulative loss	0.79	0.06	166	<.0001
Cost*Cum. loss	−0.72	0.06	134	<.0001
Panel B: Case 3. Number of Observations = 82,843				
Intercept	2.84	0.03	8092	<<.0001
Cost	−0.01	0.05	0.05	0.82
Stay surplus	0.32	0.01	1096	<<.0001
Cumulative loss	−0.74	0.18	16.8	<.0001
Cost*Cum. loss	0.73	0.18	16.3	<.0001

Notes: Dependent variable $Z = 1$ if subject clicks once more, $Z = 0$ if subject leaves the island

Table 5 reports a logistic regression with dependent variable Z, the decision to dig for another treasure (click once more) on the same island. The highly significant positive coefficient of 0.79 indicates that larger cumulative losses (more negative values) tend to decrease the probability of staying on an island reached with low travel costs, contrary to loss aversion. The last line in Panel A shows that the effect disappears when travel costs are high, as the sum of the coefficients on cumulative loss and the interaction variable Cost*Cumulative loss is close to zero.

On the other hand, panel B shows that the results are consistent with loss aversion in Case 3 for islands reached with low travel costs. Of course, since Cumulative loss has range $[-\mathrm{c}, 0]$, its absolute value remains small when $\mathrm{c} = \mathrm{c}_L$. Also, as in Case 2, the effect disappears on islands reached with high travel costs. Hence the data provide little evidence of loss aversion.

Table 6 reports our sharpest tests of Hypothesis 4. The Stay surplus coefficients in all panels confirm the huge impact of rational considerations. The second line of Panel A reports the best evidence we have found for the sunk cost effect. Although relatively small, the Cost coefficient estimate has the predicted sign and is significant at the 1% level. The main effects for the self-justification treatments are both significant at the 5% level (indeed, SOIC is significant at the 0.01% level). Thus allowing the subject to choose North or South, and showing the cost of the route not taken, both increase subjects' tendency to stay on the current island.

The treatment predictions in Hypothesis 4 concern interactions rather than main effects. Panel C covers Case 3 and hence should give the research hypothesis its best shot. Here the CNI interaction with Cost is significant at the 5% level but has the wrong sign. The other self-justification interaction, with SOIC, has the predicted sign but it is barely significant at the 10% level. The loss-aversion treatments both have highly significant interaction coefficients with the predicted sign. Panel B covers Case 2 and here the results are different. The CNI interaction coefficient flips to the predicted sign and the SOIC interaction coefficient keeps the predicted sign and becomes significant. The loss-aversion interactions both flip sign, although only the interaction with TD (the thermometer display) retains significance. All these variables become

Table 6 Logistic regressions for decision to stay

Parameter	Estimate	Standard error	Wald chi-square	Pr > ChiSq
Panel A: Main Effects, Cases 2 and 3 combined. Number of observations = 164,225				
Intercept	2.34	0.12	397.	<.0001
Cost	0.07	0.03	6.2	0.01
Stay surplus	0.35	0.01	2785.	<.0001
Choose next island	0.16	0.08	4.2	0.04
Show other island cost	0.25	0.06	19.1	<.0001
Cost pause	0.00	0.00	0.13	0.72
Thermometer displayed	0.03	0.04	0.42	0.52
Panel B: Interactions, Case 2. Number of observations = 81,382				
Intercept	2.74	0.03	10620.	<.0001
Cost	−0.89	0.24	14.0	0.0002
Stay surplus	0.37	0.01	1398.	<.0001
Cost*CNI	0.55	0.14	16.1	<.0001
Cost*SOIC	0.79	0.20	14.9	0.0001
Cost*CP	−0.00013	0.00008	2.33	0.13
Cost*TDn	−0.24	0.10	5.71	0.02
Panel C: Interactions, Case 3. Number of observations = 82,843				
Intercept	2.86	0.03	8369.	<.0001
Cost	−0.37	0.29	1.65	0.20
Stay surplus	0.32	0.01	1113.	<.0001
Cost*CNI	−0.40	0.20	3.95	0.047
Cost*SOIC	0.14	0.09	2.74	0.098
Cost*CP	0.0006	0.0002	10.3	0.001
Cost*TD	0.17	0.08	4.91	0.027

Note: Dependent variable $Z = 1$ if subject clicks once more, $Z = 0$ if subject leaves the island. The dummy variables CNI, SOIC and TD are 1 for default value y of the corresponding treatments and are 0 when the treatments have value n. All interaction terms between cost and treatment variables are insignificant when jointly included in the Panel A regression.

insignificant when all main effects and interactions are included in the same logistic regression.

Additional tests, omitted here, find no evidence that offering an asymmetric choice between North and South (CNI = a) increases the sunk cost effect. Indeed, when such a choice is offered, the effect is slightly stronger for those who choose South, contrary to prediction; the difference however at best is marginally significant. The weighted average click difference shows no consistent trend over voyage number; indeed we find no evidence that the sunk cost effect increases or decreases with experience. Also, increasing the contrast between high and low sunk costs (with or without a corresponding change in time delay, CP) seems perversely to reduce the sunk cost effect, but again the impact is not significant at the 5% level. Reducing the lower sunk cost to $c_L = 0$ has no detectable incremental impact. Neither does the intense incentive (or bonus voyage) treatment.

7 Discussion

The experiment seeks to isolate the famous but elusive sunk cost fallacy. Even though it is difficult to analyze theoretically, the treasure hunt task is simple to play. Our matched pair design can detect even very small sunk cost effects, and numerous treatments enable us to explore the proposed psychological drivers of the fallacy. The results can be summarized as follows.

1. Subjects' choices are surprisingly consistent with optimal search behavior. A large majority of choices are optimal, and actual losses in expected payoff represent less than 7% of possible losses. Losses due to stubbornness are larger than losses due to impatience, in part because subjects can be stubborn many times on a given island, but impatient only once.
2. There is evidence for the sunk cost fallacy. Stubborn errors are more frequent when sunk costs are high in Case 1 (all-or-none choice, island value displayed on arrival) and Case 3 (click-by-click choice, island value not displayed), and on average stubborn losses are larger in the remaining case 2 (click-by-click choice, value displayed). The relevant logit regressions indicate that subjects are more likely to stay on islands reached with higher sunk cost.
3. The effect is surprisingly small and inconsistent. The simple comparison (click difference) indicates a small reverse effect, and so do several variants of the stubborn error and logit specifications. Probably the strongest evidence for the fallacy is the main effect for Cost in Panel A of Table 6. The coefficient 0.07 is significant at the 1% level and has the right sign, but implies a rather small effect: even a player who would stay with probability 0.50 with low sunk cost would stay with probability only $1/(1 + \exp(-0.07)) \approx 0.52$ with high cost.
4. The treatments intended to manipulate the psychological drivers of the fallacy also have rather small and inconsistent impact. Contrary to conjecture, the variables manipulating self-justification work best in Case 2, while in Case 3 they either have the wrong sign or are hardly significant. On the other hand, the variables manipulating loss aversion have the wrong sign (and one is insignificant) in Case 2 but work better in Case 3.

The results reported here arose from an extensive design search. Persuasive anecdotes and conventional wisdom let us to anticipate finding a substantial sunk cost effect quickly, but we did not. Helpful advice from many colleagues suggested a succession of new treatments and eventually we found the small and inconsistent effects just described. In sum, we were unable to find sunk-cost tasks and treatments that reliably lead subjects to substantial departures from rational behavior.

Is the earlier evidence on the sunk cost fallacy due mainly to rational factors that we eliminated, such as reputation or real options? Or does our design somehow introduce a reverse sunk cost effect that almost negates the direct effect? Can new tasks and treatments, or more structural estimates from field data, demonstrate more substantial sunk cost effects? These open questions remain a challenge for future research.

Table A1 Sessions and treatments

Date	# of complete voyages	# periods w/cost diffs	Low cost (c_L)	High cost (c_H)	Island values, hit prob.	Click budget	Stakes (cents/point)	CP	SIV, SHP	RCU	TD	SOIC	CNI	# of people	Experience
8/1/02	52	11	2	18	35 − 65	100	4	TMT	No	No	Yes	Yes	yes	4	no
8/7/02	60	16	2	18	5–95	100	4	TMT	No	No	Yes	Yes	yes	4	no
8/15/02	16	0	2	18	20–80	100	4	800	No	No	Yes	Yes	TMT	2	no
8/20/02	57	11	2	18	20–80	100	4	800	TMT	TMT	Yes	Yes	yes	4	yes
8/21/02	26	7	2	18	20–80	200	2	800	No	No	Yes	Yes	TMT	4	no
8/22/02	14	4	2	18	20–80	200	2	800	Yes	No	Yes	Yes	TMT	4	no
9/5/02	15	4	1	13	10–90	200	2	800	Yes	No	Yes	Yes	TMT	4	no
9/11/02	30	5	1	13	10–90	200	2	800	Yes	Yes	Yes	Yes	TMT	6	yes
9/12/02	13	4	1	13	10–90	200	2	800	No	No	Yes	Yes	TMT	3	yes
1/24/03	46	9	1	13	20–80	200	2	800	TMT	No	Yes	Yes	Yes	6	no
1/31/03	36	9	1	13	10–90	200	2	800	TMT	No	Yes	Yes	Yes	4	no
2/4/03	29	0	1	13	10–90*	200	2	800	TMT	No	Yes	Yes	Yes	4	no
2/5/03	46	8	1	13	10–90*	200	2	800	TMT	No	Yes	Yes	Yes	6	yes
2/19/03	22	0	1	13	10–90*	200	2	800	TMT	No	Yes	Yes	Yes	4	yes
5/19/03	38	5	1	13	10–90*	200	2	800	TMT	No	Yes	Yes	Yes	8	no
5/23/03	21	4	1	13	10–90*	200	2	800	TMT	No	Yes	Yes	Yes	6	no
6/2/03	18	6	1	13	10–90*	200	2	800	TMT	No	Yes	Yes	Yes	4	yes
8/14/03	62	12	1	13	10–90*	200	4	800	No	No	Yes	No	Yes	6	no
8/21/03	76	15	0	10	10–90*	200	4	TMT	Yes	TMT	Yes	Yes	Yes	6	no
10/2/03	23	7	1	13	.1–.9	200	2	800	Yes	No	Yes	Yes	Yes	4	no

(*Continued on next page.*)

Table A1 (*Continued*).

Date	# of complete voyages	# periods w/cost diffs	Low cost (c_L)	High cost (c_H)	Island values, hit prob.	Click budget	Stakes (cents/ point)	CP	SIV, SHP	RCU	TD	SOIC	CNI	# of people	Exper- ience
10/3/03	24	5	1	13	.1–.9	200	2	800	Yes	No	Yes	Yes	Yes	6	no
10/10/03	34	5	1	13	.1–.9	200	2	800	TMT	No	Yes	Yes	Yes	6	yes
10/23/03	54	9	0	12	.1–.9	200	2	800	TMT	No	Yes	Yes	Yes	8	both
11/6/03	24	5	0	12	.1–.9	200	2	800	TMT	No	Yes	Yes	Yes	6	both
11/14/03	31	6	0	12	.1–.9	200	2	800	TMT	No	Yes	Yes	Yes	6	no
12/4/03	44	12	0	12	.1–.9	200	1	800	TMT	No	Yes	Yes	Yes	4	yes
12/4/03	3	1	0	12	.1–.9	100	5	800	No	No	Yes	Yes	Yes	4	yes
1/23/04	27	5	0	12	.1–.9	200	2	800	TMT	No	TMT	Yes	Yes	8	no
1/23/04	8	1	0	12	.1–.9	100	4	800	No	No	Yes	Yes	Yes	8	no
1/28/04	34	8	0	12	.1–.9	200	2	800	TMT	No	TMT	Yes	Yes	7	no
2/4/04	34	9	0	12	.1–.9	200	1	800	TMT	No	TMT	Yes	Yes	4	no
2/4/04	4	1	0	12	.1–.9	100	4	800	No	No	Yes	Yes	Yes	4	no
2/11/04	58	7	0	12	.1–.9	200	1	800	TMT	No	TMT	Yes	Yes	12	no
2/11/04	12	1	0	12	.1–.9	100	4	800	No	No	Yes	Yes	Yes	12	no
2/25/04	17	4	0	12	.1–.9	200	1	800	Yes	TMT	TMT	TMT	Yes	7	no
2/25/04	6	1	0	12	.1–.9	100	5	800	No	No	Yes	Yes	Yes	7	no

Notes: The third column reports the cases where completed voyages had complementary cost structures. All sessions use treatments $r = 5$ points per treasure, AutoDig = yes, while other treatments vary as indicated; TMT (Treatment) indicates variation within session. In the sixth column, entries less than 1.0 indicate probabilities in Replace = y sessions, and an * indicates CNI = a sessions. Island values |North = Island values |South and Prob(c = c_H | North) = Prob(c = c_H | South) = 0.5, except for the * sessions, for which Prob(c = c_H |North) = 0.7 and Prob(c = c_H |South) = 0.4 with the entry showing Island values | North while Island values | South = 10–70 (10–80 on 2/19/03).

Pilot experiments conducted before 8/1/02 are not listed and were excluded from the data analysis because different instructions were used and the data format is incompatible with later formats. Similarly, three sessions, run on 8/14/02, 2/12/03 and 2/14/03, were excluded because of technical problems with the software during the session.

Appendix A: Computation of rational decisions

The classic economic model of sequential job search (e.g., Lippman and McCall, 1976) gives insight into our treasure search task. In the classic model, a job seeker can always pay a cost $c \geq 0$ to receive another job offer $y \in [0, \infty)$, assumed an iid random variable with known distribution function F. The search terminates as soon as an offer $y = x$ is accepted, and the payoff is $x - nc$, where n is the number of offers purchased. In the simplest version, the job seeker is risk-neutral, there is no discounting (time lags are negligible, as in our experiment), and there is no bound on the number of offers that can be purchased.

With a current offer x in hand, the job seeker maximizes expected value V defined recursively by the Bellman equation

$$V(x) = \max\{x, -c + EV(y)\}. \tag{1}$$

It is well-known that this problem has a unique solution using a reservation price R. That is, the solution is of the form: accept the most recent offer x if $x \geq R$ and otherwise pay c for another offer drawn from the distribution F. The reservation price R is determined from (1) by equating the value of the current offer $x = R$ to the value of continuing an optimal search $-c + EV(y) = -c + \int(\max\{y, R\})dF(y) = -c + R\int_0^R dF(y) + \int_R^\infty ydF(y) = -c + R + \int_R^\infty (y - R)dF(y)$. Cancelling R from both sides of the equation and simplifying slightly we get the marginal condition

$$c = H(R), \quad \text{where } H(z) = \int_z^\infty (y - z)dF(y). \tag{2}$$

That is, R equates the incremental cost of search c to its incremental expected benefit $H(R)$. If F has a positive density over its support $[L, U] \subset [0, \infty)$, it is easily checked that the function H is strictly decreasing from $H(L) = Ey - L$ to $H(U) = 0$. Then (2) has a unique solution $R = H^{-1}(c) > 0$ for any search cost $c \in (0, Ey - L)$.

The classic problem can be adapted to a finite horizon. If only m more draws are possible, then the value function depends on m as well as x and the solution reservation price decreases as m decreases.

Case 1: Value displayed, uncover all or none

Our treasure search problem at first glance looks like the classic finite horizon problem, but it turns out to be a bit different. Consider first the case where each island value is known upon arrival and one must uncover all sites on the island or none. Without further loss of generality, normalize so that one click uncovers all the sites and the initial click budget is $Y > 0$. In the experiment the standard click budget is $Y = 200/20 = 10$ with this normalization.

Since the click budget is separate from earnings, there is no limit on the number of islands that can be visited and skipped. Hence the analogy is to Y classic job searchers, each with an infinite horizon. The solution is to click islands whose displayed value W is at least R, and otherwise to skip to the next island, until the click budget is exhausted.

(This assumes that the travel cost is not exorbitant. If $c \geq Ey = (U + L)/2$, a situation never seen in the experiment, then the player should quit playing rather than skip.)

To compute R, first recall that the (expected) travel cost is $c = (c_L + c_H)/2 > 0$, and that the number of treasures is uniformly distributed between $L \geq 0$ and $U \leq n_{\max} = 20$ with value $b = 5$ value per treasure. The continuous uniform distribution function for value then is $F(x) = (x - bL)/b(U - L)$ for $x \in [bL, bU]$, with $F(x) = 0$ for $x < bL$ and $F(x) = 1$ for $x > bU$, and the incremental benefit function is

$$H(z) = \int_z^{bU} (y - z)dF(y) = (bU - bL)^{-1} \int_z^{bU} (y - z)dy = \frac{(bU - z)^2}{2b(U - L)}. \qquad (3)$$

The function H is continuous and, for $z < bL$, it is linear decreasing (with slope $= \int_{bL}^{bU} -1dF(y) = -1$); quadratic decreasing for $z \in [bL, bU]$; and is 0 for $z > bU$. Solving $c = H(R)$ we obtain

$$R = bU - \sqrt{2bc(U - L)}. \qquad (4)$$

It follows that R decreases from $R = bU$ when $c = 0$ to $R = bL$ when $c = Ey - bL = b(U - L)/2$. For $c \in [b(U - L)/2, b(U + L)/2]$ every island should be clicked, but for $c \geq b(U + L)/2 = Ey$ the search should be abandoned. For $0 \leq c \leq b(U - L)/2$, the expected value of moving to the next island is R before sinking the cost and is $R + c$ after arrival.

The reservation value from (4) is not exact when, as in the experiment, the number of treasures on an island must be an integer. The uniform distribution then has discrete support $\{bL, b(L + 1), \ldots, bU\}$ with equal mass $1/(U - L + 1)$ at each point. For $z \in [bL, bU]$, write $z = b(L + i_z + r_z)$, where i_z is the unique integer between 0 and $U - L$ such that the residue r_z is in [0, 1). The H-function for this discrete uniform distribution is $H_D(z) = \int_z^{\infty}(y - z)dF(y) = \frac{b}{U-L+1}[(1 - r_z) + (2 - r_z) + \cdots + (U - L - i_z - r_z)]$. Sum the series to get

$$H_D(z) = \frac{b(U - L - i_z)(U - L - i_z + 1 - 2r_z)}{2(U - L + 1)}. \qquad (5)$$

Comparing (3) and (5), or just noting that the distributions have the same support, one can see that $H_D(z) = H(z) = 0$ for $z \geq bU$ and that $H_D(z) = H(z)$ for $z \leq bL$; in particular $H_D(bL) = H(bL) = b(U - L)/2$. Moreover, the slope of H_D increases in $U - L + 1$ equal steps from -1 at $z = bL$ to 0 at $z = bU$, while the slope of H increases linearly from -1 to 0 over the same interval. Thus H_D is a continuous, piecewise linear approximation of the quadratic function H, and (4) closely approximates the exact reservation value when $U - L$ is reasonably large, as in the experiment.

Case 2: Value displayed, uncovering discretionary

Now consider the decision problem when the player can choose click by click whether to skip to the next island. Recall that there are two ways of specifying island value. With *replacement*, Case 2a, each click hits treasure with independent constant probability p. When p is displayed on arrival and the number of remaining clicks is evenly divisible by the number of $n_{\max}$ of sites per island, then the decision problem is equivalent to the problem in the non-discretionary case (under the maintained hypothesis of risk neutrality). The expected island value is $W = bn_{\max} p = 100p$ and the optimal strategy is to click all $n_{\max} = 20$ sites on the present island if p is at least $r = R/100$, and otherwise to click none. The logic and computation of R are exactly as in the previous case.

The problem is considerably more complicated in Case 2b, *no replacement*. In this case the number of treasures w is shown on arrival and the island value therefore is $W = bw$. For the first click, the hit probability is $p_o = w/n_{\max}$, but it changes after that. The hit probability p given h hits out of $n < n_{\max}$ clicks so far on the current island is the number of remaining treasures divided by the number of remaining sites,

$$p = \frac{w - h}{n_{\max} - n}. \tag{6}$$

Hence the hit probability p typically rises after a miss and declines after a hit. This *catch-up effect* can cause an initially attractive value of $p = p_o$ to become quite unattractive after a hot streak. Thus the player may rationally click some but not all sites on an island. Likewise, if the click budget is not evenly divisible by $n_{\max} = 20$, then a detailed analysis is again necessary. The material below on Bellman equations covers these subcases.

Case 3: Value not displayed, uncover all or none
The player has no basis for distinguishing one island from another on arrival and can't sample. Thus there is never a reason to skip an island; that only increases cost without increasing expected revenue. As in case 1, the player should quit if cost exceeds average island value. Otherwise he should dig up every island in order until the budget is exhausted. This case is trivial and we don't use it in the experiment. (Warning: the text of the paper refers to the present case as case 4 and the next case as case 3 to preserve continguity.)

Case 4: Value not displayed, uncovering discretionary
This case is intricate, due to the *update effect*: after each click, a player should use Bayes theorem to update his estimate of the island value. As in case 2, there are two subcases. With no replacement (Case 4b) the catchup effect opposes the update effect and is stronger in the extreme cases (very few hits or very few misses), but is weaker in other cases. Consequently the optimal strategy here cannot be expressed in terms of a reservation price. Replacement (case 4a) eliminates the catchup effect. Here the optimal search is characterized by a reservation price (in terms of hits and clicks so far on the island) that reflects the information value of another click as well as the Bayes posterior expected values.

In both subcases, a key computation is $p(h, n)$, the posterior probability that the next click on the current island will hit treasure, given h hits out of n clicks so far on the current island. We are given the prior distribution $f(p)$ with support contained in [0, 1]. Of course, $p(0, 0)$ is simply the prior mean $\int_0^1 pf(p)dp \equiv \bar{p}$. In the experiment, subjects are told the maximum U and minimum L numbers of treasures and that the distribution is uniform, so $\bar{p} = (U + L)/(2n_{\max}) = (U + L)/40$.

By Bayes theorem, the posterior density $f(p \mid h, n)$ of the hit probability given (h, n) is the likelihood of (h, n) times the prior probability $f(p)$ and normalized so that the expression integrates to 1.0. The desired posterior probability $p(h, n)$ is the expectation $\int_0^1 xf(x \mid h, n)dx$.

With replacement as in case 4a, the likelihood is the binomial expression $\binom{n}{h}$ $x^h(1-x)^{n-h}$. With a continuous uniform prior supported on $[l, u] \subset [0, 1]$ we therefore have

$$p(h, n) = \frac{\int_l^u x^{h+1}(1-x)^{n-h}dx}{\int_l^u x^h(1-x)^{n-h}dx}. \tag{7}$$

When $u = 1$ and $l = 0$, we can integrate by parts repeatedly and the surface terms (i.e., $x^i(1-x)^j$ evaluated at 0 and 1) vanish, yielding $p(h, n) = \frac{h+1}{n+2}$.

In the experiment, the distributions are discrete uniform. Equation (7) gives a close approximation for $u = U/n_{\max}$ and $l = L/n_{\max}$. The exact expression replaces the integrals by sums over $t = L, \ldots, U$ and replaces x by $x_t = t/n_{\max}$, viz.,

$$p(h, n) = \frac{\sum_{t=L}^U x_t^{h+1}(1-x_t)^{n-h}}{\sum_{t=L}^U x_t^h(1-x_t)^{n-h}}. \tag{8}$$

In the no-replacement case 4b, the likelihood is hypergeometric instead of binomial. The likelihood that there are exactly $t \le U \le 20$ treasures on the island, given that h were found on the first $n \le n_{\max} = 20$ tries, is

$$p(h, n \mid t) = \frac{\binom{n}{h}\frac{t!}{(t-h)!}\frac{(20-t)!}{(20-t-(n-h))!}}{\frac{20!}{(20-n)!}} \tag{9}$$

if $h \le t \le u$ and otherwise is 0. In the expression for the Bayesian posterior probability, both denominator (the normalizing constant) and numerator contain the binomial coefficient $\binom{n}{h}$, the expression $\frac{20!}{(20-n)!}$, and the constant prior probability $1/(U - L + 1)$. Hence these expressions cancel and we obtain the exact posterior distribution

$$f(t \mid h, n) = \frac{G(t \mid h, n)}{\sum_{s=L}^U G(s \mid h, n)}, \quad \text{where } G(t \mid h, n) = \frac{t!}{(t-h)!}\frac{(20-t)!}{(20-t-(n-h))!} \tag{10}$$

for $h, L \le t \le U$ and $h \le n$. Finally, the desired exact probability is the expectation of the remaining number of treasures (without replacement) divided by the remaining number of sites,

$$p(h, n) = \sum_{t=\max\{L,h\}}^{U} \frac{t-h}{20-n} f(t \mid h, n). \tag{11}$$

Bellman equations
We now are prepared to derive complete solutions for cases 2a, 2b, 4a and 4b. The approach is the same in each case: we write out the Bellman equation for optimal decision, insert appropriate boundary values and state transitions, and compute the values and contingent decisions by backward induction on the number of clicks remaining.

The Bellman equations take the following form.

$$V(l, s) = \max\{C(l, s), S(l), 0\}, \tag{12}$$

$$C(l, s) = E(x \mid s) + EV(l - 1, s'), \tag{13}$$

$$S(l) = -c + V(l, s_o). \tag{14}$$

The first line says that the value, i.e., the expected payoff over the rest of the voyage given l clicks remaining and state s, is the maximum obtainable from three options: clicking on the present island (C), skipping to the next island (S), or quitting immediately (0). The second line defines the click value recursively as the expected payoff from the next click $E(x \mid s)$, plus the expected value of continuing the voyage with one less click, taking into account the transition from the current state s to a new state s'. The third line defines the skip value as the value of starting on a new island (state s_o) less expected travel cost; note that it depends only on the number l of clicks remaining, and not on the current state s.

General boundary conditions include

1. $V(0, s) = 0$, i.e., the game is over when zero clicks remain; and
2. $E(x \mid s_m) = -\infty \ \forall s_m$ such that $n = n_{\max}$, i.e., only $n = n_{\max}$ clicks are permitted on each island.

In case 4a, the relevant state s is (h, n), the number of hits and clicks so far on the current island. Using $p = p(h, n)$ from Eq. (8), the click value for $l > 0$ clicks remaining and $n < n_{\max}$ is

$$C(l, s) = E(x \mid s) + EV(l - 1, s') \tag{15}$$

$$= pb + pV(l - 1, h + 1, n + 1)) + (1 - p)V(l - 1, h, n + 1). \tag{16}$$

The skip value here is

$$S(l) = -c + V(l, s_o) = -c + \max\{0,\ C(l, 0, 0)\} \tag{17}$$

where $s_o = (0, 0)$ refers to the state on arrival at a new island, 0 hits on 0 clicks. The last term uses only the click value, because skipping at $s_o = (0, 0)$ sinks the travel cost without improving prospects and thus is dominated by clicking or by quitting. Recall that the skip value is constant across states, while the click value obviously is increasing in h for given n. Hence the optimal decision typically is of the form: Click iff $h \geq h^*(n)$, for some reservation value function $h^*(n)$.

Case 4b, no replacement, is the same as case 4a except that the expectations in the click value use the probability $p = p(h, n)$ defined in Eq. (11). For reasons noted earlier, the probability here is not monotone in h for given n, and therefore the optimal decision cannot be characterized by a reservation value.

Recall that in Case 2a the island value $bw \in \{bL, \ldots, bU\}$ is observed, so the relevant state now is $s = (w, h, n)$. Hence in this case the expression $V(l, s_o)$ in Eq. (14) expands to $E_w V(l, w, 0, 0)$. The click value is the same as in Case 4a except that the hit probability now comes from Eq. (6). The skip value is more complicated because new islands with low w should be skipped immediately. The skip value satisfies

$$S(l) = -c + V(l, s_o) = -c + E_w V(l, w, 0, 0), \tag{18}$$

where the value of arriving at a new island of value bw satisfies

$$V(l, w, 0, 0) = \max\{C(l, w, 0, 0), S(l), 0\}. \tag{19}$$

The difficulty is that Eqs. (18) and (19) do not tell us directly whether to click or skip on a new island; the skip value $S(l)$ in (18) also enters the right hand side of (19).

To work it out, recall that the skip value to be determined is independent of the new island value w, while the click value $k_w = C(l, w, 0, 0)$ is increasing in w because it takes an expectation using probability $p = w/20$. Hence there is some threshold $w^*(l)$ such that optimally one clicks at least once on an island iff the displayed $w \geq w^*(l)$. The tentative skip value $T(w_o)$ is the value obtained using an arbitrary threshold $w_o \in \{L, \ldots, U\}$. By definition, $T(w_o) = -c + \frac{1}{U-L+1}[\sum_{w=L}^{w_o-1} T(w_o) + \sum_{w=w_o}^{U} k_w]$, so

$$T(w_o) = \frac{\sum_{w=w_o}^{U} k_w - (U - L + 1)c}{U - w_o + 1}. \tag{20}$$

The optimal threshold is the smallest number of treasures for which the tentative click surplus $K_w = k_w - T(w)$ is positive, i.e., $w^*(l) = \min\{w : K_w \geq 0\}$, and the true skip value is $S(l) = T(w^*(l))$.

The solution is straightforward to compute and well behaved because K_w is an increasing sequence in w that is positive for $w = U$. To see this, use (20) to write $K_w = k_w - \overline{k_{w+}} + b_w c$. The term $k_w - \overline{k_{w+}}$ is increasing because k_w increases faster than its (upper) average $\overline{k_{w+}} = \frac{\sum_{v=w}^{U} k_v}{U-w+1}$. The cost coefficient $b_w = \frac{U-L+1}{U-w+1}$ is also increasing in w. Clearly $K_U = 0 + (U - L + 1)c > 0$.

One last subcase remains. When the number of clicks remaining on arrival at a new island is not divisible by $n_{\max} = 20$, then Case 2a no longer reduces to Case 1. For example, one should not skip to the next island when the value of a new island is slightly below R when only 10 clicks remain, because the marginal benefit of skipping to a new island is depressed but the marginal cost is not. By backward induction, this complication also affects choices when 30 clicks remain on arrival to a new island, etc. Although every subject in the experiment begins every voyage with a multiple of 20 clicks, most subjects depart from optimal strategy at some point and find themselves in the situation just described.

We compute the (henceforth) optimal strategy in this subcase essentially the same way as in Case 2b. Use the displayed hit probability p to compute the click values

k_p. Note that the possible values of p are discrete multiples of 0.01 and, since $b = 5$, they can be indexed $i = 100p_i$, where $i = 5L, 5L + 1, \ldots, 5U$. Thus in this case, Eqs. (18) and (19) reduce to

$$S(l) = -c + E_p V(l, p, 0, 0) = -c + \sum_{i=5L}^{5U} \max\{k_{p_i}, S(l), 0\}. \tag{21}$$

Now (21) can be solved to yield

$$S(l) = \frac{\sum_{i=i^*}^{5U} k_{p_i} - (5U - 5L + 1)c}{5U - i^* + 1}, \tag{22}$$

where i^* is the threshold index, for which k_{p_i} first exceeds $S(l)$.

Acknowledgments The National Science Foundation funded this work under grant IIS-9986651. It builds on intellectual contributions by undergraduate researchers Dan Goldwasser and Mike McCaffrey, graduate research assistant Ralf Hepp, and programmers Nitai Farmer and Kannan Goundan. We also benefited from comments by numerous colleagues including Joshua Aizenman, Dan Ariely, Rachel Croson, Eric Eyster, Rob Fairlie, Federico Ravenna, Nirvikar Singh, and Donald Wittman. The final version benefited from the detailed comments of two anonymous referees. The usual caveat applies with unusual force.

References

Arkes, H.R., & Ayton, P. (1999). The sunk cost and concorde effects: Are humans less rational than lower animals? *Psychological Bulletin, 125*(5), 591–600.

Arkes, H.R., & Blumer, C. (1985). The psychology of sunk cost. *Organizational Behavior and Human Decision Processes, 35*, 124–140.

Aronson, E. (1968). Dissonance theory: Progress and problems. In R. Abelson et al (Eds.), *Theories of cognitive consistency: A Sourcebook*. pp. 5–27. Rand McNally and Co.

Barron, J.M., Chulkov, D.V., & Waddell, G.R. (2001). The escalation phenomenon and executive turnover: Theory and evidence. Mimeo, Purdue University.

Bazerman, M.H. (1986). *Judgment in managerial decision making* (2nd Ed. Chapter 4). John Wiley & Sons, Inc.

Camerer, C.F., & Weber, R.A. (1999). The econometrics and behavioral economics of escalation of commitment: A re-examination of Staw and Hoang's NBA data. *Journal of Economic Behavior and Organization, 39*(1), 59–82.

Carmichael, L., & MacLeod, W.B. (2003). Caring about sunk costs: A behavioral solution to the hold-up problem. *Journal of Law, Economics and Organization, 19*(1), 106–118.

Cox, J., & Oaxaca, R. (1989). Laboratory experiments with a finite-horizon job-search model. *Journal of Risk and Uncertainty, 2*, 301–330.

Cox, J., & Oaxaca, R. (1992). Direct tests of the reservation wage property. *The Economic Journal, 102*, 1423–1432.

Cura-Juri, M., & Galiani, S. (2004). In D. Friedman & A. Cassar (Eds.), *Economics lab: An intensive course in experimental economics*, Chapter 18. London and New York, Routledge.

Dixit, A., & Pindyck, R.H. (1994). *Investment under uncertainty*. Princeton University Press.

Economist. (2003). The space shuttle: Old, unsafe and costly. *Economist* magazine. , pp. 77.

Elliott, S.R., & Curme, M. (1998). Understanding the sunk cost effect: An experimental approach. Mimeo, Miami University.

Eyster, E. (2002). Rationalizing the past: A taste for consistency. Nuffield College, Oxford, Job-Market Paper.

Festinger, L. (1957). *A Theory of cognitive dissonance*. Stanford, Stanford University Press.

Garland, H., & Conlan, D.E. (1998). Too close to quit: The role of project completion in maintaining commitment. *Journal of Applied Social Psychology, 28*(22), 2025–2048.

Garland, H., & Newport, S. (1991). Effects of absolute and relative sunk costs on the decision to persist with a course of action. *Organizational Behavior and Human Decision Processes, 48*, 55–69.
Harrison, G.W., & Morgan, P. (1990). Search intensity in experiments. *Economic Journal, 100*, 478–486.
Heath, C. (1995). Escalation and *De*-escalation of commitment in response to sunk costs: The role of budgeting in mental accounting. *Organizational Behavior and Human Decision Processes, 62*(1), 38–54.
Hey, J.D. (1982). Search for rules for search. *Journal of Economic Behavior and Organization, 3*, 65–81.
Hey, J.D. (1987). Still searching. *Journal of Economic Behavior and Organization, 8*, 137–144.
Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*(2), 263–291.
Kanodia, C., Bushman, R., & Dickhaut, J. (1989). Escalation erros and the sunk cost effect: An explanation based on reputation and information asymmetries. *Journal of Accounting Research, 27*(1), 59–77.
Klein, J., Moon, Y., & Picard, R. W. (1999). This computer responds to user frustration. MIT Media Laboratory. 1999 Vision and Modeling Group Technical Report #502.
Kogut, C. A. (1990). Consumer search behavior and sunk costs. *Journal of Economic Behavior and Organization, 14*, 381–392.
Manley, S., & Seltzer, M. (1997, December 1997). Web facts and fantasy. *Proceedings of the 1997 USENIX Symposium on Internet Technologies and Systems*, Monterey, CA.
Meyer, D. J. (1993). First price auctions with entry: An experimental investigation. *The Quarterly Journal of Economics and Finance, 33*(2), 107–122.
Milgrom, P., & Roberts, J. (1992). *Economics, organization, and management.* Prentice-Hall, Englewood Cliffs, NJ.
Northcraft, G.B., & Neale, M.A. (1986). Opportunity costs and the framing of resource allocation decisions. *Organization Behavior and Human Decision Processes, 37*, 348–356.
Offerman, T., & Potters, J. (2001). Does auctioning of entry licenses induce collusion? An experimental study. Working paper, University of Amsterdam.
Phillips, O.R., Battalio, R.C., & Kogut, C.A. (1991). Sunk and opportunity costs in valuation and bidding. *Southern Economic Journal, 58*(1), 112–128.
Schotter, A., & Braunstein, Y.M. (1981). Economic search: An experimental study. *Economic Inquiry, 19*, 1–25.
Schwartz, E.I. (1999, June 8). In E.I. Schwartz (Ed.), Digital Darwinism: 7 Breakthrough Business Strategies for Surviving in the Cutthroat Web Economy. Broadway Books.
Sonnemans, J. (1998). Strategies of search. *Journal of Economic Behavior and Organziation, 35*(3), 309–332.
Sonnemans, J. (2000). Decision and strategies in a sequential search experiment. *Journal of Economic Psychology, 21*, 91–102.
Staw, B.M. (1976). Knee deep in the big muddy: A study of escalating commitment to a chosen course of action. *Organizational Behavior and Human Performance, 16*, 27–44.
Staw, B.M., & Hoang, H. (1995). Sunk costs in the NBA: Why draft order affects playing time and survival in professional basketball. *Administrative Science Quarterly, 40*, 474–494.
Tan, H.-T., & Yates, F.J. (1995). Sunk cost effects: The influences of instruction and future return estimates. *Organizational Behavior and Human Decision Processes, 63*(3), 311–319.
Thaler, R. (1980). Toward a positive theory of consumer choice. *Journal of Economic Behavior and Organization, 1*, 39–60.
Trivers, R.L. (1972). Parental investment and sexual selection. In B. Campbell (Ed.), *Sexual selection and the descent of man 1871–1971* (pp. 136–179). Chicago: Aldine.
Whyte, G. (1993). Escalating commitment in individual and group decision making: A prospect theory approach. *Organizational Behavior and Human Decision Processes, 54*(3), 430–455.
Whyte, G., Sacks, A.M., & Hook, S. (1997). When success breeds failure: The role of self-efficacy in escalating commitment to a losing course of action. *Journal of Organizational Behavior, 18*(5), 415–432.