The Thrill of Gradual Learning*

Faruk Gul†

Paulo Natenzon‡

Erkut Y. Ozbay§

Wolfgang Pesendorfer¶

Abstract

We report on an experiment that shows subjects prefer a gradual resolution of uncertainty when information about winning yields decisive bad news but inconclusive good news. This behavior is difficult to reconcile with existing theories of choice under uncertainty, including the Kreps-Porteus model. We show how the behavioral patterns uncovered by our experiment can be understood as arising from subjects' special emphasis on their best (peak) and worst (trough) experiences along the realized path of uncertainty.

*Date: May 4, 2022. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper. This research was supported by the NSF grant SES-1729021.

†Princeton University. Email: fgul@princeton.edu

‡Washington University in St. Louis. Email: pnatenzon@wustl.edu

§University of Maryland. Email: ozbay@umd.edu

¶Princeton University. Email: pesendor@princeton.edu

1 Introduction

Consider the following situation: a decision maker places a bet that pays off if the event W occurs. At some specified time T, the state of the world is revealed and the decision maker receives the prize if the state is in W. Prior to time T, the decision maker may learn about the likelihood of W. For example, she may learn whether or not W has occurred in some period prior to T; alternatively, she may receive partial information or no information at all.

Important determinants of a person's utility, such as the person's health, wealth or educational attainment, take the form of a lottery that will resolve at some future date. Information about these outcomes, even if decision irrelevant, will affect the person's utility and it seems plausible that utility would also depend on how such information is disclosed. However, the manner in which decision-irrelevant information is disclosed has no effect on the utility of a standard expected utility maximizer; this agent would be indifferent between learning the outcome immediately, not learning anything prior to period T or any other form of disclosure.

Contrary to the predictions of the standard model, experimental evidence suggests that many decision makers value non-instrumental information; they often prefer early over late disclosure (e.g. Ahlbrecht and Weber, 1997; Eliaz and Schotter, 2010; Falk and Zimmermann, 2016); though in some situations they have a clear preference for late resolution (e.g. Oster, Shoulson and Dorsey, 2013; Ganguly and Tasoff, 2017). Golman, Hagmann and Loewenstein (2017) provide an extensive survey of the experimental literature on preferences for timing of resolution of uncertainty.

This paper adds to this evidence and uncovers new patterns of demand for non-instrumental information. In our experiment, subjects exhibit strict preferences with respect to *gradual* information disclosure. Specifically, subjects prefer a process in which the chance of winning a monetary prize goes up gradually or drops to zero all at once, to a process that resolves all uncertainty in a single moment. Thus, subjects opt for gradual over immediate information disclosure if the gradual option reveals *partial good news or decisive bad news.* By contrast, we find that subjects are more inclined to choose immediate disclosure when the gradual option is of the *decisive good news* or *inconclusive bad news* variety.

Our findings are likely to be relevant for welfare analysis in information design problems outside the lab. For example, "up or out" HR policies determine that, at specific time intervals, employees that are on track for a large future promotion (such as making partner in a firm) must either get a smaller promotion or be fired. Such policies resemble the partial good news and decisive bad news disclosure policy favored by the majority of our subjects. In healthcare, patients may opt for a more invasive procedure that provides immediate disclosure and forgo safer or less invasive diagnostic procedures when they provide information of the gradual bad news or decisive good news variety. Oster (2014) provides a pre-natal testing example: "If I was probably going to do that [the more invasive procedure] anyway, why should I go through the anxiety of being told I had a bad screening result, then worry for weeks before getting a final answer? (p. 108)"

Models that extend standard expected utility framework to account for the value of non-instrumental information can be useful for guiding the design of information disclosure policies in the examples above. They are also important for understanding the link between information feedback and the behavior of agents in financial markets. Our experimental design, described in detail below, is extremely simple and offers an ideal setting for calibrating and testing such models.

The behavior observed in our experiment, however, is difficult to reconcile with standard theories of choice and even with many of the alternatives that have been developed more recently. For example, preferences over temporal lotteries in Kreps and Porteus (1978), preferences for one-shot resolution of uncertainty in Dillenberger (2010), and surprise and suspense utility in Ely et al. (2015) require the direction of preference to be the same across our decision problems —either always preferring gradual or always preferring immediate disclosure— and therefore cannot account for our main findings.

To interpret the observed behavior, we adopt the framework developed in Gul et al. (2021) which departs from expected utility in two ways: first, as in Kreps and Porteus (1978), subjects may have a categorical preference for early or late resolution of uncertainty. Second, subjects may place special emphasis on their best and worst experiences. The resulting utility, *peak-trough utility*, allows us to rationalize and interpret the observed patterns of behavior.

Kreps and Porteus (1978) develop the first theory of preference for timing of resolution of uncertainty. Their model is well-suited for addressing a categorical preference for early or late resolution of uncertainty; it is less well-suited for dealing with more nuanced attitudes to information. As noted by Nielsen (2020) and confirmed in our experiment, many subjects have a preference for *gradual resolution of uncertainty*, a behavior that is difficult to reconcile with the Kreps-Porteus model. Ely, Frankel and Kamenica (2015) offer an alternative to the Kreps-Porteus model with different choice objects. They provide formal definitions of suspense and surprise and show how the two determine the decision maker's attitude toward non-instrumental information. While the Ely et al. model is consistent with a preference for gradual resolution of uncertainty, it cannot accommodate the strict preference for gradual good news over gradual bad news that plays an important role in our experiment. We discuss the predictions of these models in detail below.

Our framework builds on Caplin and Leahy (2001) who present the first model of anticipatory utility. Palacios-Huerta (1999) shows that a two-stage nonexpected utility model can lead to a preference for one-shot resolution of uncertainty. Building on this insight, Dillenberger (2010) develops a general theory of preference for one-shot resolution of uncertainty. Kőszegi and Rabin (2009) formulate a theory of reference-dependent utility for analyzing attitudes towards non-instrumental information. They identify conditions guaranteeing a preference for one-shot resolution of uncertainty and conditions that yield a preference for early resolution.

Experimental research that analyzes the type of non-instrumental information favored by subjects includes Masatlioglu, Orhun and Raymond (2017) who find a preference for positively skewed information. Nielsen (2020) observes that more subjects chose one-shot late resolution of uncertainty when the options are presented as two-stage lotteries than when they are presented as information structures. Conversely, more subjects chose one-shot early resolution of uncertainty when the options are presented as information structures than when they are presented as two-stage lotteries. Nevertheless, Nielsen concludes that "gradual resolution generally is preferred to one-shot resolution in both treatments."

The rest of the paper proceeds as follows. In Section 2, we describe the experimental design. In Section 3, we provide a framework for understanding the choice options in our experiments as risk consumption paths. We obtain the theoretical predictions of existing models in Section 4. We then present the results of the experiment in Section 5, and make our concluding remarks in Section 6.

2 Experimental Design

A total of 125 University of Maryland undergraduate students participated in the experiment conducted in the Experimental Economics Laboratory at the University of Maryland. We had 8 sessions in total (6 sessions with 16 subjects, 1 session with 15 subjects and 1 session with 14 subjects).

Subjects were recruited through ORSEE (Greiner, 2015). The experiment was programmed in zTree (Fischbacher, 2007). A typical session lasted around 30 minutes. The instructions were incorporated into the experiment. (Complete instructions and the screenshots may be found in the Online Appendix.) On average, a total of 10 minutes of a typical session was instructional. Subjects earned an average of $12.40, including a $7 show-up fee, paid in cash privately at the end of the experiment.

Subjects were asked to choose one of three boxes on their screen and then choose the manner in which the content of the boxes are to be revealed. Every subject confronted four decision problems:

G1. Gradual resolution, one box contains a prize of $10;

G2. Gradual resolution, two boxes contain a prize of $10 each;

O1. One-shot resolution, one box contains a prize of $10;

O2. One-shot resolution, two boxes contain a prize of $10 each.

One of these decision problems was chosen for implementation at the end of the experiment. The subjects earned $10 if the box that they selected contained $10, and $0 otherwise. The decision problems were presented in random order during the sessions.

In decision problems G1 and G2, the contents of the boxes were revealed to the subjects one after the other, with a 60-second delay between boxes. After choosing their boxes, subjects decided whether they wanted their box to be opened early or late. For concreteness, we will call the box that the subject chose box 1 and the others box 2 and box 3. The option *early* means that the experiment reveals the content of box 1 first, one minute later reveals the content of box 2, and one minute after that the content of box 3. Choosing *late* means that the experiment reveals the content of box 2 first. Then, after a 60-second delay, the experiment reveals the content of box 3 and 60 seconds after that the content of box 1.

Note that subjects learn the outcome once the first two boxes are opened. The experiment reveals the content of the third box for the sake of transparency. In decision problems O1 and O2, the contents of all boxes are revealed at the same time. After selecting their boxes, subjects chose whether they wanted all of the boxes to be opened at the start or at the end of a 120-second waiting period. To check the strength of the preference, we employed the standard choice list procedure (e.g. Epstein and Halevy, 2018). Ten different amounts of compensation ranging from 1 cent to 50 cents were offered. This procedure is identical to the willingness to switch elicitation procedure used in Masatlioglu, Orhun and Raymond (2017).

Once the answers were collected, the computer randomly picked one of the four decision problems and one of the 10 price list questions. If the DM had stated that she is unwilling to switch her initial choice at the randomly picked level of compensation, the boxes were opened in the manner that she had chosen initially and the DM received no additional compensation. Otherwise; that is, if the DM had stated that she would accept the randomly picked level of compensation, the boxes were opened in the manner that she had not chosen initially and the DM received the additional compensation. All subjects waited for 120 seconds until the experiment ended.

3 A Framework for Risk Consumption

In this section, we provide a framework for risk consumption; that is, the idea that agents derive utility from the evolution of a risky prospect over time. The decision maker (DM) receives, in period N, a prize from a finite set, A, of prizes. In each period, $t = 1, \ldots, N$, the DM faces a lottery α_t; that is, a probability distribution over prizes. This lottery evolves over time as the DM receives new information. Hence, we call the resulting path of lotteries, $(\alpha_1, \ldots, \alpha_N)$, an *evolving lottery*. Below, we use the terms "evolving lottery" and "path" interchangeably.

Each decision problem in our experiment presents the subject with two options. Both options yield the same probability of winning \$10. However, they lead to different *random evolving lotteries*, that is, probability distributions over evolving lotteries. In each decision problem, let P denote the random evolving lottery associated with the early option, and let Q denote the random evolving lottery associated with the late option.

Consider decision problem G1: the initial probability of winning the \$10 prize is 1/3 and, therefore, in the first period of every path, the probability of winning the prize is $\alpha_1 = 1/3$. Suppose the subject chooses the early option; that is, the random evolving lottery P. Then, in period 2, the subject learns whether or not they won the prize and therefore, α_2 is either zero or one. Hence, P has two paths $(1/3, 1, 1)$ and $(1/3, 0, 0)$ and assigns to them the following probabilities:

$$P(1/3, 1, 1) = 1/3 \text{ and } P(1/3, 0, 0) = 2/3.$$

If the subject chooses late option; that is, chooses Q, information is revealed gradually. The initial probability of winning is, again, $\alpha_1 = 1/3$. If box 2, the box opened in period 2, contains the prize, then the subject will learn that she has lost once it is opened and $\alpha_2 = \alpha_3 = 0$; if box 2 does not contain the prize, then the probability of winning rises to $\alpha_2 = 1/2$. Since each box is equally likely to contain the prize, the probability that $\alpha_2 = 1/2$ is 2/3. If $\alpha_2 = 1/2$, all uncertainty is resolved when box 3 is opened. Therefore, $\alpha_2 = 1/2$ will be followed either by $\alpha_3 = 1$ or $\alpha_3 = 0$. Thus, Q, the random evolving lottery associated with the gradual resolution of uncertainty is as follows:

$$Q(1/3, 1/2, 1) = Q(1/3, 1/2, 0) = Q(1/3, 0, 0) = 1/3.$$

Notice that P and Q are the only two random evolving lotteries that can be generated by varying the order in which the boxes are opened.

Decision problem G2 is identical to G1 except that two of three boxes contain a prize. In this case, the random evolving lottery, P associated with the early option assigns the following probabilities:

$$P(2/3, 1, 1) = 2/3 \text{ and } P(2/3, 0, 0) = 1/3,$$

while opening late option results in the random evolving lottery Q such that

$$Q(2/3, 1/2, 1) = Q(2/3, 1/2, 0) = Q(2/3, 1, 1) = 1/3.$$

Decision problem O1 offers a simple timing trade-off. One choice reveals all information in period 2 while the other reveals all information in period 3. Since only one box contains a prize, the first choice leads to a random evolving lottery P below:

$$P(1/3, 1, 1) = 1/3 \text{ and } P(1/3, 0, 0) = 2/3$$

while the second choice leads to a random evolving lottery Q:

$$Q(1/3, 1/3, 1) = 1/3 \text{ and } Q(1/3, 1/3, 0) = 2/3.$$

Finally, decision problem O2 offers the same simple timing trade-off as O1 but with two boxes containing a prize. The early option leads to the random evolving lottery P:

$$P(2/3, 1, 1) = 2/3 \text{ and } P(2/3, 0, 0) = 1/3$$

while "late" leads to the random evolving lottery Q:

$$Q(2/3, 2/3, 1) = 2/3 \text{ and } Q(2/3, 2/3, 0) = 1/3.$$

4 Predictions from theory

A standard expected utility maximizer whose only concern is the ultimate outcome and who does not care about how uncertainty resolves would identify each path $\alpha = (\alpha_1, \ldots, \alpha_N)$ with α_1 and assign to α the expected utility of α_1. Hence, a standard DM with von Neumann-Morgenstern utility should be exactly indifferent in every choice problem of our experiment.

Kreps-Porteus preferences

Kreps and Porteus (1978) develop the first model that permits a preference for early or late resolution of uncertainty. Their choice objects are *temporal lotteries* rather than random evolving lotteries. A one-stage temporal lottery is simply a probability distribution over prizes. Then, we define a t-stage temporal lottery inductively as a lottery over $(t-1)$-stage temporal lotteries.

In our experimental setting, each choice can be mapped to a temporal lottery. For example, gradual resolution in G1 corresponds to the following temporal lottery: in period 2, the agent can encounter two possible one-stage lotteries; ℓ_1 yields the prize with probability 0; while ℓ_2 yields the prize with probability 1/2. In period 1, the agent has a 2-stage temporal lottery that yields one-stage lottery ℓ_1 with probability 1/3 and one-stage lottery ℓ_2 with probability 2/3. Let L denote this temporal lottery.

Normalize the utility of \$10 to 1 and the utility of \$0 to zero so that $u_2(\ell_1) = 0, u_2(\ell_2) = 1/2$ are the two possible payoff realizations in period 2. Then, a period-1 aggregator $u_1 : [0, 1] \to [0, 1]$ such that $u_1(1) = 1, u_1(0) = 0$ determines the utility of a Kreps-Porteus agent as follows:

$$U(L) = \frac{1}{3}u_1(u_2(\ell_1)) + \frac{2}{3}u_1\left(u_2(\ell_2)\right) = \frac{1}{3}u_1(0) + \frac{2}{3}u_1\left(\frac{1}{2}\right) = \frac{2}{3}u_1\left(\frac{1}{2}\right)$$

Table 1, below, shows the Kreps-Porteus utilities of all choices in our experiment.

Table 1: Kreps-Porteus utilities

	G1	G2	O1	O2
P *(early)*	$\frac{1}{3}$	$\frac{2}{3}$	$\frac{1}{3}$	$\frac{2}{3}$
Q *(late)*	$\frac{2}{3}u_1\left(\frac{1}{2}\right)$	$\frac{2}{3}u_1\left(\frac{1}{2}\right) + \frac{1}{3}$	$u_1\left(\frac{1}{3}\right)$	$u_1\left(\frac{2}{3}\right)$

The curvature of u_1 governs the behavior of Kreps-Porteus agents. If u_1 is convex, then the agent *always* prefers early resolution while if u_1 is concave the agent always prefers late resolution. More nuanced behavior is possible if u_1 is neither convex nor concave. For example, if $u\left(\frac{1}{3}\right) > \frac{1}{3}$ and $\frac{2}{3} > u\left(\frac{2}{3}\right)$, then the agent prefers early resolution in O1 but late in O2. Note that Kreps-Porteus subjects prefer the early choice over the late choice in G1 and G2 if and only if $\frac{1}{2} \geq u\left(\frac{1}{2}\right)$. Thus, the main testable prediction of the Kreps-Porteus model in our experiment is that the decision maker should behave identically across problems G1 and G2: either P is strictly preferred in both problems, or Q is strictly preferred in both problems, or else the decision maker must be indifferent in both problems.

Preference for one-shot resolution

Palacios-Huerta (1999) shows that a two-stage nonexpected utility model can lead to a preference for one-shot resolution of uncertainty. Building on this insight, Dillenberger (2010) develops a general theory of preference for one-shot resolution of uncertainty, which provides two testable implications in our experiment. The first one is time neutrality, which says that, as long as all uncertainty resolves in a single moment, the decision maker does not care when it resolves. Therefore, the decision maker must be indifferent in problems O1 and O2. The second is preference for one-shot resolution, which requires that the decision maker find the early option at least as good as the late option in both G1 and G2.

Surprise and Suspense

Ely, Frankel and Kamenica (2015) propose two measures of the entertainment value of gradual information revelation: *suspense* and *surprise*. Suspense is the expected belief variation next period while surprise is the realized belief variation from period to period. Specifically, the surprise utility of a belief path $\alpha = (\alpha_1, \alpha_2, \alpha_3)$ in our experiment is given by

$$w_{\text{surprise}}(\alpha) = \sum_{t=2}^{3} u\left(\|(\alpha_t, 1-\alpha_t) - (\alpha_{t-1}, 1-\alpha_{t-1})\|^2\right)$$

where u is increasing and strictly concave with $u(0) = 0$. The suspense utility of a belief path α depends on the random evolving lottery P that assigns probabilities to each path in the decision problem faced by the agent,

$$w_{\text{suspense}}(\alpha, P) = \sum_{t=1}^{2} u\left(\mathbb{E}_P \|(\alpha_{t+1}, 1-\alpha_{t+1}) - (\alpha_t, 1-\alpha_t)\|^2\right)$$

where again u is increasing and strictly concave with $u(0) = 0$. The surprise/suspense utility of a random evolving lottery P is the expected value of the surprise/suspense utility of a path, with probabilities given by P.

In Tables 2 and 3 we calculate the suspense utility and the surprise utility of the options in each of our decision problems. The decision problems G1 and G2 are mirror images of one another: we can derive the random evolving lottery associated with each option in G2 from the corresponding choice in G1 by replacing every probability α_t in the former with $1 - \alpha_t$. An immediate implication of this symmetry is that suspense and surprise utility cannot distinguish between decision problems G1 and G2 and between decision problems O1 and O2.

Proposition 1. *Surprise and suspense utility agents prefer the early option P to the late option Q in problem G1 if and only if they prefer P to Q in problem G2.*

Proposition 2. *Surprise and suspense utility agents are indifferent in problems O1 and O2.*

The proof of Propositions 1 and 2 is immediate from the comparison of utilities calculated in Tables 2 and 3. For surprise utility, we obtain a sharper prediction for G1 and G2:

Proposition 3. *Surprise utility agents strictly prefer the late option Q in both G1 and G2.*

Proof. The index u in surprise utility is increasing and strictly concave, with $u(0) = 0$. Hence,

$$\begin{aligned} u(1/18) + u(1/2) &= [u(1/18) - u(0)] + [u(1/2) - u(1/9)] + u(1/9) \\ &> [u(1/9) - u(1/18)] + [u(8/9) - u(1/2)] + u(1/9) \\ &> u(2/9) - u(1/18) + u(8/9) - u(1/2) \end{aligned}$$

which yields $\frac{2}{3}u\left(\frac{1}{18}\right) + \frac{2}{3}u\left(\frac{1}{2}\right) + \frac{1}{3}u\left(\frac{2}{9}\right) > \frac{2}{3}u\left(\frac{2}{9}\right) + \frac{1}{3}u\left(\frac{8}{9}\right)$ as desired. □

Table 2: Suspense utilities

	G1 and G2	O1 and O2
P *(early)*	$u\left(\frac{4}{9}\right)$	$u\left(\frac{4}{9}\right)$
Q *(late)*	$u\left(\frac{1}{9}\right) + \frac{2}{3}u\left(\frac{1}{2}\right)$	$u\left(\frac{4}{9}\right)$

Table 3: Surprise utilities

	G1 and G2	O1 and O2
P *(early)*	$\frac{2}{3}u\left(\frac{2}{9}\right) + \frac{1}{3}u\left(\frac{8}{9}\right)$	$\frac{2}{3}u\left(\frac{2}{9}\right) + \frac{1}{3}u\left(\frac{8}{9}\right)$
Q *(late)*	$\frac{2}{3}u\left(\frac{1}{18}\right) + \frac{2}{3}u\left(\frac{1}{2}\right) + \frac{1}{3}u\left(\frac{2}{9}\right)$	$\frac{2}{3}u\left(\frac{2}{9}\right) + \frac{1}{3}u\left(\frac{8}{9}\right)$

In an extension, Ely, Frankel and Kamenica (2015) allow the agent to care more about the suspense and surprise associated with one outcome than the suspense and surprise associated with another. This modification generalizes the model when there are more than two outcomes but has no effect in our setting with only two outcomes. In the binary setting, changes in the probability of winning must coincide with changes in the probability of losing and, thus, as long as utility depends only on belief-variation the model cannot capture differences in subjects' behavior between decision problems G1 and G2 and between decision problems O1 and O2.

In another extension of their model, Ely, Frankel and Kamenica (2015) introduce discounting and allow the agent to give different weights to surprise and suspense delivered in different periods. This extension breaks the prediction of indifference between early and late resolution, and can accommodate a strict preference in problems O1 and O2, as long as that preference is the same across both problems.

Peaks and Troughs

Fredrickson and Kahneman (1993) argue that, in retrospective evaluations, subjects neglect the duration of experiences and emphasize extremes. Peak-trough utility theory (Gul et al., 2021) offers a simple theoretical framework in which agents place special emphasis on their best and worst experiences. A path utility, w, assigns a value to each path and the utility of a random evolving lottery, P, is the expected utility of these paths:

$$W(P) = \sum_{\alpha} w(\alpha) P(\alpha). \tag{1}$$

To define the path utility function of peak-trough utility, let u be the DM's expected utility function over lotteries. Then, define the peak, $\overline{u}$, and trough, $\underline{u}$, utilities of any path $\alpha = (\alpha_1, \ldots, \alpha_N)$ as follows:

$$\begin{aligned} \overline{u}(\alpha_1, \ldots, \alpha_N) &= \max_t u(\alpha_t) \\ \underline{u}(\alpha_1, \ldots, \alpha_N) &= \min_t u(\alpha_t) \end{aligned}$$

Let $v : [0,1] \to [0,1]$ be a strictly increasing, continuous and onto function and let θ_h, θ_ℓ be weights such that

$$\begin{aligned} 1-\theta_h-\theta_\ell &> 0 \\ (1-\theta_h-\theta_\ell)/N+\theta_h &> 0 \\ (1-\theta_h-\theta_\ell)/N+\theta_\ell &> 0 \end{aligned}$$

Then, the utility of any path $\alpha = (\alpha_1, \ldots, \alpha_n)$ is

$$w(\alpha) = \frac{1-\theta_h-\theta_\ell}{N} \sum_{t=1}^{N} v(u(\alpha_t)) + \theta_h v(\overline{u}(\alpha_1, \ldots, \alpha_N)) + \theta_\ell v(\underline{u}(\alpha_1, \ldots, \alpha_N)) \tag{2}$$

The function W defined in (1) is a peak-trough utility whenever w is as described in (2). Compared to the standard expected utility model, a peak-trough utility has three new parameters; the index v which determines the agent's preference for early or late resolution of uncertainty and the weights θ_h, θ_ℓ which determine the agent's sensitivity to the best and worst experiences. If we assume v is linear and set the weights $\theta_h = \theta_\ell = 0$, then we are back to the standard model.[1]

It is easy to relate the parameters of peak-trough utility to the decision problems in our experiment. Consider, for example, the decision problem G1. Table 4 shows the distribution of peaks and troughs for the early resolution option P and the late resolution option Q. Note that the distribution of path troughs is identical for P and Q: both offer a trough of zero with probability 2/3 and a trough of 1/3 with probability 1/3. Therefore, the value of θ_ℓ plays no role in their comparison.

Table 4: Distribution of path peaks and path troughs in decision problem G1

Path	Peak	Trough	P	Q
(1/3, 0, 0)	1/3	0	2/3	1/3
(1/3, 1, 1)	1	1/3	1/3	–
(1/3, 1/2, 1)	1	1/3	–	1/3
(1/3, 1/2, 0)	1/2	0	–	1/3

[1]For simplicity, we assumed that each period has the weight $1/N$. A more general model would permit discounting. Gul et al. (2021) provide an axiomatic foundation for this, more general, model.

The following proposition relates the optimal choice in decision problem G1 to the parameters v and θ_h of peak-trough utility.

Proposition 4. *Let W be a peak-trough utility with parameters $(v, \theta_h, \theta_\ell)$, and let P and Q be the early and late options, respectively, in decision problem G1.*

(i) *If v is linear, then $W(Q) > W(P)$ if and only if $\theta_h > 0$.*

(ii) *If v is convex, then $W(Q) > W(P)$ implies $\theta_h > 0$*

(iii) *If v is concave, then $W(P) > W(Q)$ implies $\theta_h < 0$.*

Proof. The early option P assigns probability 1/3 to the path $(1/3, 1, 1)$ and probability 2/3 to the path $(1/3, 0, 0)$. On the other hand, the late option Q gives probability 1/3 to each of the three paths $(1/3, 1/2, 1)$, $(1/3, 1/2, 0)$ and $(1/3, 0, 0)$. Applying the path utility formula (2) and using the fact that $v(0) = 0$ and $v(1) = 1$ we obtain

$$\begin{aligned} W(Q) - W(P) &= \frac{1}{3}\left[w\left(\frac{1}{3}, \frac{1}{2}, 1\right) + w\left(\frac{1}{3}, \frac{1}{2}, 0\right) + w\left(\frac{1}{3}, 0, 0\right) - w\left(\frac{1}{3}, 1, 1\right) - 2w\left(\frac{1}{3}, 0, 0\right)\right] \\ &= \frac{1 - \theta_h - \theta_\ell}{9}\left[2v\left(\frac{1}{2}\right) - 1\right] + \frac{\theta_h}{3}\left[v\left(\frac{1}{2}\right) - v\left(\frac{1}{3}\right)\right] \end{aligned}$$

If v is linear then $W(Q) - W(P)$ above simplifies to $\frac{\theta_h}{18}$ and, therefore, part (i) follows. If v is convex, then $2v(1/2) \leq 1$ and, therefore, $W(Q) - W(P) > 0$ implies $\theta_h > 0$ which proves part (ii). If v is concave, then $2v(1/2) \geq 1$ and, therefore, $W(P) - W(Q) > 0$ implies $\theta_h < 0$ which proves part (iii). □

Subjects with $\theta_h < 0$ dislike getting their hopes up. Therefore, they tend to prefer the early option in G1. Subjects with $\theta_h > 0$ enjoy paths that look promising even if things don't pan out in the end. These subjects tend to prefer the late option in G1. As Table 4 shows, the late option Q is equally likely to yield the path peaks 1, 1/2 and 1/3. The early option P has the path peaks 1/3 and 1 with the former being twice as likely as the latter. Thus, Q offers a better distribution of path peaks (in the sense of first order stochastic dominance) than P.

Next, consider decision problem G2, which is similar to G1 but now two of the three boxes contain the prize. Table 5 shows the early option P and the late option Q generate the same distribution of path peaks: the peak is 2/3 with probability 1/3 and it is 1 the remaining 2/3 of the time. Therefore, θ_h plays no role in the comparison P versus Q.

Table 5: Distribution of path peaks and path troughs in decision problem G2

Path	Peak	Trough	P	Q
(2/3, 0, 0)	2/3	0	1/3	–
(2/3, 1, 1)	1	2/3	2/3	1/3
(2/3, 1/2, 1)	1	1/2	–	1/3
(2/3, 1/2, 0)	2/3	0	–	1/3

The next proposition relates the remaining parameters v and θ_ℓ of peak-trough utility to the choice in decision problem G2.

Proposition 5. *Let W be a peak-trough utility with parameters $(v, \theta_h, \theta_\ell)$, and let P and Q be the early and late options, respectively, in decision problem G2.*

(i) *If v is linear, then $W(Q) > W(P)$ if and only if $\theta_\ell < 0$.*

(ii) *If v is convex, then $W(Q) > W(P)$ implies $\theta_\ell < 0$*

(iii) *If v is concave, then $W(P) > W(Q)$ implies $\theta_\ell > 0$.*

Proof. The peak-trough utility agent prefers the late option Q whenever the difference

$$W(Q) - W(P) = \frac{1 - \theta_h - \theta_\ell}{9}[2v(1/2) - 1] + \frac{\theta_\ell}{3}[v(1/2) - v(2/3)] \tag{3}$$

is positive. If v is linear, then (3) simplifies to $-\frac{\theta_\ell}{18}$ and, therefore, part (i) follows. If v is convex, then $2v(1/2) \leq 1$ and, therefore, $W(Q) - W(P) > 0$ implies $\theta_\ell < 0$ which proves part (ii). If v is concave, then $2v(1/2) \geq 1$ and, therefore, $W(P) - W(Q) > 0$ implies $\theta_\ell > 0$ which proves part (iii). □

Subjects with $\theta_\ell > 0$ enjoy paths with comebacks; that is, they like paths that end well despite the good outcome seeming unlikely at an earlier stage. Those with $\theta_\ell < 0$ dread such paths. The random evolving lottery Q differs from P in the distribution of troughs. As Table 5 shows, Q is equally likely to yield the troughs 2/3, 1/2, and 0, while P yields the trough 2/3 with probability 2/3 and 0 with probability 1/3. As we noted above, the two random evolving lotteries offer the same distribution of peaks and therefore, θ_h plays no role in their comparison.

In decision problem O1 (and O2), the agent considers a simple timing trade-off. One choice reveals all information in period 1 while the other reveals all information in period 3. Table 6 shows the distribution of peaks and troughs is not affected by the decision. Therefore the weights θ_ℓ and θ_h play no role in the comparison. And, as the following proposition shows, only the curvature of v matters.

Table 6: Distribution of path peaks and path troughs in decision problem O1

Path	Peak	Trough	P	Q
(1/3, 0, 0)	1/3	0	2/3	–
(1/3, 1, 1)	1	1/3	1/3	–
(1/3,1/3,1)	1	1/3	–	1/3
(1/3,1/3,0)	1/3	0	–	2/3

Proposition 6. *Let W be a peak-trough utility with parameters $(v, \theta_h, \theta_\ell)$ and let P and Q be the early and late options, respectively, in decision problem O1. If v is convex (concave) then $W(P) > W(Q)$ ($W(P) < W(Q)$).*

Proof. The difference in utility between the early option P and the late option Q in problem O1 is given by,

$$W(P) - W(Q) = \frac{1 - \theta_h - \theta_\ell}{3}\,[c - v(c)] \tag{4}$$

Since $v(c) \geq c$ if v is concave and $v(c) \leq c$ if v is convex, the result follows from (4). □

5 Experimental Results

We first analyze the subjects' strength of preferences. Figure 1 shows the distribution of the required compensation for switching away from the preferred mode of information disclosure. We exclude 3 of the 125 subjects because their responses to the strength of preference question were not monotone. As Figure 1 shows, depending on the decision problem, 16–25% of subjects are *indifferent*; that is, are willing to switch their choice if they are offered just one cent. Furthermore, only 16% of subjects are indifferent in all of the decision problems.

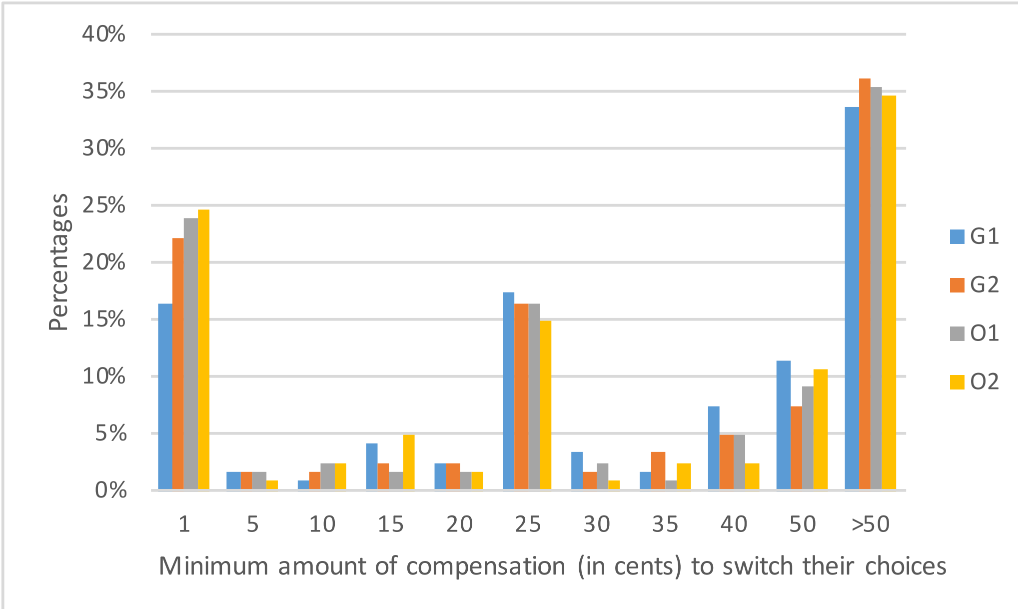

Figure 1: Distribution of the minimum required compensation for switching their choices for each decision problem

As can be seen from Figure 1, approximately 1/3 of the subjects are unwilling to switch away from their preferred mode of information disclosure even if they are offered \$0.50. In other words, these subjects are willing to give up more than 5% of the surplus to obtain their preferred mode of information disclosure.

Result 1: The majority of the subjects have a strict preferred mode of information disclosure regarding the resolution of uncertainty.

Since the standard model predicts indifference between P and Q in all four decision problems, Result 1 reveals that the predictions of the standard model fail for the vast majority of our experimental subjects.

Next, we investigate the strict preference of the subjects. Table 7 shows the aggregate choice percentages in the four decision problems for subjects who exhibit strict preference for resolution of uncertainty.

As it can be seen in Table 7, subjects' strict preference is toward early resolution of uncertainty when all boxes are opened simultaneously. In the decision problems O1 and O2, 65.59% and 65.22% of the subjects prefer early resolution of uncertainty, respectively. When the boxes

Table 7: Aggregate choice percentages for early resolution of uncertainty for subjects who exhibit strict preference

Problem	% Early	p-value	Conclusion
G1	39.22	0.015	$Q \succ P$
G2	58.95	0.041	$P \succ Q$
O1	65.59	0.001	$P \succ Q$
O2	65.22	0.004	$P \succ Q$

The numbers of subjects who exhibit strict preference in G1, G2, O1, and O2 are 102, 95, 93, and 92, respectively. p-values are from one-sample test of proportions.

are opened sequentially, for decision problem G2; that is when two boxes contain a prize, still 58.95% of the subjects prefer early resolution of uncertainty. However, for decision problem G1; that is, when there is only one prize, the preference is reversed: 60.78% of the subjects prefer the late resolution of uncertainty. As an additional test to compare the distribution of choices in each decision problem, we look at the choices of 88 subjects who exhibit strict preference in all of the decision problems. Based on sign tests, there is no significant difference when the boxes are opened simultaneously (O1 vs O2, p=1.000) but there is a significant difference when the boxes are opened sequentially (G1 vs G2, p=0.005). Furthermore, whether opening boxes simultaneous or sequentially matters only when there is a single prize (O1 vs G1, p=0.001; O2 vs G2, p=0.286).

Result 2: While the majority of the subjects strictly prefer early resolution of uncertainty in O1, O2, and G2, the late resolution of uncertainty is strictly preferred in G1.

A categorical preference for early resolution in O1 and O2, where uncertainty always resolves in a single moment, is consistent with the results in the previous literature. Result 2, however, identifies a novel pattern of information demand when gradual information disclosure is involved. Subjects still opt for early resolution in problem G2, where the late option provides gradual information with conclusive good news and partial bad news. However, subjects opt for late over immediate disclosure in problem G1, where the late option provides gradual good news coupled with conclusive bad news.

Heterogeneous types and its implications to theory

The observed difference in behavior between problems G1 and G2 is incompatible with standard theory and with several alternative models of demand for non-instrumental information that have been developed more recently. Section 4 showed that preferences over temporal lotteries in Kreps and Porteus (1978), preferences for one-shot resolution of uncertainty in Dillenberger (2010), and surprise and suspense utility in Ely et al. (2015) all predict identical behavior across G1 and G2, and are incompatible with Result 2.

Section 4 shows that a difference in behavior between G1 and G2 can be understood as arising from subjects' special emphasis on their best (peak) and worst (trough) experiences along their realized path of uncertainty. In particular, the late option Q in G1 offers a better distribution of path peaks than the early option P, while offering the same distribution of path troughs (Table 4). Conversely, in problem G2 the early option P offers a better distribution of path troughs than the late option Q, while offering the same distribution of path peaks (Table 5). Result 2 provides evidence that the emphasis on peaks and troughs captured by the peak-trough utility model (Gul et al., 2021) may be a quantitatively important aspect of decision making behavior.

The averages reported in Table 7 conceal substantial heterogeneity among subjects. The most prevalent types of demand for non-instrumental information fall into four categories[2]:

Thrill seeker: 29% of non-indifferent subjects would never pay for late resolution of uncertainty in O1 and O2, where all uncertainty is resolved at once, but switch to a strict preference for late resolution in problem G1 (and, in some cases, in both G1 and G2) when the late option provides gradual information. This behavior is consistent with peak-trough utility agents with a convex v and sufficiently high weight θ_h given to path peaks. Among them, those with a sufficiently low θ_ℓ will switch to strict preference for late resolution in both G1 and G2. For example, let $v(\alpha) = \alpha^\gamma$ for $\gamma > 1$. Then, the agent is a thrill seeker with strict preference for

[2]The distribution of all decisions is in Appendix B.

late resolution in G1 and G2 if and only if

$$\begin{aligned} \theta_h &> 3^{\gamma-1}\frac{2^\gamma - 2}{3^\gamma - 2^\gamma} \\ \theta_\ell &< 3^{\gamma-1}\frac{2 - 2^\gamma}{4^\gamma - 3^\gamma} \end{aligned}$$

The parameters $\gamma = 1.1, \theta_h = .2, \theta_\ell = -.2$, for instance, satisfy these conditions.

Information seeker: 22% of non-indifferent subjects prefer early resolution of uncertainty in each of the four decision problems. Peak-trough agents are information seeking types if their v is convex, θ_h is not too large (or negative) and θ_ℓ is not too negative (or positive).

Information avoider: 13% of non-indifferent subjects prefer late resolution of uncertainty in each of the four decision problems. Peak-trough agents are information avoiding types if their v is strictly concave, $\theta_h \geq 0$ and $\theta_\ell \leq 0$.

Thrill avoider: 8% of non-indifferent subjects prefer early resolution of uncertainty when uncertainty is resolved gradually, i.e. in problems G1 and G2, but they strictly prefer late resolution in at least one of the decision problems where all uncertainty is resolved at once, i.e. in problem O1, problem O2, or both. For example, an individual who prefers early resolution in G1 and G2 but prefers late resolution in O1 and O2 is compatible with peak-trough utility when v is concave, θ_ℓ is sufficiently high and θ_h is sufficiently negative.

6 Conclusion

Our simple experimental design offers an ideal setting for calibrating and testing models of demand for information. Our experiment offers strong evidence that subjects prefer gradual resolution of uncertainty over early or late resolution, when information about winning yields decisive bad news but inconclusive good news. Our main findings are difficult to reconcile with standard theory and with many of the existing models of demand for non-instrumental information, including the Kreps-Porteus model. We showed that these findings can be understood as arising from the decision maker's emphasis on their best (peak) and worst (trough) experiences along the realized path of uncertainty.

References

Ahlbrecht, Martin and Martin Weber, "An empirical study on intertemporal decision making under risk," *Management Science*, 1997, *43* (6), 813–826.

Caplin, Andrew and John Leahy, "Psychological expected utility theory and anticipatory feelings," *The Quarterly Journal of Economics*, 2001, *116* (1), 55–79.

Dillenberger, David, "Preferences for one-shot resolution of uncertainty and Allais-type behavior," *Econometrica*, 2010, *78* (6), 1973–2004.

Eliaz, Kfir and Andrew Schotter, "Paying for confidence: An experimental study of the demand for non-instrumental information," *Games and Economic Behavior*, 2010, *70* (2), 304–324.

Ely, Jeffrey, Alexander Frankel, and Emir Kamenica, "Suspense and surprise," *Journal of Political Economy*, 2015, *123* (1), 215–260.

Epstein, Larry G and Yoram Halevy, "Ambiguous correlation," *Review of Economic Studies*, 01 2018, *86* (2), 668–693.

Falk, Armin and Florian Zimmermann, "Beliefs and utility: Experimental evidence on preferences for information," *Working paper*, 2016.

Fischbacher, Urs, "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, 2007, *10* (2), 171–178.

Fredrickson, Barbara L and Daniel Kahneman, "Duration neglect in retrospective evaluations of affective episodes.," *Journal of Personality and Social Psychology*, 1993, *65* (1), 45.

Ganguly, Ananda and Joshua Tasoff, "Fantasy and Dread: The Demand for Information and the Consumption Utility of the Future," *Management Science*, 2017, *63* (12), 4037–4060.

Golman, Russell, David Hagmann, and George Loewenstein, "Information Avoidance," *Journal of Economic Literature*, 2017, *55* (1), 96–135.

Greiner, Ben, "Subject pool recruitment procedures: organizing experiments with ORSEE," *Journal of the Economic Science Association*, 2015, *1*, 114–125.

Gul, Faruk, Paulo Natenzon, and Wolfgang Pesendorfer, "Random Evolving Lotteries and Intrinsic Preference for Information," *Econometrica*, 2021, *89* (5), 2225–2259.

Kőszegi, Botond and Matthew Rabin, "Reference-dependent consumption plans," *American Economic Review*, 2009, *99* (3), 909–36.

Kreps, David M and Evan L Porteus, "Temporal resolution of uncertainty and dynamic choice theory," *Econometrica*, 1978, pp. 185–200.

Masatlioglu, Yusufcan, A Yesim Orhun, and Collin Raymond, "Intrinsic information preferences and skewness," *Ross School of Business Paper*, 2017.

Nielsen, Kirby, "Preferences for the resolution of uncertainty and the timing of information," *Journal of Economic Theory*, 2020, *189*, 105090.

Oster, Emily, *Expecting Better: Why the Conventional Pregnancy Wisdom is Wrong–and what You Really Need to Know*, Penguin, 2014.

_ , Ira Shoulson, and E Dorsey, "Optimal expectations and limited medical testing: evidence from Huntington disease," *American Economic Review*, 2013, *103* (2), 804–30.

Palacios-Huerta, Ignacio, "The aversion to the sequential resolution of uncertainty," *Journal of Risk and Uncertainty*, 1999, *18* (3), 249–269.

A Instructions

Welcome and thank you for coming today to participate in this experiment. This is an experiment in decision making. Your earnings will depend on your own decisions and chance. It will not depend on the decisions of the other participants in the experiment. Please pay careful attention to the instructions as a considerable amount of money is at stake. The entire experiment is expected to finish within 30 minutes. At the end of the experiment you will be paid privately. At this time, you will receive $7 as a participation fee (simply for showing up on time).

In this experiment, you will participate in four independent decision questions that share a common form. At the end of the the experiment, the computer will randomly select one decision question. The question selected depends solely upon chance, and each one is equally likely. The question selected, your choice and your payment in that question will be shown. Your final earnings in the experiment will be your earnings in the selected question plus $7 show-up fee.

During the experiment it is important that you do not talk to any other subjects. Please turn off your cell phones. If you have a question, please raise your hand, and the experimenter will come by to answer your question. Failure to comply with these instructions means that you will be asked to leave the experiment and all your earnings will be forfeited.

A.1 Decision Questions

A.1.1 G1

In the next screen, you will be shown three identical looking boxes. One of the boxes contain a prize of $10, the other two boxes do not contain any prize. Your task is to select one of the boxes by clicking on the box of your choice. If the box you selected contains a prize, you will earn $10 in this decision question. If the box you selected does not contain a prize, you will not earn or lose any amount in this decision question.

In the screen after, you will make a selection to learn the content of the boxes. The boxes will be opened sequentially. First, one of the boxes will be opened. 60 seconds later, another box will be opened. The last box will be opened 60 seconds after the second box is opened. You will see the time counter in the upper-right corner of your screen.

A.1.2 G2

In the next screen, you will be shown three identical looking boxes. Two of the boxes contain a prize of $10, the other box does not contain any prize. Your task is to select one of the boxes by clicking on the box of your choice. If the box you selected contains a prize, you will earn $10 in this decision question. If the box you selected does not contain a prize, you will not earn or lose any amount in this decision question.

In the screen after, you will make a selection to learn the content of the boxes. The boxes will be opened sequentially. First, one of the boxes will be opened. 60 seconds later, another box will be opened. The last box will be opened 60 seconds after the second box is opened. You will see the time counter in the upper-right corner of your screen.

A.1.3 O1

In the next screen, you will be shown three identical looking boxes. One of the boxes contain a prize of $10, the other two boxes do not contain any prize. Your task is to select one of the boxes by clicking on the box of your choice. If the box you selected contains a prize, you will earn $10 in this decision question. If the box you selected does not contain a prize, you will not earn or lose any amount in this decision question.

In the screen after, you will make a selection to learn the content of the boxes. The boxes will be opened simultaneously. All three boxes can be opened immediately, or all three boxes can be opened after 120 seconds. You will see the time counter in the upper-right corner of your screen.

A.1.4 O2

In the next screen, you will be shown three identical looking boxes. Two of the boxes contain a prize of $10, the other box does not contain any prize. Your task is to select one of the boxes by clicking on the box of your choice. If the box you selected contains a prize, you will earn $10 in this decision question. If the box you selected does not contain a prize, you will not earn or lose any amount in this decision question.

In the screen after, you will make a selection to learn the content of the boxes. The boxes will be opened simultaneously. All three boxes can be opened immediately, or all three boxes can be opened after 120 seconds. You will see the time counter in the upper-right corner of your screen.

Sample screenshots

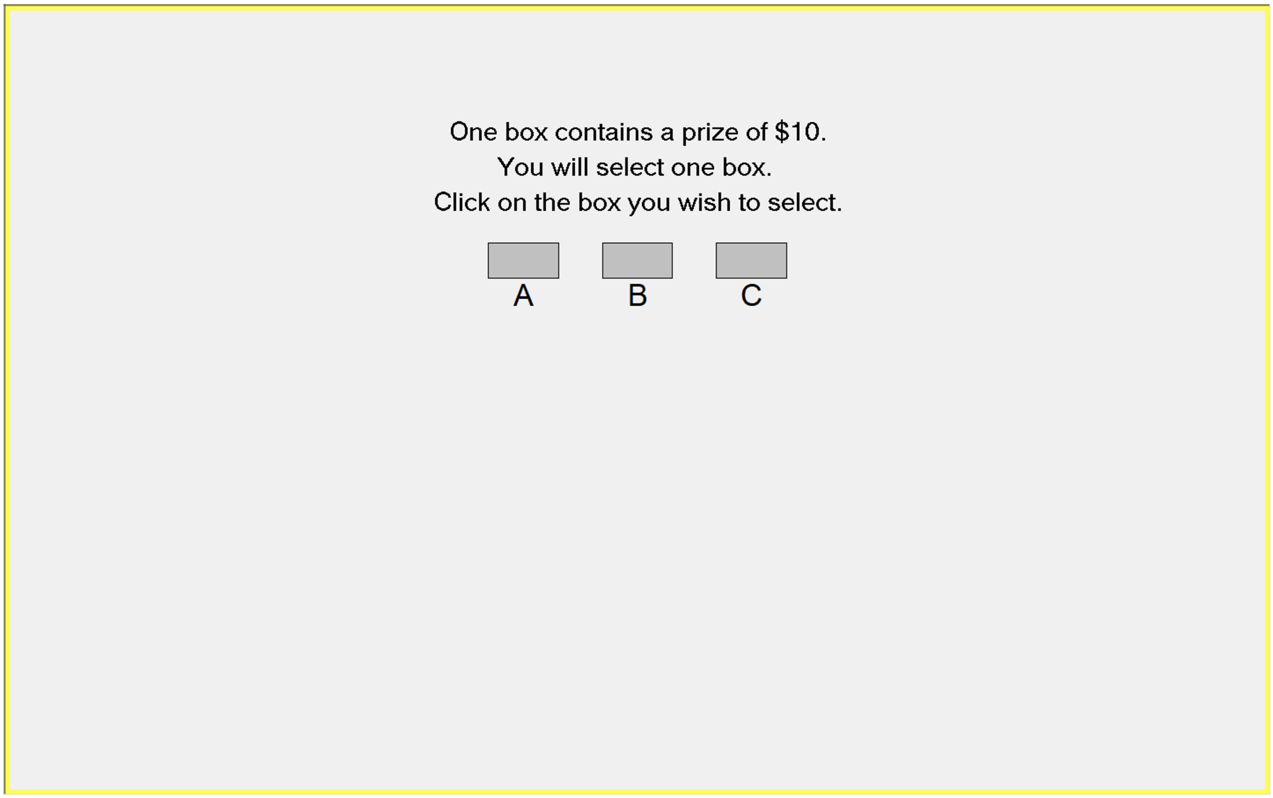


You selected box A.
Now, you will make a selection to learn the content of the boxes. The boxes will be opened sequentially: First, one of the boxes will be opened. 60 seconds later, another box will be opened. The last box will be opened 60 seconds after the second box is opened. You will see the time counter in the upper-right corner of your screen.
When do you want to learn the content of your box?
First
Last
OK

A.2 Willingness to Switch Elicitation

For the decision question where there {is one prize/ are two prizes} and the boxes will be opened {sequentially / simultaneously}, you choose your box to be opened {*subject's response*}. Now, you will see 10 questions, each of which will ask you whether you would change your choice from opening your box {*subject's response*} to opening your box {*unchosen response*} if we compensated you for the amount specified in that question. You will answer by selecting Yes or No.

If this decision question is randomly selected to be played, then one of the 10 questions for this decision question will be randomly selected by the computer. Each question is equally likely, and your choice in the selected question will determine whether your box will be opened first or last. If you select Yes, you will receive the monetary compensation specified in that question but you will change your choice, so your box will be opened {*unchosen response*}. If you select No, you will keep your choice, so your box will be opened {*subject's response*}.

The more you want the option you chose (opening your box {*subject's response*}) over the option you rejected (opening your box {*unchosen response*}), the higher compensation you should require to give up your choice and switch to the option you did not want. Think about what compensation is too little for you to switch your choice, and what compensation would be enough. Accordingly, click Yes or No for each question.

B Distribution of decisions in each question

G1	G2	O1	O2	Percentage
Early	Early	Early	Early	19%
Indiff	Indiff	Indiff	Indiff	16%
Late	Early	Early	Early	12%
Late	Late	Late	Late	11%
Late	Late	Early	Early	7%
Early	Late	Early	Late	4%
Early	Early	Late	Late	3%
Late	Early	Late	Early	3%
Early	Early	Late	Early	2%
Late	Indif.	Indif.	Indif.	2%
Late	Late	Indif.	Indif.	2%
Late	Late	Early	Late	2%
Late	Late	Late	Early	2%
Late	Early	Early	Late	2%
Late	Early	Late	Late	2%
Early	Indif.	Early	Indif.	1%
Early	Indif.	Early	Late	1%
Early	Early	Late	Indif.	1%
Early	Late	Early	Early	1%
Early	Late	Late	Early	1%
Late	Indif.	Indif.	Early	1%
Late	Indif.	Late	Indif.	1%
Late	Early	Indif.	Late	1%
Late	Early	Early	Indif.	1%
Late	Late	Indif.	Late	1%