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: October 19, 2020. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

†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

In this paper, we report on an experiment that shows subjects prefer gradual resolution of uncertainty when information about winning yields decisive bad news but inconclusive good news. This behavior is difficult to reconcile with standard theories of choice and even with many of the alternatives that have been developed more recently. We assess the usefulness of a recently proposed model, peak-trough utility (Gul, Natenzon and Pesendorfer, 2019), for analyzing the observed patterns of behavior.

Consider the following scenario: a decision maker (DM) places a bet that pays off if the event W occurs. At some pre-specified time T, the state of the world is revealed and the DM receives the prize if the state is in W. Prior to time T, the DM receives a flow of information about the likelihood of W. For example, she may learn immediately whether or not W has occurred; alternatively, the DM may receive no information before time T. Since this information is decision irrelevant, a standard DM whose preferences depend only on the probability distribution over outcomes would be indifferent between learning immediately and not learning at all or any other disclosure rule.

We provide evidence that subjects prefer to learn gradually if information is of the *decisive bad news* or *inconclusive good news* variety. Thus, subjects prefer a process in which the chance of winning goes up gradually or drops to zero all at once, to a process that resolves all uncertainty in a single moment. By contrast, we find that subjects are less inclined to choose gradual information disclosure when information is of the *decisive good news* or *inconclusive bad news* variety.

It is well known that many decision makers value noninstrumental information; they often prefer early over late disclosure (e.g. Ahlbrecht and Weber, 1997; Falk and Zimmermann, 2016; Eliaz and Schotter, 2010); though in some situations they have a clear preference for late resolution (e.g. Oster, Shoulson and Dorsey, 2013). Golman, Hagmann and Loewenstein, 2017 provide an extensive survey of the experimental literature on preferences for timing of resolution of uncertainty.

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 information. Nevertheless, Neilsen concludes that "gradual resolution generally is preferred to one-shot resolution in both treatments."

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 DM's attitude toward noninstrumental information. The Ely et al. model is consistent with a preference for gradual resolution of uncertainty. However, it cannot accommodate the strict preference for gradual good news over gradual bad news that plays an important role in our experiment. We discuss these two papers in detail in section 6 below.

Caplin and Leahy (2001) present the first model of anticipatory utility. Their model is suitable for studying attitudes towards noninstrumental information. They do not, however, address the issue. 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 noninstrumental information. They identify conditions guaranteeing a preference for one-shot resolution of uncertainty and conditions that yield a preference for early resolution.

2 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 DM receives, in period N, a prize from a finite set, A, of prizes. In each period $t = 1, \dots, N$, 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, $(\alpha_1, \dots, \alpha_N)$, an *evolving lottery*. Below, we use the terms "evolving lottery" and "path" interchangeably.

The carrier of utility is a *random evolving lottery*; that is, a probability distribution, P, over paths. 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 $\alpha = (\alpha_1, \dots, \alpha_N)$ with α_1 and assign to α the expected utility of α_1. Hence, a standard DM with von Neumann-Morgenstern utility function u would assign utility

$$U(P) = \sum_{\alpha} u(\alpha_1)P(\alpha) \tag{1}$$

to every random evolving lottery P. Note that random evolving lotteries must satisfy the martingale property: if P assigns non-zero probability to a sequence $(\alpha_1, \dots, \alpha_t)$ up to a time $t < N$, then the expected value of α_{t+1} at time t, given $(\alpha_1, \dots, \alpha_t)$ must be equal to α_t. Therefore, the law of iterated expectation and the linearity of u imply,

$$\sum_{\alpha_t} u(\alpha_t)P(\alpha) = \sum_{\alpha_s} u(\alpha_s)P(\alpha).$$

for all P and all t, s.

Unlike the standard agent of equation (1), an expected utility maximizing DM who cares about how uncertainty resolves would have a von Neumann-Morgenstern utility index w that assigns a value to each path. Such a DM's utility function over random evolving lotteries would take the following form:

$$W(P) = \sum_{\alpha} w(\alpha)P(\alpha). \tag{2}$$

Gul et al. (2019) study a subset of the utility functions in (2); this subset includes the standard model described in equation (1). Ely et al. (2015) study a set of nonexpected utility functions over random evolving lotteries. Both of these papers are motivated by the growing literature showing that decision makers have preferences over how uncertainty resolves. The experimental evidence we present below identifies patterns of gradual resolution that decision makers like and for which they are willing to pay.

3 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. 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 different 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 the 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.

After making their choices in the four decision problems, subjects completed four choice lists. These lists offered various levels of compensation to the subjects for switching their initial choices. Ten different amounts of compensation ranging from 1 cent to 50 cents were offered. This is identical to the 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.

Each decision problem above presents the subject with two options. Both options yield the same probability of winning \$10. However, they lead to different random evolving lotteries. Let P denote the random evolving lottery associated with the early option 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:

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

The standard model described in (1) predicts indifference between P and Q in all four decision problems. The results below reveal that the predictions of the standard model fail for the vast majority of our experimental subjects.

4 Experimental Results

Table 1, below, shows the aggregate choice frequencies in the four decision problems of the 125 subjects in the experiment:

For decision problem G1; that is, when only one box contains a prize and boxes are opened sequentially, Table 1 shows that —in the aggregate— nearly 60% of subjects prefers late resolution of uncertainty. This preference is reversed in decision problem O1 when all boxes are opened simultaneously. In that case, only 36% of subjects prefer late resolution and 64% prefer early resolution. When two boxes contain a prize, this asymmetry disappears. In decision

Table 1: Aggregate choice frequencies

Problem	% late	p-value	Conclusion
G1	59.2	0.024	Q $\succ$ P
G2	40.0	0.016	P $\succ$ Q
O1	36.0	0.006	P $\succ$ Q
O2	38.3	0.001	P $\succ$ Q

problem G2, 40% prefer late resolution of uncertainty when the boxes are opened sequentially while 38% prefer late resolution in the one-shot case O2.

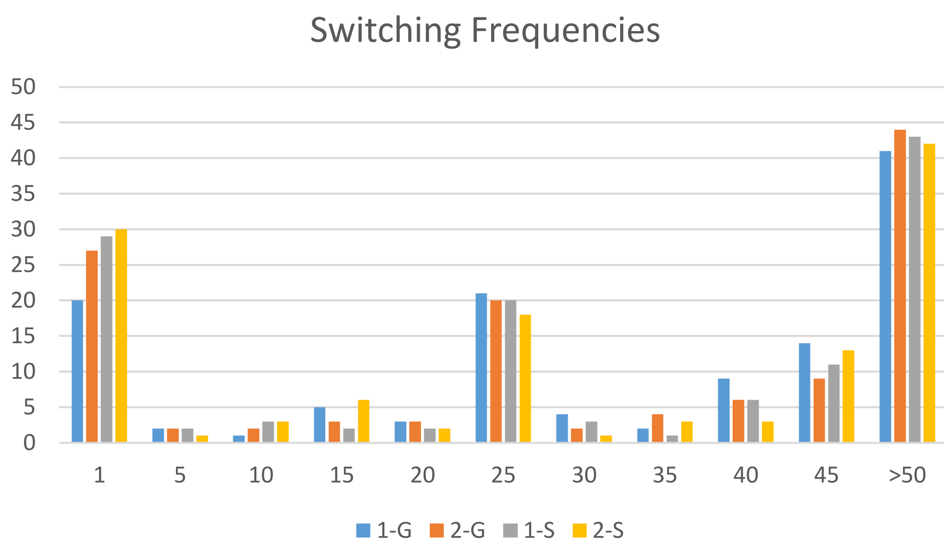

Figure 1

In the second part of the experiment, we elicit subjects' strength of preference. For each decision problem, we ask subjects whether they would change their choice if we paid them X to do so, where X varies from 1 cent to 50 cents. We exclude 3 of the 125 subjects because their responses to the strength of preference question were not monotone. Figure 1, above, shows the frequency distribution of the required compensation for switching away from the preferred mode of information disclosure.

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. As can be seen from

Table 2: Distribution of choices in each decision problem

Problem	% Early	%Indiff.	% Late
G1	32.79	16.39	50.82
G2	45.90	22.13	31.97
O1	50.00	23.77	26.23
O2	49.18	24.59	26.22

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. Hence, these subjects are willing to give up more than 5% of the surplus to obtain their preferred mode of information disclosure.

Table 2, below, summarizes the distribution of choices once the indifferent subjects are excluded. The results are qualitatively similar to those of Table 1: the fraction of subjects choosing late resolution is again significantly higher in problem G1 than in the three other choice problems.

The averages reported in Tables 1 and 2 conceal substantial heterogeneity among subjects. Table 3 presents the five most prevalent types in the subject sample.[1] The most common type, encompassing around 19% of subjects, is the *always early* type, who prefers early resolution of uncertainty in each of the four choice problems. The second largest group of subjects (16%) are *indifferent* types who do not care how uncertainty is revealed. *Hopeful* types (12%) are subjects who prefer early resolution of uncertainty except in problem G1. Next come *always late* subjects which comprise around 11% of the sample. Finally, around 7% of subjects are *thrill seekers*, who prefer the late option when uncertainty resolves gradually and the early option when all uncertainty is resolved at once.

Note that only a minority of subjects, the indifferent types, are compatible with the standard model in (1). Below, we assess the ability of recent theories to model the rich variety of attitudes towards resolution demonstrated above. First, we show how the patterns identified above can be understood through the lens of the Peak-Trough utility model.

[1]The distribution of all decisions are in Appendix B.

Table 3: The five most prevalent decision maker types

	G1	G2	O1	O2	Percentage
Always Early	Early	Early	Early	Early	19%
Indifferent	Indiff	Indiff	Indiff	Indiff	16%
Hopeful	Late	Early	Early	Early	12%
Always Late	Late	Late	Late	Late	11%
Thrill-Seeker	Late	Late	Early	Early	7%

5 Peaks and Troughs

Fredrickson and Kahneman (1993) argue that, in retrospective evaluations, subjects neglect the duration of experiences and emphasize extremes. Hence, a useful way to understand the evidence presented in the last section is to examine the role of a subject's best (peak) and worst (trough) experiences as the uncertainty resolves. Peak-trough utility theory (Gul et al., 2019) offers a simple theoretical framework in which agents place special emphasis on their best and worst experiences. Consider, again, the general class of utility function over random evolving lotteries of (2): 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:

$$W(P) = \sum_{\alpha} w(\alpha)P(\alpha) \tag{2}$$

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 δ_h, δ_ℓ be weights such that

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

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

$$
w(\alpha) = \frac{1-\delta_h-\delta_\ell}{N}\sum_{t=1}^{N} v(u(\alpha_t)) + \delta_h v(\overline{u}(\alpha_1,\ldots,\alpha_N)) + \delta_\ell v(\underline{u}(\alpha_1,\ldots,\alpha_N)) \tag{3}
$$

The function W defined in (2) is a peak-trough utility whenever w is as described in (3). Compared to the standard model of equation (1), 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 δ_h, δ_ℓ which determine the agent's sensitivity to the best and worst experiences. If we assume v is linear and set the weights $\delta_h = \delta_\ell = 0$, then we are back to the standard model.[2]

We can calibrate the parameters of peak-trough utility to our experimental results. 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 δ_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, 1)	1	1/3	1/3	1/3
(1/3,1/2,1)	1	1/3	–	1/3
(1/3,1/2,0)	2/3	0	–	1/3

The following proposition relates the optimal choice in decision problem G1 to the param-

[2]For simplicity, we assumed that each period has the weight $1/N$. A more general model would permit discounting. Gul et al. (2019) provide an axiomatic foundation for this, more general, model.

eters v and δ_h of peak-trough utility.

Proposition 1. *Let W be a peak-trough utility with parameters $(v, \delta_h, \delta_\ell)$, 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 $\delta_h > 0$.*

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

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

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 (3) and using the fact that $v(0) = 0$ and $v(1) = 1$ we obtain

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

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

Subjects with $\delta_h < 0$ dislike getting their hopes up. Therefore, they tend to prefer the early option in G1. Subjects with $\delta_h > 0$ 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 latter being twice as likely than the former. 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, δ_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 δ_ℓ of peak-trough utility to the choice in decision problem G2.

Proposition 2. *Let W be a peak-trough utility with parameters $(v, \delta_h, \delta_\ell)$, 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 $\delta_\ell < 0$.*

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

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

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

$$W(Q) - W(P) = \frac{1 - \delta_h - \delta_\ell}{9}[2v(1/2) - 1] + \frac{\delta_\ell}{3}[v(1/2) - v(2/3)] \tag{4}$$

is positive. If v is linear, then (4) simplifies to $-\frac{\delta_\ell}{18}$ and, therefore, part (i) follows. If v is convex, then $2v(1/2) \leq 1$ and, therefore, $W(Q) - W(P) > 0$ implies $\delta_\ell < 0$ which proves part (ii). If v is concave, then $2v(1/2) \geq 1$ and, therefore, $W(P) - W(Q) > 0$ implies $\delta_\ell > 0$ which proves part (iii). □

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

In decision problems 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 δ_ℓ and δ_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 3. *Let W be a peak-trough utility with parameters $(v, \delta_h, \delta_\ell)$ 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,

$$W(P) - W(Q) = \frac{1 - \delta_h - \delta_\ell}{3} [c - v(c)] \tag{5}$$

Since $v(c) \geq c$ if v is concave and $v(c) \leq c$ if v is convex, the result follows from (5). □

Propositions 1-3 relate peak-trough utility parameters to the behaviors described in Table 3. Peak-trough agents are always early types (19% of subjects) if their v is convex, δ_h is not too large (or negative) and δ_ℓ is not too negative (or positive). Always indifferent types (16% of subjects) are consistent with standard theory; that is, with a linear v and $\delta_h = \delta_\ell = 0$. Peak-trough utility agents with a convex v and δ_h sufficiently high are hopeful (12%). Peak-trough utility agents with a concave v, δ_h not too negative (or positive) and δ_ℓ not too positive (or negative) are always late types (11%). Finally, a peak-trough utility agent with a convex v, $\delta_h > 0$, and $\delta_\ell < 0$ is a thrill seeker (7%).

To see how the peak-trough utility parameters interact in the propositions above, consider the following example: let $v(\alpha) = \alpha^\gamma$ for $\gamma > 1$. Then, the agent is a thrill seeker if and only if

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

The parameters $\gamma = 1.1, \delta_h = .2, \delta_\ell = -.2$, for example, satisfy these conditions.

6 Discussion

Our experiment offers strong evidence that subjects prefer gradual resolution of uncertainty over early or late resolution. Peak-trough utility provides a rationale for this behavior but it is not the only theory that predicts a preference for gradual resolution. Ely, Frankel and Kamenica (2015) offer *suspense* and *surprise* as an alternative rationale for a preference for gradual resolution. Suspense is the expected belief variation next period while surprise is the realized belief variation from period to period.

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 x in the former with $1 - x$. An immediate implication of this symmetry is that suspense and surprise utility cannot distinguish between decision problems G1 and G2. Thus, suspense and surprise utility cannot accommodate agents who choose gradual resolution in G1 *and* early resolution in G2 (or the reverse).

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 in decision problems G1 and G2. A theory that focuses on belief variation cannot distinguish between "good news" and "bad news," a distinction that is at the heart of our experimental results.

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

Table 7: Kreps-Porteus utilities

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

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 Kreps-Porteus model cannot capture the differences in behavior between decision problems G1 and G2.

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. In our experiment, few subjects exhibit this pattern of behavior; the combination early in G1 and late in G2, on the other hand, is common yet inconsistent with the Kreps-Porteus model.

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.

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.

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

Gul, Faruk, Paulo Natenzon, and Wolfgang Pesendorfer, "Random Evolving Lotteries and Intrinsic Preference for Information," *working paper*, 2019.

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, 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

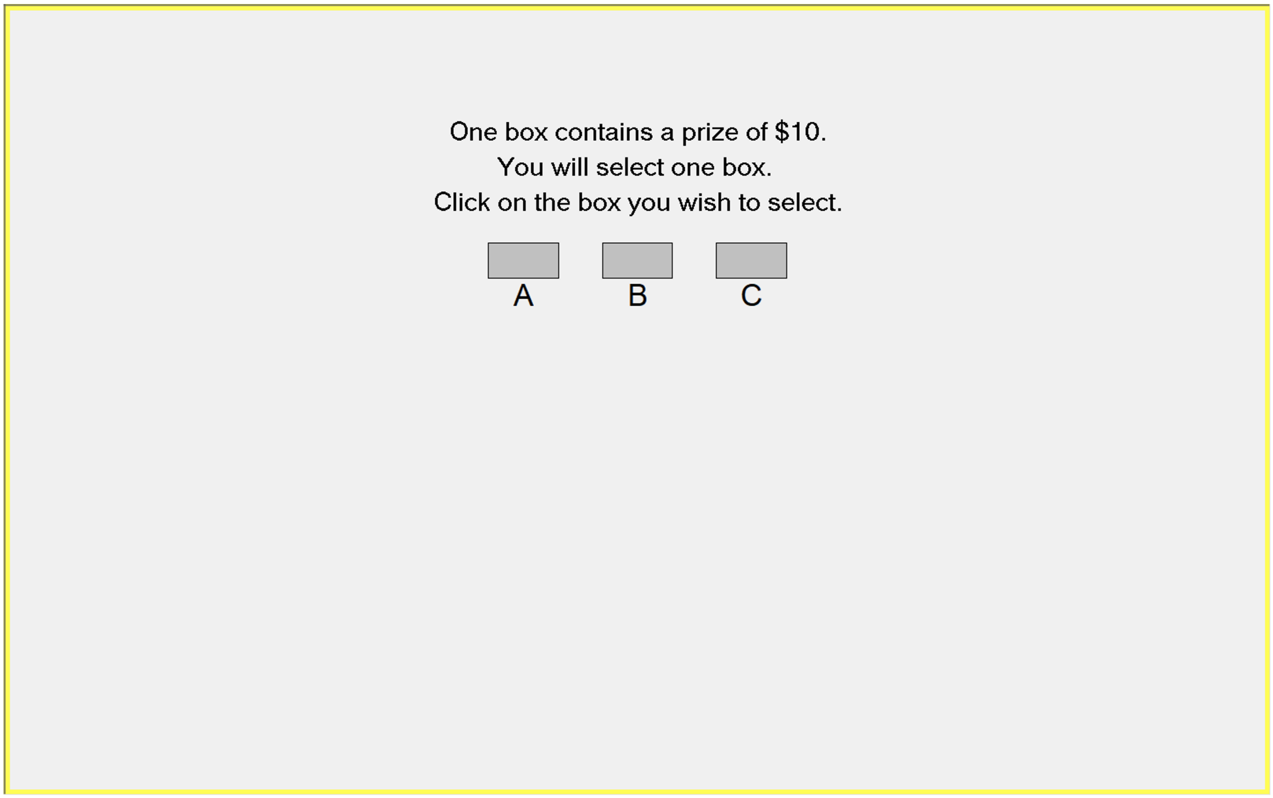


You selected box A.
Now, you will make a selection to learn the content of the boxes. The boxes will be opened sequentially: First, one of the boxes will be opened. 60 seconds later, another box will be opened. The last box will be opened 60 seconds after the second box is opened. You will see the time 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%