The Resolution of Uncertainty in the Value and Probability Domains*

Eungik Lee
NYU

Kathleen Ngangoue
UCLA Anderson

Andrew Schotter
NYU

March 2023

Abstract

We compare preferences for resolution of uncertainty when the uncertainty is resolved over a probability rather than a value. In various existing frameworks–e.g., Kreps and Porteus (1978)–, preferences over gradual versus one-shot resolution do not depend on whether values or probabilities define the main object of uncertainty. Yet, in our experiment, a large majority of subjects preferred to resolve uncertain values gradually but uncertain probabilities all at once–both with uncertainty defined over gains and losses. We investigate the possible determinants of this discrepancy and propose an explanation for it using what we call "process utility". Finally we investigate this idea on a set of experiments, which confirm our theoretical expectations.

Keywords: resolution of uncertainty, probability, gradual resolution, one-shot resolution, process utility, non-instrumental information, Kreps-Porteus.

*The paper has benefited from comments from Jay Lu, Erkut Ozbay, Ariel Rubinstein, Emmanuel Vespa, Georg Weizsaecker, and the seminar audiences at Berlin Behavioral Economics seminar, NYU CESS internal seminar, Briq Belief 22, HBS, UCSD, UCLA, LA Theory Day, UCLA PhD seminar, University of Zurich, AES '22, SITE Experimental economics '22, PhD EI 22, EWMES '22, the MiddExLab Virtual seminar. We thank Xuemeng Yin for excellent research assistance. The authors declare that they have no relevant or material financial interests that relate to the research conducted in this paper. Financial support is acknowledged by the Center for Experimental Social Science at New York University. The data collection was approved by the IRB of New York University.

1 Introduction

Consider the following two lotteries over three equally likely prizes: Lottery A, a value lottery, offers a decision maker a chance of receiving either \$40, \$60, or \$80 each with a 1/3 probability, while Lottery B, a probabilistic lottery, offers the DM either a 40%, 60%, or 80% chance of receiving \$100 (and otherwise nothing). As should be obvious, the lotteries differ in that while the first one is a lottery defined over value prizes, Lottery B is one defined over the probability of winning a fixed prize of \$0 or \$100.

There are two ways these lotteries can be resolved. We can resolve the uncertainty all at once by simply performing the lotteries as described above, or we can resolve them gradually by first randomly removing one of the three prizes leaving two prizes (each equally likely) and then removing a second prize leaving a final prize. Note that the information gained by the gradual resolution of uncertainty is intrinsic and not instrumental in that it does not change the ultimate outcome but simply informs the DM of the way uncertainty is resolved.

The main question we ask in this paper is whether subjects have different preferences over the temporal resolution of uncertainty across these two domains. In other words, do subjects choose to resolve value and probabilistic uncertainty in the same way, or is the object of uncertainty (values or probabilities) an important determinant of these preferences?

Models dealing with the temporal resolution of uncertainty do not predict different choices across our two domains. In other words, if a decision maker (DM) chooses to resolve uncertainty all at once in one domain, she should do so in the other one. For example, Kreps and Porteus (1978, henceforth KP) were first to model preferences for resolution of uncertainty but their model predicts consistent choices across our two domains as do models based on disappointment aversion (Palacios-Huerta, 1999), anticipatory utility (Caplin and Leahy, 2001), suspense and surprise (Ely et al., 2015), information avoidance (Golman et al., 2017). Of course some differences may arise from the way the DM values the binary prospects that exist in the probability domain relative to deterministic payoffs that exist in the value domain, thereby affecting the strength but not the direction of these preferences (Koszegi and Rabin, 2009; Dillenberger, 2010; Strzalecki, 2013; Gul et al., 2021).

Although many theories predict no difference in choices across these two domains, it is not clear whether empirically subjects treat them equally. Our main contribution is to investigate this possibility in a simple experiment. We present an experiment to compare preferences for temporal resolution when uncertainty is in values versus probabilities. Our subjects face lotteries similar to the ones presented above where a lottery is presented with a set of cards and each card has a different prize written on it. In one scenario, the prizes correspond to different money values, creating thereby uncertainty about values, while in the other the numbers represent probabilities, giving rise to uncertainty about the distribution of an underlying prospect. Subjects are asked if they want to resolve uncertainty either all at once or gradually. If they choose all at once, one prize card is selected randomly and immediately, and that determines their payoff in a future period. In the value setting, the card selected is their payoff while in the probability setting the card selected is the probability with which they receive the high payoff. In contrast, if the subjects choose to resolve gradually, then cards are sequentially removed one by one until one card remains. Our main finding is that decisions are only weakly correlated across the two problems (correlation coefficient of 0.17). Our subjects were significantly more likely to resolve uncertain values gradually than uncertain probabilities. This poor correlation cannot be accommodated by many deterministic choice models, including KP.

We investigate the robustness of our results along two dimensions. As preferences for one-shot resolution are often motivated with negative consumption events (e.g., disappointment aversion in Palacios-Huerta (1999), loss aversion in Koszegi and Rabin (2009)), we study preferences as we shift from the gain to the loss domain. Here we find that while preferences for one-shot resolution substantially increase, framing lotteries as losses does not alter the main conclusion relative to the gains treatments. Second, we embed the choices in a more dynamic setting where uncertainty is resolved over time and incoming news may potentially impact preferences. We find this to be the case as subjects altered their preference for gradual resolution to one-shot resolution once they learned that the best outcome (card) was no longer available, making the object of uncertainty less relevant.

Given our results the obvious question to ask is why subjects care more about the resolution of values than probabilities. To explain our results we propose a new approach to the temporal resolution of uncertainty which we call process utility. Its main idea is that subjects value the random process with which their payoff will be determined. As subjects contemplate how to resolve uncertainty, they contrast the immediate payoff lottery with the possible future lotteries that they will face if they choose to resolve gradually. If they find future lotteries substantially more attractive than the current one, they will prefer to resolve gradually. A final experiment confirms our hypothesis.

The paper is organized as follows. In Section 2 we discuss the extant literature on the temporal resolution of uncertainty. In Section 3 we provide some background for the type of decision problems we deal with in our experiment and offer some predictions for our experiment based on the model of KP. We start by describing the three-prize (three-card) problem discussed above where the DM needs to only decide whether to resolve uncertainty immediately or sequentially. In Section 4 we present our results, while

Section 5 investigates the idea of Process Utility experimentally. Section 6 presents our conclusions.

2 Related Literature

To the best of our knowledge, existing experiments model the uncertainty in utility payoffs, mostly in form of money. While a larger literature has provided some evidence on preferences for early versus late resolution (Chew and Ho, 1994; Arai, 1997; Ahlbrecht and Weber, 1997; Kocher et al., 2014), our experiments shed light on the understudied preferences for gradual versus one-shot resolution (Zimmermann, 2015; Falk and Zimmermann, 2017; Masatlioglu et al., 2017; Nielsen, 2020).[1]

One of the first experimental studies testing subjects' aversion to gradual resolution was Zimmermann (2015). Zimmermann (2015) finds little support for an aversion toward gradual resolution. With the exception of Falk and Zimmermann (2017) who find that subjects prefer one-shot resolution with negative consumption events (in their case the occurrence of an electro-shock), subsequent experiments corroborate the finding that with positive consumption events subjects tend to prefer gradual over one-shot resolution. Masatlioglu et al. (2017) investigate primarily subjects' preferences for different information structures, but overall in 75% of the cases their subjects preferred to receive information gradually. Relatedly, Gul et al. (2021) find that preferences for gradual resolution interact with the information structure. In their experiment, subjects liked gradual good news (and decisive bad news), but disliked gradual bad news. The closest paper to ours in spirit is that of Nielsen (2020). Her experiments expose an interaction between the preference for the resolution of uncertainty and the framing of uncertainty. She finds that individuals prefer to delay uncertainty resolution when the choice is framed as a compound lottery where the uncertainty is resolved in real time. In contrast, when uncertainty is framed as information structure in that the outcome has already been determined and is simply being revealed gradually, subjects prefer to learn the outcome earlier. Yet, overall subjects preferred gradual over one-shot resolution. Nielsen's study is dealing more with differences in the source of uncertainty (compound lotteries versus information structures) rather than

[1] The interest in understanding preferences for early versus late resolution was motivated by their potential to distinguish risk attitudes from behavior toward intertemporal substitution in dynamic choice problems (e.g., Weil, 1990). In contrast, preferences for gradual versus full resolution also help explain phenomena related to information acquisition. For instance, agents' intrinsic preferences for information have been used to model investor attention and the ostrich effect (Karlsson et al., 2009; Sicherman et al., 2016; O'Donoghue and Sprenger, 2018) consumption events (Alvarez et al., 2012; Andries and Haddad, 2020), subsequent risk taking (Gneezy and Potters, 1997; Thaler et al., 1997; Anagol and Gamble, 2013), the equity premium puzzle in the presence of myopic loss aversion (Benartzi and Thaler, 1995; Gneezy et al., 2003; Bellemare et al., 2005; Haigh and List, 2005).

the object of uncertainty (values or probabilities). Still, the spirit of our two experiments is similar in that they expose a systematic correlation between these preferences and the underlying uncertainty.

3 The basic decision problem

3.1 The three-card problem

Consider the following stylized decision problem. There are three cards in a deck facing up such that the DM can see what is on each card. Each of them has a different number written on it, referring either to a value or a probability. To determine the DM's payoff, two out the three cards will be randomly removed, and the remaining one will dictate her payoff.

The DM's main task is to choose how to remove two cards. She can either have two cards removed at once and immediately learn her final payoff, or she may choose to have two cards eliminated one-by-one, thereby resolving the uncertainty gradually.

Note that the information revealed by resolving the uncertainty gradually is noninstrumental in that the DM's choices do not affect which cards will be removed. It only affects how she learns which cards have been removed, which is sometimes referred to as *news utility*.

We implement this decision problem in what we call the three-card problem with three different time periods, $t \in \{0, 1, 2\}$. The DM receives her payoff in $t = 2$, but makes her decision on how to resolve uncertainty in $t = 0$. If she decides to resolve the uncertainty in one shot, i.e., have two cards removed simultaneously, then in $t = 1$ she will learn which card will determine the payoff that she will obtain in $t = 2$. If, however, she prefers to have uncertainty resolved gradually, a first card will be removed in $t = 1$, leaving her facing two remaining cards, and then in $t = 2$ a second card will be removed, revealing her payoff-relevant card.

Our main interest is whether subjects' preferences for gradual or one-shot resolution change systematically as we change the meaning of the numbers written on the cards. In one variant of the decision problem, the numbers on the cards correspond to different amounts of money in credits (₵), which was our experimental currency. In $t = 0$, the DM faces three different cards with, for example, ₵40, ₵60 or ₵80. In other words, the DM faces a uniform distribution over three possible monetary prizes. We denote this variant of the decision problem with **M3** (short for *Money with 3 Cards*). In the other variant of the decision problem, the cards depict probability prizes. These probabilities describe the chance of getting ₵100 and otherwise nothing. The equivalent example to the one above would be the situation where the DM faces three different cards with a probability of 40%, 60% or 80% of getting ₵100, or simply put, a uniform distribution over three possible probability prizes of getting ₵100. We denote this variant of the decision problem with **P3** (short for *Probability with 3 Cards*). Our main interest is understanding to what extent the object of uncertainty (values versus probabilities) matters for preferences. Do subjects treat these two decision problems similarly?

In our decision problem one-shot resolution coincides with early resolution. Previous experiments suggest that subjects prefer early over late resolution, and thus, for a more stringent test in the comparison between gradual and one-shot resolution we framed one-shot resolution as early resolution. Our choice also implies that, in our setting, decisions are not only captured by models of preferences for the *form* of resolution (gradual versus one-shot)–i.e., the choice of *how* to resolve uncertainty– but also by models of preferences for the *timing* of resolution (early versus late)–i.e., the choice of *when* to resolve uncertainty.

3.2 Theory and Hypothesis

Different models capture preferences for resolution of uncertainty, but without loss of generality we constrain our attention to one classic model to derive our hypotheses:[2] Kreps and Porteus (1978, henceforth KP) have proposed the first model to capture the difference to the timing of resolution.

We denote $\omega_j \in \Omega = [0, 100]$ the label of card $j \in \{1, 2, 3\}$. In the decision problem M3, subjects receive some money amount ω_j in credits. The problem P3 differs in that ω_j refers to the percentage probability of receiving ₵100. For a better comparison between the problems M3 and P3, we present the predictions as functions of monetary valuations of the cards (i.e., certainty equivalents). We denote $\bar{z}_j$ the monetary value of prize ω_j such that $u(\bar{z}_j) = u(\omega_j)$. In the money domain, $\bar{z}_j$ is simply the monetary prize on card j ($\bar{z}_j = \omega_j$), but in the probability domain where cards denote probability prizes, we assume that the DM values each card with its certainty equivalent such that $\bar{z}_j = u^{-1}(\omega_j u(100))$. Focusing on certain monetary valuations in the model will allow us to abstract from the DM's attitudes toward the risk inherent in the probability prizes, because it is the uncertainty across prizes–not the one underlying the individual prizes–that matters for attitudes toward temporal resolution of uncertainty.

[2]For instance, preferences for one-shot over gradual resolution can be captured with aversion to being disappointed (Palacios-Huerta, 1999), to fluctuations in beliefs (Koszegi and Rabin, 2009) or to compound lotteries (Dillenberger, 2010). Other models modify the utility function to account for emotional reactions to intrinsic information (Caplin and Leahy, 2001; Ely et al., 2015; Loewenstein, 1987). For instance, Caplin and Leahy (2001)'s anticipatory utility model captures the case where agents experience emotions like anxiety or excitement in anticipation of an event.

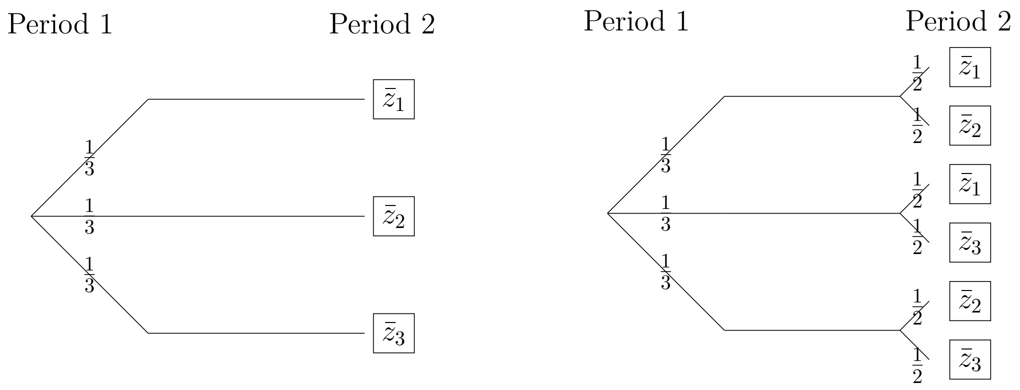

Figure 1: Uncertainty resolution

Following KP, we represent the utility of a temporal lottery $\ell \in \ell$, $U : \ell \to \mathbb{R}$, with:

$$U(\ell) = E[h(E_z u(\bar{z}_j)]$$

The utility function $u(\cdot)$ (with $u(0) = 0, u(100) = 1$) captures attitudes toward risk, while the function $h(\cdot)$ captures separately preferences toward the timing of resolution. The inner expectation operator E_z computes the expected lottery outcome within a period, while the outer expectation operator E computes the expectation over all future contingencies. Next, we illustrate this formula in the context of our decision problem.

The choice between one-shot ("Remove two cards at once") and gradual resolution ("Remove one card first") is tantamount to a choice between two temporal lotteries. The temporal lottery ℓ_O that corresponds to one-shot resolution resolves all uncertainty in period 1 and provides the following utility (see Figure 1(a)):

$$U(\ell_O) = \frac{1}{3}\sum_{j=1}^{3} h(u(\bar{z}_j)).$$

In comparison, the temporal lottery ℓ_G for gradual resolution resolves the uncertainty across periods 1 and 2. In period 1, as one card is removed, subjects face a binary lottery over the two possible outcomes they may get in period 2 (see Figure 1(b)). The utility from resolving uncertainty across two periods is then:

$$U(\ell_G) = \frac{1}{3}\Big[h\Big(\frac{1}{2}(u(\bar{z}_1) + u(\bar{z}_2))\Big) + h\Big(\frac{1}{2}(u(\bar{z}_1) + u(\bar{z}_3))\Big) + h\Big(\frac{1}{2}(u(\bar{z}_2) + u(\bar{z}_3))\Big)\Big]$$

By Jensen's inequality, if the function $h(\cdot)$ is strictly concave, then $U(\ell_G) > U(\ell_O)$. Vice-versa, if $h(\cdot)$ is strictly convex, then $U(\ell_G) < U(\ell_O)$. Intuitively, a concave function

$h(\cdot)$ translates into disliking a spread in outcomes when resolving uncertainty in the early periods. In our design, the only way to reduce this spread in earlier periods is to shift some of the uncertainty resolution to later periods by resolving the uncertainty gradually; Analogously, a convex function $h(\cdot)$ translates into preferring to face a high spread early. Notice that the utility function $u(\cdot)$ does not affect the preference relation and, thus, preferences over temporal lotteries are pinned down by the functional form of $h(\cdot)$. This proposition defines our first hypothesis.[3,4]

Hypothesis 1. *For a given $h(\cdot)$ that is concave or convex over the domain of $u(0)-u(100)$, the agent will make the same choice in M3 and P3.*

One important assumption inherent in Hypothesis 1 is the ability to reduce compound lotteries. In our setting, we need to distinguish between the ability to reduce compound risk 1) when the lotteries are compounded within a single period like a multiple toss of a coin or the lottery ℓ_O in P3 (*intratemporal* recursivity) and 2) when lotteries are compounded across multiple periods like temporal lotteries (*intertemporal* recursivity). Concretely, *intratemporal* recursivity refers to preferences that evaluate the same lotteries in their reduced and compound form equivalently when lotteries are played out within a single period. *Intertemporal* recursivity, on the other hand, implies that the DM is indifferent between the simple lottery describing the full resolution and the multi-stage lotteries inherent in gradual resolution. To capture non-indifference toward temporal lotteries we relax intertemporal recursivity, but assume intratemporal recursivity. We discuss the validity of this assumption in Appendix Section A.2, where we show that, on average, we do not find any significant aversion to compound risk and any correlation between preferences for compound risk and preferences for temporal resolution at the individual level.[5]

[3]Risk aversion will matter in this environment in two ways. First, risk attitudes determine whether subjects prefer the decision problem M3 over P3 depending on how much the certainty equivalents $\bar{z}_j$ of the degenerate lotteries on the M3 cards differ from the ones of binary lotteries on the P3 cards. Second, risk aversion may affect how strong preferences for resolution are. The nonlinearity of the utility function impacts the difference in the valuations $|U(\ell_O) - U(\ell_G)|$, which may result in a DM being closer to indifference in one versus the other domain. Section A.1 in the Online Appendix discusses the strength of these preferences in more details.

[4]If the function $h(\cdot)$ is not strictly concave or convex over $u(0) - u(100)$, the utility function $u(\cdot)$ may matter in that it determines the difference in certainty equivalents between domains. These differences in certainty equivalents may then introduces differences in preferences across our domains if the function $h(\cdot)$ is not strictly concave or convex, preventing clear testable predictions. In our data, we estimate the $h(\cdot)$ functions nonparametrically and find them to be either concave (in the money domain) or convex (in the probability domain) over the entire utility domain.

[5]Segal (1993)and Dillenberger (2010) assume recursivity and time neutrality as replacement for the reduction of compound lottery axiom. In contrast to Segal (1993)and Dillenberger (2010) we do not necessarily assume time neutrality.

In a nutshell, an implication of KP preferences is that subjects' choices between full and gradual resolution should be consistent across M3 and P3. Choices should also be consistent within the framework of different models such as Palacios-Huerta (1999); Caplin and Leahy (2001); Ely et al. (2015); Loewenstein (1987). Of course, differences in the strength of these preferences may stem from the uncertainty referring to different objects, in particular with models that, in P3, allow us to also model uncertainty over final utils ($u(0)$ to $u(100)$) rather than just uncertainty over probabilities (e.g., Koszegi and Rabin, 2009; Dillenberger, 2010; Gul et al., 2021). In Appendix Section A.1 we discuss to what extent our empirical results are consistent with differences in the strength of preferences.

4 The Experiments

Our study consists of three primary experiments that we refer to as the Main, the Loss, and the Process experiments. The Main experiment focuses on our main research question and compares preferences for uncertainty resolution across the money and probability domains. To do this, we study revealed preferences in two different decision problems across both domains: a three-card problem (M3 and P3) in which we elicit choices in a static framework and a four-card problem (M4 and P4) that embeds the decision problem in a more dynamic framework. Concretely, the four-card problem adds one more card to the three-card problem, allowing us to check whether the preferences that we elicit in the three-card problem are robust to resolution history as cards in the four-card problem are gradually removed.

The Loss experiment studies the robustness of the Main experiment's findings as we moved from the gain to the loss domain. Lastly, the Process experiment investigates the procedures that subjects use to determine their resolution choices. This experiment is designed to allow us to test whether these choices are consistent with Process Utility, a concept we offer to explain the resolution choices made by our subjects.

In this section, we describe the design and the corresponding results of each treatment one by one. We first describe the three-card problem (M3 and P3) of the Main experiment, and then discuss whether our findings are robust to history by adding a fourth card and to losses by framing the problem in the loss domain.

The Main experiment had a total of 200 participants. The experiment was computerized with o-Tree (Chen et al., 2016) and was run on Zoom across 13 different sessions. Sessions lasted approximately 75 minutes and subjects earned, on average, \$28.53 including a \$10 show-up fee. The currency used in the experiment was experimental credits (₵) with ₵8 corresponding to \$1.

4.1 The Main Experiment

The experiment was implemented as a within-subject design. That is, the same subjects were presented with both variants of the decision problem, facing both value and probabilistic uncertainty. In addition, subjects faced all decision problems with and without history, i.e., (M3 and P3) and (M4 and P4). Subjects always faced the simpler three-card problem before the corresponding more dynamic four-card problem, but we alternated the order of the money and probability lotteries across sessions. After submitting their resolution choices for all decision problems, we asked our subjects whether they would rather have their experimental payoff determined by a problem with monetary cards or a problem with probabilistic cards (see Online Appendix A.5 for more details). As a last part of the Main experiment, we elicited attitudes toward risk, compound risk, ambiguity, and time preferences (see Online Appendix A.2). At the end of the experiment, subjects learned their payoffs in the three and four-card tasks and answered an unincentivized questionnaire. In the questionnaire, they provided some information on their socio-demographic background and took Frederick's (2005) Cognitive Reflection Test.

Before presenting our results for the four-card experiment, let us pause and explain the three-card experiment and its results. We will then proceed to the four-card experiment as a robustness check.

4.1.1 The Three-card Problem

Design. The three-card task in the Main experiment was run as follows. Subjects saw three different cards on their computer interface, and had to select their preferred method of resolving uncertainty over the three cards.[6] Subjects made 16 different decisions in M3 and P3, each. Each decision problem corresponded to the basic three-card problem described above where the numbers on the cards were varied to present our subjects with a variety of lotteries to (see Table 1). Remember, in M3 subjects received the amount of credits on the last remaining card; in P3, the probability on the last card determined the subject's chance of getting ₵100, in which case the computer would then play out the lottery of getting ₵100 with the corresponding probability. It should be noted that in P3 subjects learned the outcome of the lottery (₵0 or ₵100) at the same time they learned their final probability with which the lottery would be played out. In other words, the resolution of the lottery outcome was not in a separate period. This was a deliberate design choice to discourage subjects to perceive the realization of the probability lottery as a separate resolution stage that may contribute to their preference.

[6] Examples of the interface can be found in the experimental instructions in Online Appendix C.

For each of the problems they were presented with, our subjects had to choose between the following three options:

1. **"Remove one card first"**: the first card is randomly removed leaving two remaining cards. After a pause of 15 seconds, the computer randomly removes the second card, leaving a single card that determines the payoff in the round.[7]

2. **"Remove two cards at once"**: the computer randomly removes two of the cards at once, leaving one card which determines the payoff in the round.

3. **"Do not care"**: one of two options above is randomly implemented.

We used different sets of cards for each of the 16 decisions within domains, but use equivalent sets of cards across M3 and P3. The sequence of sets was randomized at the individual level. Table 1 reports the 16 different sets of cards in M3 and P3. The numbers in the tables are normalized to lie in the interval [0,100] so that any number represents either a monetary value in experimental credits or a probability of receiving a ₵100 prize. Before making payoff-relevant decisions, subjects could experiment with two practice decisions. After the two practice decisions in which they learned the final outcome with their chosen resolution, subjects made their 16 payoff-relevant choices without any feedback to prevent potential effects of outcomes on subsequent choices.

Results. Figure 2(a) shows the percentage of each chosen option. Only a relatively small proportion of choices expressed indifference (4.7% in M3, 7% in P3). In M3, subjects preferred, on average, gradual resolution. More precisely, of the total number of choices made by all subjects over all lotteries, 76.2% of the time subjects chose to resolve the uncertainty gradually, while one-shot resolution was chosen only 19.1% of the time. In P3, on the other hand, choices differed visibly: There was a mixed preference between gradual and one-shot resolution. Subjects chose one-shot (gradual) resolution 48.7% (44.3%) of the time. The proportion of gradual resolution between M3 and P3 after excluding the indifferent choices is significantly different at a 1% level (z-statistics: 7.12; Binomial test clustered at subject level).

[7]The experimental literature has experimented with various time periods, ranging from minutes, days, weeks to (hypothetical) years. How preferences vary with the period length is still poorly understood, but our data suggest that seconds are sufficient to capture non-indifference to temporal resolution of uncertainty. In fact, one may perceive preferences for temporal resolution as a preference for (or aversion toward) the processing of compounded information irrespective of the period length. In our setting, we believed that 15 seconds gave subject sufficient time to process these different pieces of information in isolation. This is corroborated by experiments in neuroscience where a time span of thousands of milliseconds are sufficient to study anticipatory neural processes in the brain (e.g., Huettel et al., 2005; Bruhn et al., 2014).

Table 1: 3 cards Lottery Parameters

#	Card 1	Card 2	Card 3	Mean	S.D.	Skewness
1	40	60	80	60.0	20.0	0
2	20	60	80	53.3	30.6	-0.38
3	20	40	80	46.7	30.6	0.38
4	20	40	60	40.0	20.0	0
5	25	75	95	65.0	36.1	-0.47
6	5	75	95	58.3	47.3	-0.56
7	5	25	95	41.7	47.3	0.56
8	5	25	75	35.0	36.1	0.47
9	40	70	90	66.7	25.2	-0.23
10	5	70	90	55.0	44.4	-0.54
11	5	40	90	45.0	42.7	0.21
12	5	40	70	38.3	32.5	-0.09
13	30	60	95	61.7	32.5	0.09
14	10	60	95	55.0	42.7	-0.21
15	10	30	95	45.0	44.4	0.54
16	10	30	60	33.3	25.2	0.23

Notes: In M3 (P3) the cards depict the credits (probability of getting ₵100). The set of cards were chosen to include variation along three different dimensions. First, we wanted sufficient variation in the expected values of the cards. The expected values vary from 33 to 67. Second, we varied the standard deviation (S.D.) of the card values. The standard deviation varies from 20 to 47. Lastly, we varied the skewness in the card values. There are seven positively skewed, seven negatively skewed, and two symmetric distributions of values.

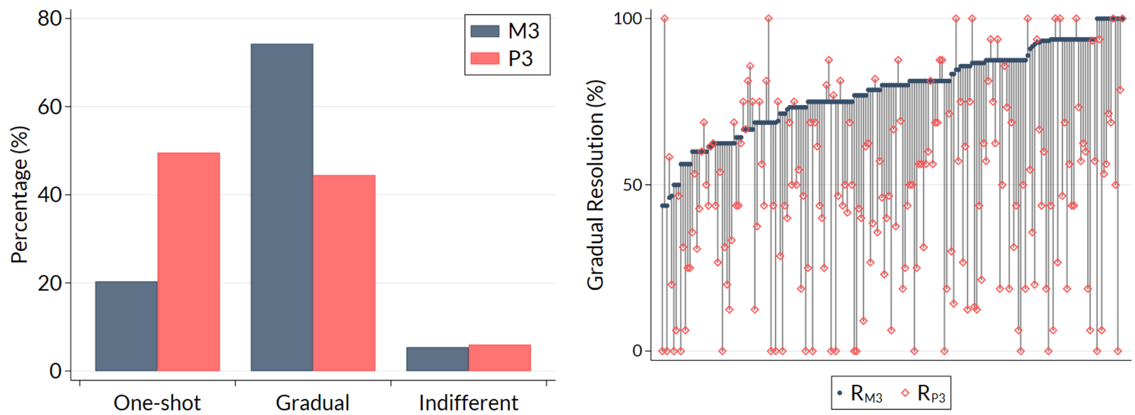

(a) Choices at the aggregate level

(b) Choices at subject level

Notes: In panel (b), each vertical line represents one subject. R_{M3} (R_{P3}) represents the proportion of times the subject chose gradual over one-shot resolution in the money (probability) domain.

Figure 2: Choices in M3 and P3

Result 1. *Subjects preferred gradual resolution in M3, but had mixed, revealed preferences in P3.*

This difference in preferences is also present at the subject level. In Figure 2(b), each vertical line represents one subject and has two points: one dark dot for M3 and one empty diamond for P3. Each of these points represents the ratio $R_k = \frac{d_k^G}{d_k^O + d_k^G}$, where $k \in \{\text{M3, P3}\}$, d_k^O and d_k^G are the number of times a subject chose to resolve the uncertainty fully and gradually, respectively, over the 16 problems. This ratio measures the proportion of times a subject chose gradual over one-shot resolution, excluding the rounds in which she was indifferent.As we can see, most subjects chose gradual resolution more often in M3 than in P3. The mean R_{M3} (R_{P3}) is 83% (49%), confirming that the aggregate result of Figure 2(a) also holds at the individual level. Furthermore, there was a substantial difference in the intensity of the subjects' preference for gradual resolution in M3 vs. P3 in that a substantial fraction of 43% of subjects has R_{M3} strictly greater than 0.5 but at the same time R_{P3} strictly lower than 0.5, thereby exhibiting a reversal in their revealed preferences.[8]

In the following we refer to the mean difference $\frac{1}{n}\sum_i (R_{i,M3} - R_{i,P3})$ as the average treatment effect (ATE) which indicates the average difference in behavior of our subjects facing the value and probability problems, respectively. An ATE of 32 percentage points is found consistently across all compositions of cards. Figure 3(a) presents the ATE as the lotteries presented in these treatments varied with respect to their mean, standard deviation, and skewness. As can be seen the ATE does not substantially vary with the expected value of the three cards, nor with their skewness (Figure 3(c)). In contrast, Figure 3(b) shows that differences in revealed preferences increase with the ex-ante uncertainty. As subjects faced wider spreads in the labels of cards, they disproportionately preferred gradual resolution with uncertain values: In M3, choices for gradual resolution increase from 60% to 90% as the standard deviation increases by 27.2 units. In P3, there is a smaller 15% increase in the relative preference for gradual resolution (from 39% to 54%). Thus, more choices for gradual resolution are generally associated with a higher variance of the uncertain object, but because the correlation is higher in the money domain the ATE increases with uncertainty. In a nutshell, with increasing uncertainty, the object over which the uncertainty is defined becomes relatively more important.

[8]To check whether these reversals would be consistent with stochastic or erratic choice, we asked additional 51 subjects how much we would need to pay them to change their choice in each M3 and P3. We find that subjects on average ask ₵5.1 in M3 versus ₵3.2 in P3 (two-sided t-statistics:4.59). This finding allows concluding that subjects have indeed stronger preferences for gradual resolution in the money than in the probability domain. The details are in Online Appendix A.1.

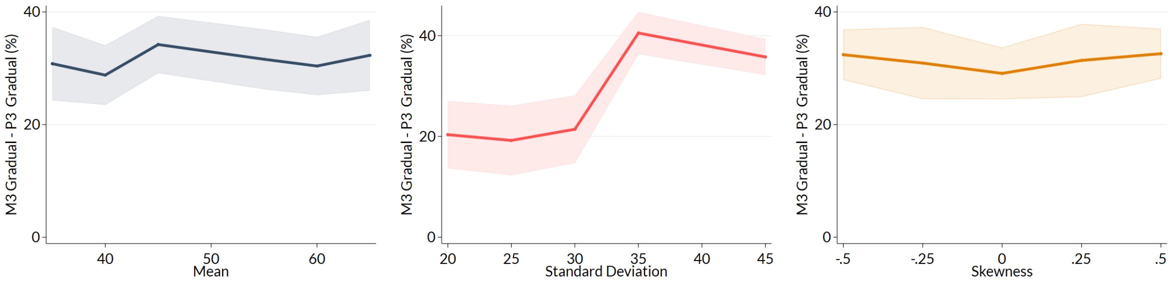

(a) Mean (b) Standard Deviation (c) Skewness

Note: The confidence interval shows 95% level.

Figure 3: Effect of cards composition

4.2 Robustness check 1: History in the four-card problem

Design. In the three-card problem we study subjects' preferences in a setting without history, in which subjects make only one decision and can only look forward when comparing the two options of gradual and one-shot resolution. Yet, in decision problems with longer horizons uncertainty is often resolved dynamically. As uncertainty is resolved over time, the DM is exposed to incoming news. A natural question that arises is to what extent subjects' preferences are stable and not affected by previous events. Put differently, are these preferences for temporal resolution only forward-looking, or does experiencing a negative (disappointing) or positive (elating) outcome alter subjects' preferences for future resolution decisions?[9]

In the four-card experiment, we study the effect of history with the problems M4 and P4. These experiments are almost identical to M3 and P3, except that, as the name suggests, they were modified by adding a fourth card. The four-card problem has two different stages: In stage 1, subjects see four different cards on their computer screen.[10]

[9] A literature studying information avoidance suggests that news has an effect on subsequent information acquisition decisions. For instance, Golman et al. (2017) find evidence for the ostrich effect, according to which investors are more reluctant to look up their portfolio after incoming bad news. In our framework, subjects have no possibility to avoid information, they can only delay some part of it. To the best of our knowledge, this history aspect has not been investigated in experiments on intrinsic preferences for resolution of uncertainty.

[10] In the Main experiment we implemented two versions of history treatment. The two versions of the four-cards game differ in whether subjects made an active resolution choice in the first stage of the four-card problem. This first stage is the stage that generates the history event for the subsequent stage. In the active choice version of the history treatment, subjects would be presented with four cards and then be asked whether they wanted to resolve gradually or all at once. In contrast, in the version presented in this section gradual resolution at the first stage is enforced. Among 200 subjects who participated in the Main experiment, 103 subjects were presented with the version without resolution choice in the first stage (i.e., gradual resolution was enforced). Here we present our results for these 102 subjects. The results of the version with an active first-stage resolution choices are qualitatively and quantitatively similar to what we describe in this section. The design and results for the active choice version can be found in Online Appendix Section A.3.

Table 2: 4 cards Lottery Parameters

#	Card 1	Card 2	Card 3	Card 4	Mean	S.D.	Skewness	Corresponding 3 cards in Table 1
1	20	40	60	80	50.0	25.8	0	#1-4
2	5	25	75	95	50.0	42.0	0	#5-8
3	5	40	70	90	51.3	37.1	-0.26	#9-12
4	10	30	60	95	48.8	37.1	0.26	#13-16

Notes: In M4 (P4) the cards represent the credits (probability to get ₵100). Corresponding 3 cards shows the possible set of 3 cards in stage 2 after one card is randomly removed.

After subjects see the cards' labels, the computer flips the cards over to conceal their labels and randomly shuffles them. Subjects then have to choose one card to be removed without knowing its identity. In stage 2, after learning which card was removed, they are presented with the three-card problem as in M3 and P3.

Table 2 shows the list of cards used in M4 and P4. It is important to note that after the first card is randomly removed in stage 1, the sets of cards in stage 2 correspond exactly to the chosen sets of cards in M3 and P3. For example, after one card is removed in the first stage, the #1 set of cards in Table 2 (row 1) corresponds to #1-#4 set of cards in Table 1 in the three-card problem. This design allows us to investigate the effect of news by comparing the decisions of subjects in identical three-card decision problems with and without history. Also in this part, subjects experimented with two practice decisions before making four payoff-relevant decisions. Likewise, there was no feedback about the final payoff between the four payoff-relevant decisions.

Our question here is whether the discrepancy in revealed preferences that we observe in M3 and P3 persists in a more dynamic setting where subjects saw one card removed.

Under KP, preferences for resolution of uncertainty are captured by the concavity of the function $h(\cdot)$. Assuming that this $h(\cdot)$ does not vary with the number of cards remaining, the DM should resolve uncertainty in the same way regardless of the feedback she receives over time. This is true for both uncertainty in values or probabilities. Learning which card was removed in the first stage of M4 and P4 should not reverse subjects' preferences revealed in M3 and P3. Therefore, the choices with history in stage 2 of M4 and P4 and the ones without history in the M3 and P3 treatments should be identical. Although they should behave in the same way, after a card is removed subjects may not be in the same emotional situation as they were in M3 and P3 where they simply faced a three card problem with no prior history.

Results. Figure 4 compares the ATE between the money and probability problems in stage 2 of the history treatment with the ATE that was observed with the same set of cards in the no-history three-card problems.

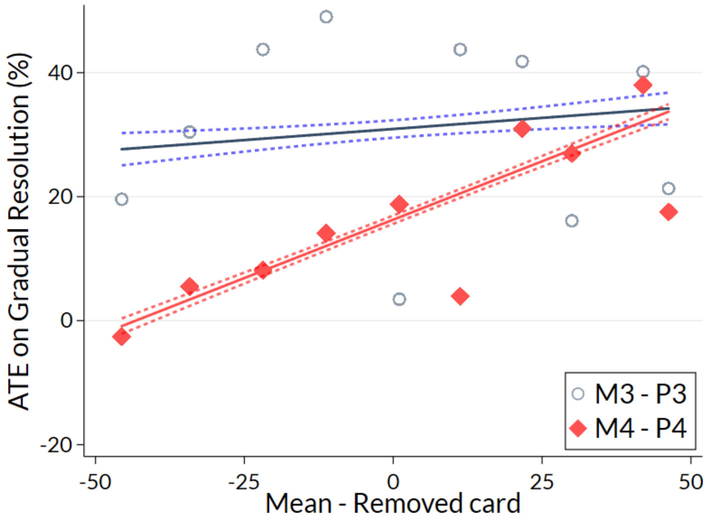

Notes: The diamonds represent the ATE after we divided the sample into 10 equal-sized bins using the difference between the ex-ante mean of the four cards and the number on the removed card. The empty dots represent the benchmark, which is the ATE in the no-history (three-card) treatment. The lines are obtained separately by regressing the dummy variable for gradual resolution on the Mean-Removed cards, M3 or M4 dummy variable, and the interaction term. We then use the regression estimates to predict the treatment effect across different values of (Mean- Removed cards) and draw the lines and 95% confidence interval.

Figure 4: History Dependence

We organize the data as a function of the difference between the ex-ante mean of the lottery and the number on the removed card. We use a binned-scatter plot to represent the data: Concretely, we group the data into 10 equal-sized bins according to their difference on the x-axis, and on the y-axis, for each bin we plot the ATE on resolution choices (i.e., how much more subjects chose gradual resolution in the money relative to the probability domain) as red diamonds. As a comparison benchmark, we also plot the corresponding ATE in the no-history treatments as blue empty dots. The lines are separately constructed through linear regressions with the dependent variable as a gradual resolution dummy and estimated treatment effect using the Mean-Removed cards, M3 or M4 dummy variable, and the interaction term. We predict the treatment effect across different Mean-Removed cards and draw the lines and 95% confidence interval. The higher this difference is (the further to the right of the x-axis), the more the removal of the card can be interpreted as positive news. For instance, if the subject saw originally the four cards $\{20, 40, 60, 80\}$ whose mean is 50 and, if the 20 card is removed, the difference is 30 which is good news since the worst card was removed.

We find that preferences for the resolution of uncertainty are history-dependent, but with a meaningful asymmetry. Learning that the worst outcome is no longer possible does not affect our previous result: subjects continue to prefer gradual resolution substantially more often with value compared to probabilistic uncertainty. However, if we compare the ATE at the extreme left side of the graph, we see that the ATE shrinks to insignificant levels as better outcomes are removed. In Appendix Figure 4, we show the proportion of gradual choices separately by domains, making apparent that the reduction of the ATE is mainly driven by preferences in the money domain: Although in both domains subjects are more likely to choose one-shot resolution after bad news, the effect is particularly strong with uncertain values. In M4, the relative preference for gradual resolution is 86% when the lowest card is removed, but falls sharply to 42% when the highest card is removed. This difference is significant at 1% level (t-test statistics: 7.20). In P4 some history-dependence is discernible but, compared to M4, substantially weaker and not significantly different even when the highest card is removed. The proportion of gradual resolution choices is 56% when the lowest card is removed compared to 40% when the highest card is removed.

In a nutshell, the history-dependence in the ATE can be attributed to two behavioral patterns. First, bad news moves preferences from gradual to one-shot resolution in both domains. Second, changes in preferences are larger with value than with probabilistic uncertainty, thereby reducing the ATE.

We draw two conclusions from our results. First, the resolution of uncertainty is evaluated relative to some reference lottery in the past. As options available in our four-card lottery are eliminated, preferences for the remaining lotteries change. Second, in our setting preferences for gradual resolution are disproportionately driven by high versus low outcomes, where high value outcomes have a particular valence compared to high probability outcomes. Hence, this treatment points to high-label cards as important decision factors that drive subjects' preferences for gradual resolution, and we find that the object of uncertainty matters less when these decision factors are eliminated. In Section 5, we propose a mechanism that emphasizes this comparative nature of the decision-making process.

Result 2. *Preferences for the resolution of uncertain values and probabilities become indistinguishable after negative news (desirable cards removed). Preferences in the value domain converged to those in the probability domain when high value outcomes were no longer available.*

4.3 Robustness check 2: The Loss Experiment

Design. In the Main experiment, subjects start out the decision problems with either a zero value or a zero probability of winning the big prize and at the end wind up with something nonnegative. While they probably hope for a big gain, they are never worse off materially at the end than they are at the beginning. Compared to how they start out, whatever happens to them is a gain (or at least not a loss). Yet, it might be just as natural to frame the problem in the loss domain where subjects start out with a ₵100 in the value setting or a 100 percent chance of receiving ₵100 in the probability setting, and then have that value or probability reduced as cards are removed. Note that our main objective here is to frame the two different objects of uncertainty in the loss domain, as opposed to introducing a negative consumption event. In the domain of gains we reported a greater proclivity for subjects to resolve uncertainty gradually in the value setting. The question is now whether this ATE persists when we move from the gain to the loss domain. Our translation from the gain to the loss domain does not alter any of the predictions we discussed above, but it very well may be that subjects care about resolving uncertainty differently across gains and losses. To test this hypothesis, we conducted identical three- and four-card problems in the loss domain. (For the sake of brevity, the results of the four-card loss treatment is presented in Appendix Section A.4.)

The Loss experiment was conducted on Zoom with 82 subjects across five sessions. These subjects earned, on average, $26.74 including $10 show-up fee for approximately 75 minutes.

The experiment's protocol was similar to the Main experiment. Subjects were endowed with ₵100 in **M3 Loss** and a 100% chance of getting ₵100 in **P3 Loss** prior to any decision. Our subjects then saw cards that represented how much they could lose from their initial value or probability endowment. To induce the same distribution of outcomes in the gain and loss treatments, we adjusted the labels of cards. For instance, in M3, the #1 set of cards consists of ₵40, ₵60, and ₵80. In the treatment **M3 Loss** this corresponds to losing from a ₵100 endowment the cards -₵60, -₵40, and -₵20. Likewise, in P3, subjects faced 40%, 60%, and 80%, while in the treatment **P3 Loss**, they started with a chance of 100% and then they faced a possible probability reduction of -60, -40, and -20 percentage points. In summary, the only difference between the gain and loss treatments is the framing of final outcomes relative to the starting position, whereas the distribution over final outcomes is identical across the gain and loss treatments.

Results. Fig 5(a) replicates Fig 2(a) by presenting subjects' choices in the loss domain. Relative to the gain frame, the proportion of gradual resolution choices in M3 (P3) loss decreases significantly by approximately 21 (6) percentage points (t-test; M3 test statistics: 15.43; P3 test statistics: 3.81). Yet, gradual resolution remains the preferred option in M3 Loss, while in P3 Loss one-shot resolution is now significantly preferred. Hence, as shown in Fig 5(b), the ATE persists in the loss frame. That is, while the loss frame attenuates preferences for gradual resolution in both domains, there still is a stronger desire to resolve

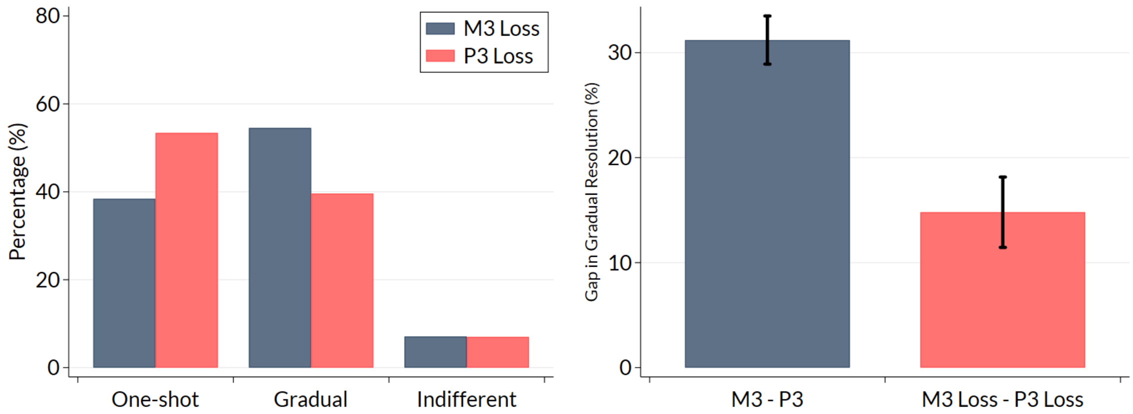

(a) Loss: Aggregated Choices (b) The ATE in the Gain vs. Loss Frame

Note: Panel (a) shows the distribution of choices in the loss treatment. Panel (b) shows the average difference in gradual resolutions in the Main (blue bar) and the Loss experiment (red bar). The confidence interval in panel (b) is at a 95% level.

Figure 5: 3-cards: Loss Treatment

gradually in the value compared to the probability domain.

Result 3. *In the aggregate, subjects preferred to resolve uncertain loss values gradually, but uncertain loss probabilities all at once. The treatment effect seen in the gain frame persists in the loss domain although it is weaker.*

Summary. Our experiments have presented us with several pieces of evidence. First, subjects' choices on how to resolve uncertainty depend on what is uncertain: subjects generally preferred to resolve value uncertainty gradually, but not probabilistic uncertainty, and this is true in both the gain and loss domains. Second, the object of uncertainty matters more with increasing uncertainty (large lottery variances) but less after bad news, suggesting that the best outcomes have a predominant role in governing these preferences. In the following, we propose an alternative explanation for these results using what we call Process Utility.

5 A Process Utility Approach

One question that arises when thinking about resolution of uncertainty is why anyone (let alone so many subjects in our experiment) would rather resolve uncertainty gradually when they could find out their payoff immediately. Many of the theories that discuss the choice between early and late resolution Palacios-Huerta (1999); Caplin and Leahy (2001); Dillenberger (2010); Koszegi and Rabin (2009) offer good reasons to resolve early;

perhaps to cut down on the anxiety of waiting or because of disappointment aversion. So, the onus seems to be on justifying gradual as opposed to immediate resolution.

In the context of our experiment we offer one such explanation which we call "Process Utility". Its main idea is that one values the process through which one's ultimate payoff is determined and not just the expected utility of the payoffs themselves. The decision to resolve uncertainty immediately (all at once) is a decision to allow the lottery one faces today to be the lottery or process that determines one's payment. (We use the terms *lottery, process* and *mechanism* interchangeably since we view lotteries as nothing more than a mechanism used to determine a DM's payoff.) It is the DM's preference over which **process** will be used to determine her payment that defines her choice over when to resolve uncertainty. The reason to resolve gradually, therefore, is a decision to reject the current payment lottery in the hope that, when new information arrives (and outcomes are eliminated as in our experiment), they will face a lottery that they consider better.

Note, this explanation differs from some of the other theories mentioned above in that the decision to resolve gradually may be driven by a cognitive rather than an affective consideration: Even in the absence of distressed or elated emotions in anticipation of what will happen in the future, the way in which subjects process information and compare different payoff mechanisms across different periods may induce them to strictly prefer one way of resolving uncertainty over the other one.

5.1 A Simple Illustration

One question that arises is how our DMs compare the different processes to decide whether it is worth waiting for a "better" lottery in the future. Many different methods could be used to compare processes, but our Process Experiments suggest that Generalized Expected Utility theory (GEU) (i.e., the classic Expected Utility theory and its modern day equivalents) is not responsible for this choice. While with GEU, we assume that DMs compare lotteries by taking expectations of utility outcomes, process utility allows the DM to value the process beyond its utility outcomes. As discussed by psychologists and economists in the past in different contexts, a reasonable way of comparing different processes with multiple attributes (such as lotteries) is one in which the DM compares the processes' attributes one by one.[11] In some sense, this approach looks for similarities and differences between lotteries or at least their attributes (Tversky and Simonson (see also 1993); Rubinstein (see also 1988) and Payne (1973) for a survey). Other alternative models that allow for contrasting lotteries to their alternatives include salience theory

[11]This approach is what psychologists call the "information processing" approach Payne (1973), see Slovic and Lichtenstein (1968)).

as in Bordalo et al. (2012) or preferences for simplicity as in Puri (2022) and Mononen (2022).[12] While we are agnostic as to which of these theories one uses to compare lotteries, we highlight that this comparison may not be done by taking expectations over the utility of prizes using either linear or non-linear decision weights.

To illustrate, consider again our M3 problem in which the DM faces three cards with the potential payoffs \$40, \$60, or \$80. The DM must choose between resolving the uncertainty over the three prizes gradually or all at once. Notice that this choice boils down to comparing different sets of lotteries. If she choose to resolve all at once, her payoff will be determined by the lottery $\ell_0 = \{\$40 : 1/3, \$60 : 1/3, \$80 : 1/3\}$. If she chooses to resolve uncertainty gradually, one card will be removed and her payoff will be determined by one the following binary lotteries $\ell_1 = \{\$40 : 1/2, \$60 : 1/2, \$80 : 0\}$, $\ell_2 = \{\$40 : 1/2, \$60 : 0, \$80 : 1/2\}$ or $\ell_3 = \{\$40 : 0, \$60 : 1/2, \$80 : 1/2\}$. Thus, when choosing how to resolve uncertainty, the DM possibly contemplates which lottery she would rather have determine her payoff: Does she want a three-card lottery or take a chance and have one of three equally likely two-card lotteries? Under standard expected utility, the DM would be indifferent as to when to resolve since the utility of the three-card lottery ℓ_0 equals the expected utility of the three possible future two-cards lotteries ℓ_1, ℓ_2 and ℓ_3, generating the same distribution over terminal prizes. However, this indifference may be broken if the DM contrasts the *process* of having her payoff determined by a specific random mechanism (lottery ℓ_0) relative to the possibility of an alternative mechanism (e.g., lottery ℓ_1). The decision to resolve gradually is then a decision to search for what the DM considers a better lottery to determine her final outcome.

We define the function $\delta(\ell_i, \ell_j)$ as the measure that captures how much more the DM values the lottery ℓ_i over the lottery ℓ_j. In our example above, the DM will strictly prefer lottery ℓ_0 over lottery ℓ_1 if $\delta(\ell_0, \ell_1) > 0$. Here, we do not impose a functional form for $\delta(\cdot)$ but will let the data speak in the following subsection. For instance, under GEU, $\delta(\ell_0, \ell_1) > 0$ if and only if, in expectation, the DM values the prizes in lottery ℓ_0 more than the ones in lottery ℓ_1 (i.e, $V(\ell_0) > V(\ell_1)$ where $V(\ell)$ is the DM's value function for the lottery ℓ). Alternatively, a DM may depart from GEU in that she does not make her choice by calculating expected utility (i.e., by taking an expectation over the utility of outcomes using either linear or non-linear (weighted) probabilities), but rather evaluates lotteries by comparing their various attributes or features one by one. For example, she may view the three-card lottery $\ell_0 = \{\$40 : 1/3, \$60 : 1/3, \$80 : 1/3\}$ as consisting of 3 main attributes: a) the 1/3 chance of getting the high prize \$80; b) the 1/3 chance of getting \$60; and c)

[12] The idea that, when presented with lotteries, subjects may derive a utility that is different from the utility over consequences was brought to attention in a literature on gambling (see e.g., Conlisk, 1993; Menestrel, 2001; Diecidue et al., 2004), but to the best of our knowledge has not been put to a test.

the 1/3 chance of getting \$40. So, the attributes she contemplates and derives utility from are probability-outcome pairs. From each of these attributes she derives utility $v(x,p)$ where x and p denote the possible prize and the corresponding probability of receiving the prize, respectively, with $v(0,p)=0, v(x,0)=0, \frac{\partial v(x,p)}{\partial x}\geq 0, \frac{\partial v(x,p)}{\partial p}\geq 0$. In contrast to standard expected utility, $v(x,p)$ may not be separable in x and p.

Importantly, in comparing which payoff process (ℓ_0 vs. ℓ_1) she prefers, she compares the lotteries attribute by attribute. For instance, assume that lottery ℓ_0 is the three card lottery described above while lottery ℓ_1 is a particular lottery derived from ℓ_0 by eliminating the \$80 card. This yields a lottery with a 0% probability of getting \$80 but a 50% probability each of receiving either \$60 and \$40, respectively. When the DM compares these lotteries, she contrasts the attributes by computing for each attribute the difference $\phi(v_{\ell_0}(x,p)-v_{\ell_1}(x,p))$, where $\phi(\cdot)>0$ is increasing in the value difference between two attributes.[13] If the function ϕ is convex ($\phi'(\cdot)\geq 0$), preferences over lotteries are disproportionately driven by the most different attributes.

To illustrate, assume that attributes are compared in an ordered fashion in that the most preferred attributes are compared to each other, then the second most etc., until the DM reaches the least preferred. Other pairwise comparisons are conceivable.[14] The comparison of attributes $a_i, a_j \in \mathcal{A}$ is then aggregated to a value $\delta(\cdot)$:

$$
\begin{aligned}
\delta(\ell_0,\ell_1) = & \sum_{a_i,a_j\in\mathcal{A}} \phi(v_{\ell_0}(a_i)-v_{\ell_1}(a_j)) \\
& = \phi(v(\$80,1/3)-v(\$80,0))+ \\
& \phi(v(\$60,1/3)-v(\$60,1/2))+ \\
& \phi(v(\$40,1/3)-v(\$40,1/2))
\end{aligned}
$$

The DM will strictly prefer lottery ℓ_0 over lottery ℓ_1 if $\delta(\ell_0,\ell_1)>0$, i.e., if in the aggregate, she prefers the attributes of ℓ_0 to the ones of ℓ_1.

Finally, as the DM considers whether or not to resolve uncertainty all at once, she compares the three-card lottery ℓ_0 to each of the possible, future two-card lotteries

[13] More complex is the case where two lotteries have a different number of attributes. One possibility would be to rank the attributes along one specific dimension and collapse similar attributes of the multiplex lottery to equalize the number of attributes. An alternative would be to focus on the k most relevant attributes of each lottery, where k is weakly smaller than the minimum number of attributes. At this point, we do not propose a specific procedure, but emphasize that there are multiple ways of departing from an expected utility comparison when evaluating lotteries.

[14] For example, the framing or the response mode may determine which attributes are compared to each other and, hence, trigger a preference reversal across different decision contexts.

(ℓ_1, ℓ_2, ℓ_3). She will prefer to resolve the lottery all at once if, in expectation, she values the attributes of the lottery ℓ_0 more:

$$\ell_0 \succcurlyeq \{\ell_1 : 1/3, \ell_2 : 1/3, \ell_3 : 1/3\} \texttt{ if } \quad \frac{1}{3} \sum_{j \in \{1,2,3\}} \delta(\ell_0, \ell_j) \geq 0. \tag{1}$$

To summarize, one possibility is that the DM first contrasts lottery ℓ_0 to each of the other possible lotteries $\{\ell_1, \ell_2, \ell_3\}$, and then aggregates her pairwise preferences. She will then choose gradual resolution if, in the aggregate, she finds one of the binary lotteries $\{\ell_1, \ell_2, \ell_3\}$ disproportionately more attractive than the immediate lottery ℓ_0.

5.2 Process Utility

The approach outlined above is one of many conceivable ways to compare lotteries. In this section we generalize the concept of process utility by defining its two main properties. The first idea that we posit is that the DM has a preference over the different random processes that may determine her payoff (here the possible lotteries in the different time periods). Importantly, the DM may want to resolve gradually because gradual resolution offers the possibility of facing different (and more desirable) payoff processes than the one given by one-shot resolution. For instance, a DM who finds the binary lottery $\ell_3 = \{60 : 1/2, 80 : 1/2\}$ far more appealing than the trinary lottery $\ell_0 = \{40 : 1/3, 60 : 1/3, 80 : 1/3\}$ may be inclined to choose gradual resolution in the hope of having her payoff determined by the more attractive lottery ℓ_3. In a nutshell, choosing gradual resolution allows one to face different payoff processes with the hope of having one of the better processes eventually determine one's outcome.

Hypothesis 2 tests our main idea that process utility underlies choices for temporal resolution. If subjects find certain two-card lotteries sufficiently desirable compared to the three-card lottery associated with resolving immediately, they will be more likely to resolve gradually in the hope of having one of these desirable lotteries determine their outcome. That is, subjects' main motivation to resolve gradually is driven by their aspiration to have their payoff determined by a preferred payoff lottery. In our setting, it follows that the DM will want to resolve gradually if she appreciates the best possible two-card lotteries (ℓ_3 in our example above) disproportionately more than the other lotteries. If not, she would rather find out her payoff at once, or at least be indifferent.

Hypothesis 2. *The more intensely subjects value desirable two-card lotteries (ones with the worst card removed), the more likely they are to resolve uncertainty gradually.*

Our last hypothesis is inspired by our findings in the Main experiment. If the valuation of lotteries determines preferences for gradual resolution, and these preferences are stronger in the value domain, subjects will exhibit stronger utility differential between value lotteries than between probability lotteries.

Hypothesis 3. *The comparison of lotteries is more sensitive to differences in values than differences in probabilities suggesting, as we have seen, that subjects should have stronger preferences for resolution in the M3 relative to the P3 treatment.*

While we plan to remain agnostic about the exact procedure subjects use to compare lotteries or the functional form and arguments of $\delta(\cdot)$, the validity of our hypothesis tests will clearly depend on the way we measure preferences over lotteries. A straightforward way of eliciting preferences over lotteries is to compare subjects' certainty equivalent for the respective lotteries. However, this method reaches its limits if preferences cannot be represented with GEU. For example, suppose that our subjects view lotteries as multi-attribute objects and compare them on the basis of their salient attributes or properties. Various theories exist on how people might do this. Slovic and Lichtenstein (1968) suggest that lotteries can be compared attribute by attribute. When faced with two lotteries a DM may look at each lottery's big prize and the probability of receiving it, small prize and its probability of occurring, and compare them attribute by attribute in order to define their preference between them. In some sense this approach looks for similarities between lotteries or at least their characteristics (Tversky and Simonson, 1993; Rubinstein, 1988; Payne, 1973). Other alternative models that allow for contrasting the characteristics of lotteries to the ones of an alternative process include salience theory as in Bordalo et al. (2012) or preferences for simplicity as in Puri (2022) and Mononen (2022). To test these ideas we ran our Process Experiment.

5.3 Testing process utility experimentally

Design. The Process Experiment has two objectives. First, we elicit subjects' preferences over the different lotteries that may determine their payoff. Importantly, these preferences are elicited in an atemporal context in which we refrain from mentioning the timing of payments. Second, to test our idea of process utility we correlate preferences over payoff lotteries with subjects' resolution choices. Our conjecture is that the decision to resolve gradually rather than all at once is a decision to search and hope for a better lottery to determine one's payment.

Sixty subjects participated in the Process Experiment. The experiment was conducted in person in the NYU CESS laboratory in October 2022 across four sessions. Subjects earned, on average, \$29.4 including \$7 show-up fee for approximately 60 minutes. The experiment has two treatments in a between-subject design, in which we present the subjects with decision problems either in the money or in the probability domain. Each treatment consists of four different parts. In Parts 1 to 3 we elicit subjects' preferences over lotteries, and in Part 4 their preferences for temporal resolution.

Our attempt to investigate subjects' preferences over lotteries contrasts two approaches. Specifically, if subjects' preferences adhere to GEU, subjects will evaluate each lottery in isolation, and this valuation will be reflected in a certainty equivalent. Comparing the certainty equivalents will then inform us about subjects' preferences over lotteries and their willingness to switch between them. However, if subjects evaluate these lotteries differently, for instance, as suggested above with an attribute-by-attribute comparison, how attractive they find a lottery will depend on how this lottery compares to another one. In that case, subjects' willingness to switch between pairs of lotteries may substantially differ from their willingness to switch based on GEU. In other words, if GEU does not hold, the information elicited in subjects' WTS between lotteries may substantially deviate from the information elicited in the CE. We will be more specific below.

To obtain WTS, we elicit preferences in a pairwise comparison. In Part 1, we present subjects with pairs of lotteries where one of them is a three-card lottery and the other one is a two-card lottery. The two-card lottery is derived from the three-card lottery by eliminating one of the three cards. When faced with such a pair of lotteries subjects must state whether they would prefer to have their payoff determined by the three-card or two-card lottery (for instance, in our example above we would ask them to choose between ℓ_0 versus ℓ_1). Subjects specify their choices for three different pairs of lotteries $\{(\ell_0, \ell_1), (\ell_0, \ell_2), (\ell_0, \ell_3)\}$ where one of the lotteries, ℓ_0, is always the same three-card lottery. Subjects stated their preference for three such sets of lotteries comparing ℓ_0 to (ℓ_1, ℓ_2, ℓ_3). In the second part of the experiment, after subjects have made their choices for the three pairs of lotteries, we present them again with their choice for each pair and ask them how much money we would have to pay them to have them switch their choice from the lottery they chose to their less preferred lottery. This willingness to accept to switch choice reflects their intensity of preference to have their payoff determined by a particular type of lottery. Subjects submitted their decisions for, again, the three sets of lotteries with $(\ell_0, \ell_1, \ell_2, \ell_3)$, leading to a total of nine elicitations of willingness to accept to switch choices.

In Part 3, we ask subjects for their CE for each lottery that they saw in the first two parts of the experiment in isolation, resulting in 12 CEs. We elicit their CE with a multiple price list in which subjects chose between the lottery and an increasing sure payoff.

The comparison between responses in parts 2 and 3 allows us to test whether subjects evaluate lotteries in a manner different than that prescribed by GEU. More precisely, if subjects' valuations of lotteries depend only on the utility of expected outcomes, then their willingness to switch between two lotteries should correspond to the difference in their respective CEs. That is, subjects' required compensation to switch between two lotteries can be retrieved by computing the differences in their respective CEs (see Appendix Section B.1). For the lottery-comparison procedure we outlined above, this need not be the case.

Finally, in the last part of the experiment, we present subjects with the same three-card lotteries used for the previous parts and ask them to choose whether to resolve uncertainty all at once or gradually. This part allows us to test Hypothesis 2 by assessing the relevance of process utility for the resolution of uncertainty. If process utility is at work in determining subjects' choices, subjects who value some two-card lotteries disproportionately more would also choose to resolve gradually. Hence, their lottery valuation (measured either in CE or WTS) should be a predictor of their resolution choices.

Results. We start our analysis with a comparison of CEs, which is the more prevalent elicitation method for lottery valuation. We consider four CEs, where CE_0 represents the CE for the thee-card lottery (ℓ_0) and CE_1, CE_2, CE_3 represent the CEs for ℓ_1, ℓ_2, ℓ_3, respectively. One main question that arises from the previous analyses is whether resolution choices can be explained by a difference in the way subjects value the lottery prizes across domains. For instance, if subjects engage in probability weighting, they may perceive the lotteries in the probability treatment more similarly and, hence, care less about the resolution of uncertainty. While our data is consistent with some inverse S-shaped probability weighting, we find that our subjects do not value equivalent lotteries significantly differently across domains. The aggregated CE is 48.77 in M3 and 46.55 in P3 (t-statistics:1.60). In Online Appendix A.6, we compare the CEs for the different types of lotteries and find a similar result.

We also compute three "CE Diff" measures CE_1 -CE_0, CE_2 -CE_0, and CE_3 -CE_0. Again, these measures based on CEs capture how much subjects value each of the two-card lotteries compared to the three-card lottery. If expected utility was the basis of comparison, the subjects' WTS between two lotteries should coincide with the difference between the corresponding CEs.

Figures 6(a) and 6(b) compare the two measures WTS and CE Diff for M3 and P3 respectively. Figure 6(a) with a darker (lighter) blue bar shows the mean WTS (CE Diff) across three different cases. The x-axis depicts three different cases depending on which card was removed from the three-card lottery. In the money treatment, we observe a systematic pattern. First, in all tasks subjects' preferences reflect the objective dominance relation between lotteries. When the best card was removed, subjects preferred the three-card to the two-card lottery with both measures. Conversely, when the worst card was removed, they preferred the two-card lottery. When the middle card was removed, they were generally close to indifference (except for the WTS measure in the money treatment, which shows a preference for the three-card lottery).[15] Second, preferences are more exaggerated with the WTS measure. For instance, after the worst card was removed, subjects required a compensation that is significantly higher than the difference in CEs (34.1 in WTS versus 17 in CE diff; t-statistics: 9.10). When the middle card was removed the average WTS is -7.9, while the average difference in CE is -2.8 (t-statistics: 2.18). Lastly, when the best card was removed the average WTS is -24.2, almost double the average difference in CE of -12.5 (t-statistics: 6.99).

Figure 6(b) shows the results for the probability treatment. The darker (lighter) red bar shows the mean of the WTS (CE diff) measure in each case. While we find a similar pattern with the probability lotteries, the gap between the two measures WTS and CE diff is substantially smaller. When the worst card is removed, the mean WTS is 20.5 and the mean of CE diff is 11.43 (t-statistics: 3.79). If the middle card is removed, the mean of WTS is -1.1 and the mean of CE diff is -1.6 (t-statistics: 0.17). Lastly, when the best card is removed, the mean WTS is -18.2 versus -14.5 for CE diff (t-statistics: 2.02).[16]

These results show that subjects evaluated lotteries differently, depending on the elicited measure. Turning to our main inquiry, the main question that arises is whether either of these atemporal measures of preferences can be linked to subjects' choices for temporal resolution. We now correlate these lottery valuations with subjects' resolution choice to test Hypothesis 2. With process utility, a DM will strictly prefer gradual resolution if she values the future two-card lotteries disproportionately more than the

[15] Most lotteries can be ranked according to first-order stochastic dominance. Whether subjects recognized this dominance may be an indicator of their understanding of or attention to the task. None of the subjects violated dominance in the elicitation of WTS between monetary lotteries. In the probability treatment, 6% of choices violate the rational choice, which means they preferred strictly worse two-card lotteries to the three-card lottery (or vice versa). The elicitation of CEs induced more violations: 11% (16%) of choices in M3 (P3) violated dominance. This means subjects reported higher (lower) CEs for a two-card lottery that was strictly worse (better) than the three-card lottery.

[16] One potential reason for why the WTS differ from differences in CE may be an artificial endowment effect. Technically, subjects are not endowed with a concrete object when we elicit their WTS, but some subjects may make their choice as if they were endowed with their preferred lottery. However, we doubt that the endowment effect is the main driver of the differences between measures. First, it is not clear why the endowment effect should be stronger with monetary than with probabilistic uncertainty and, more importantly, it is difficult to rationalize the correlation between resolution choices and WTS that we discuss below with an endowment effect. If an endowment effect would matter in the comparison of immediate and future lotteries, we should observe a preference for one shot resolution as subjects would be more inclined to resolve their current lottery. This is not what we observe.

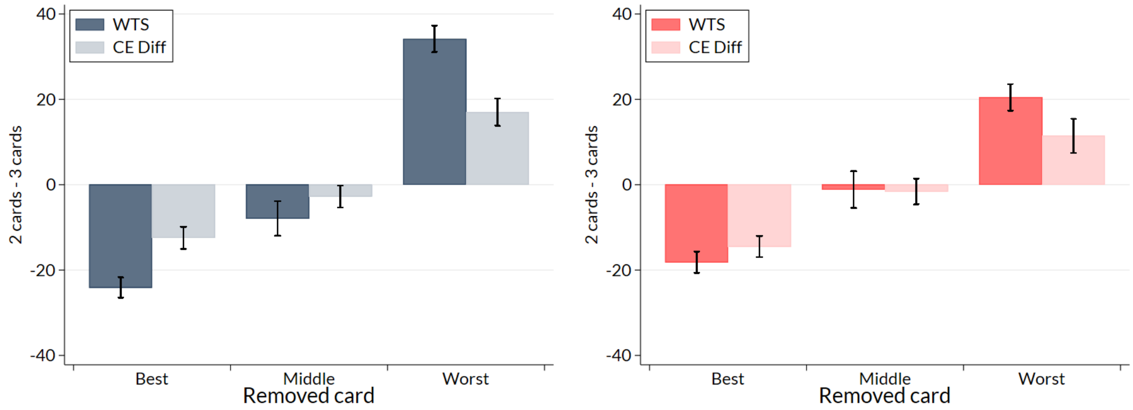

(a) M3 (b) P3

Notes: The bar represents 95% confidence interval.

Figure 6: WTS and differences in CEs

immediate three-card lottery. That is, we expect subjects' WTS to be positively correlated with the gradual resolution choice because a higher level of WTS means subjects value the two cards more than the baseline three cards. To confirm this, we put three WTSs as explanatory variables and run a logit regression with, as a dependent variable, a dummy that takes the value one if the subject chose gradual resolution. For notation, we define $\text{WTS}_{\text{R=r}}$ where r∈{Best, Mid, Worst}. For instance, $\text{WTS}_{\text{R=Best}}$ refers to the WTS when two cards evolve from three cards after removing the best card (which, on average, is negative for both M3 and P3 in Figure 6). We use the same notation for CE $\text{diff}_{\text{R=r}}$ where r∈{Best, Mid, Worst}.

Columns (1) and (2) in Table 3 show the results of the logistic regression for M3. As shown in column (1), choices for gradual resolution correlate positively with all three $\text{WTS}_{\text{R=Best}}$, $\text{WTS}_{\text{R=Mid}}$, $\text{WTS}_{\text{R=Worst}}$. Positive coefficients mean here that, when subjects prefer the two cards over the three cards, they tend to choose gradual resolution. However, the marginal effects of the WTSs vary across different pairs of lotteries. One unit increase in $\text{WTS}_{\text{R=Best}}$ and $\text{WTS}_{\text{R=Worst}}$ is associated with an increased propensity of 1.7% and 1.8%, respectively, to choose gradual resolution. While $\text{WTS}_{\text{R=Mid}}$ has a relatively small effect of 0.7%. Column (1) contrasts with column (2) where preferences measured with CEs cannot explain resolution choices. Focusing on the GEU framework would have led us to reject the idea of process utility. However, preferences measured with WTS strongly correlate with subjects' preferences for resolution. In other words, considering a preference model that is more general than GEU allowed us to detect a determinant in preferences for resolution and, most importantly, this determinant is not linked to the temporal nature of the decision problem. Subjects who tend to prefer the two-cards lotteries disproportionately more than the three-cards lotteries (maybe because they were engaging in an attribute comparison) were also more likely to resolve gradually.

Columns (3) and (4) in Table 3 show comparable, albeit weaker, results for P3. In column (3), we again find positive coefficients between WTSs and gradual resolution, but two of them are nonsignificant, while column (4) resembles column (1) in that we do not find any significant effect of preference measures based on CEs.

Our analysis confirms our intuition that preferences from temporal resolution are related to the manner in which subjects contrast future to current payoff lotteries. Our subjects' resolution choices correlate with their preferences over lotteries based on a WTS measure. The more they prefer the two-card lotteries compared to the three-card lotteries, the more they tend to choose gradual resolution.

Now that we have identified a determinant in resolution choices, let us get back to our main puzzle of understanding why subjects' preferences for temporal resolution differed across domains. Our analysis provides us with two separate reasons for this. The first one aligns with our idea of Process Utility. Our subjects have more pronounced preferences over pairs of lotteries in the money than in the probability domain. For each set of lotteries, the standard deviation of the WTS across the three different contingencies (Best, Middle, Worst) is significantly higher for money than for probability lotteries (M3: 32.541 and P3: 24.371; t-statistics:4.74). A potential reason for this might be probability insensitivity in that eliminating a probability outcome has a weaker effect on subjects' perception of the payoff process than eliminating a monetary outcome. Hence, because in the probability domain subjects perceive the future, two-cards lotteries to be more similar to the current, 3-card lotteries, they care less about the temporal resolution of uncertainty. This is what Figure 6 suggests.

In addition, our regression analysis in Table 3 also shows that subjects respond stronger to their lottery preferences in the money domain. The coefficient for the WTS measures are higher in the money relative to the probability domain, implying that preferences over lotteries have a stronger effect on preferences for temporal resolution in the money domain. At this point, we can only speculate as to why this is the case, but our suspicion is that subjects care more about the outcomes they understand better. We presume that subjects have a better grasp of what to expect in the future when they are able to conceive concrete, monetary outcomes, whereas the probabilities constitute a more abstract concept. It is easier to imagine oneself holding \$70 tomorrow, but learning that one's outcome tomorrow will be a 70% chance of being successful does not come with a clear mental image. In line with this, a recent strand of experimental studies in neuroscience has highlighted the importance of certainty for anticipatory neural processes (e.g., Bruhn et al., 2014; Johnen and Harrison, 2020). In the M3 and M4 tasks, the resolution occurs about future certain payoffs, while in the P3 and P4 tasks, resolution is over prospects with uncertain payoffs. In other words, imagining receiving a certain payoff of ₵40 in the future may trigger more intense anticipatory reactions than imagining receiving a prospect with its residual uncertainty over final payoffs. This is what our regression estimates in Table 3 suggest.

We believe these two effects come together when it comes to resolving uncertainty over time. Our subjects have stronger preferences for the gradual resolution of monetary uncertainty because, in the money domain, they perceive stronger differences between the different payoff processes they may be exposed to over time, and because they care more about these perceived differences in monetary prizes than in probability prizes.[17]

[17]Note that our findings in the Process Experiment also provide a potential explanation for the findings in the History treatment. If subjects evaluate lotteries by comparing their attributes, how much they value a future lottery will depend on the original set of outcomes. A future lottery will be valued less if it is contrasted to an initial lottery with high outcomes. Then, removing the best outcomes from the set of possible outcomes is equivalent to removing the best possible future lotteries, and leads to a preference for immediate resolution.

Table 3: The Effect of WTS and CE diff on Resolution Choices

	(1)	(2)	(3)	(4)
Sample	M3		P3	
Dependent Variable: dummy for gradual resolution				
$WTS_{R=Best}$	0.086***		0.020	
	(0.031)		(0.021)	
$WTS_{R=Mid}$	0.037**		0.010	
	(0.016)		(0.013)	
$WTS_{R=Worst}$	0.093***		0.032*	
	(0.034)		(0.018)	
CE $Diff_{R=Best}$		-0.007		0.020
		(0.018)		(0.023)
CE $Diff_{R=Mid}$		-0.020		-0.002
		(0.027)		(0.034)
CE $Diff_{R=Worst}$		0.017		0.028
		(0.015)		(0.017)
Constant	0.260	0.352	-0.473	-0.222
	(0.846)	(0.354)	(0.534)	(0.570)
Marginal effect R=Best	0.017***	-0.002	0.005	0.005
	(0.005)	(0.004)	(0.005)	(0.006)
Marginal effect R=Mid	0.007**	-0.004	0.003	-0.006
	(0.003)	(0.006)	(0.003)	(0.008)
Marginal effect R=Worst	0.018***	0.004	0.008*	0.007*
	(0.005)	(0.003)	(0.004)	(0.004)
Log-Likelihood	-43.281	-54.105	-56.890	-56.715
Observations	88	88	87	87

Notes: Logit regression with standard errors clustered at individual levels.

6 Discussion

Our experiments uncover a discrepancy in the way subjects choose to temporally resolve uncertain payoff values and uncertain probabilities. While our subjects largely preferred to resolve uncertainty over values gradually, a substantial fraction of them opted for resolving uncertain probabilities in one shot. This result stems from the fact that preferences are more pronounced when uncertainty is over monetary payoffs rather than probabilities. Clearly, variations in values and probabilities do not induce the same perception of uncertainty.

To understand why subjects' preferences vary with what is uncertain, we need to identify the main drivers of these preferences for temporal resolution. We use the discrepancy of behavior in the probability and money domains to better understand these preferences' determinants. Two robustness checks provided us with additional information. The first one is that this discrepancy persists in the loss domain. The second one is that in our four-card experiment, where removing cards affected our subjects' preferences, our subjects seem to engage in counterfactual thinking by thinking about what can and could have been.

Inspired by these additional insights, we propose to connect preferences for temporal resolution to an atemporal concept of information processing that we call process utility. Subjects possibly have a preference over the random process that determine their payoff, and since resolution in different periods entail different payoff lotteries, they may prefer to resolve gradually in the hope of facing a more attractive payoff lottery in the future. We allow such an approach to deviate from GEU framework and be based on the comparative evaluations of lotteries. While several theories modelling preferences for uncertainty resolution rely on some sort of counterfactual reasoning (Gul et al., 2021; Koszegi and Rabin, 2009; Palacios-Huerta, 1999; Caplin and Leahy, 2001), our explanations differ in that it is completely detached from the timing aspect of the decision problem. The resolution choice boils down to a comparison of lotteries that may determine one's payoff. In our design, one-shot resolution corresponds to a multi-outcome lottery, while gradual resolution imply that one's payoff is determined by a simpler lottery. Consequently, a DM who values some of these lotteries disproportionately more than others will not be indifferent toward the resolution of uncertainty.

Our experiments have provided us with several insights. First, our findings points to the limitations of the GEU framework. When we elicit subjects' valuations for the prizes within the GEU framework, we do not find significant differences in the way they value monetary versus probability prizes, and these valuations are not predictive of resolution choices. However, deviating from the GEU framework toward a more general information processing approach allows us to expose a link between resolution choices and the way subjects compare payoff lotteries to each other– both within and across domains. Subjects perceived a bigger contrast between the future, two-card and the immediate, three-card lotteries in the money than in the probability domain, and the more they preferred the future lotteries, the more they would choose gradual resolution. In a nutshell, we provide the first evidence of a correlation between temporal preferences for resolution of uncertainty and an atemporal concept that is based on the evaluation of lotteries.

In addition, we also find that, holding constant subjects' valuations for lotteries, the link between process utility and preferences for temporal resolution is stronger in the money than in the probability domain. Our presumption is that the extent to which process utility weighs in as a decision factor in the resolution of uncertainty depends on the extent to which these processes generate a concrete mental image; but this hypothesis remains to be tested in future research. A more important question that our experimental findings raise is: Do preferences for temporal resolution of uncertainty depend on the timing aspect at all, or are they pinned down by the way subjects process information when evaluating the available options?

References

Ahlbrecht, M. and M. Weber (1997): "Preference for Gradual Resolution of Uncertainty," *Theory and Decision*, 43, 167–185.

Alvarez, F., L. Guiso, and F. Lippi (2012): "Durable Consumption and Asset Management with Transaction and Observation Costs," *American Economic Review*, 102, 2272–2300.

Anagol, S. and K. J. Gamble (2013): "Does Presenting Investment Results Asset by Asset Lower Risk Taking?" *Journal of Behavioral Finance*, 14, 276–300.

Andries, M. and V. Haddad (2020): "Information Aversion," *Journal of Political Economy*, 128.

Arai, D. (1997): "Temporal resolution of uncertainty in risky choices," *Acta Psychologica*, 96, 15–26.

Bellemare, C., M. Krause, S. Kröger, and C. Zhang (2005): "Myopic loss aversion: Information feedback vs. investment flexibility," *Economics Letters*, 87, 319–324.

Benartzi, S. and R. H. Thaler (1995): "Myopic Loss Aversion and the Equity Premium Puzzle," *The Quarterly Journal of Economics*, 110, 73–92.

Bordalo, P., N. Gennaioli, and A. Shleifer (2012): "Salience theory of choice under risk," *Quarterly Journal of Economics*, 127, 1243–1285.

Bruhn, P., S. Huette, and M. Spivey (2014): "Degree of certainty modulates anticipatory processes in real time," *Journal of Experimental Psychology: Human Perception and Performance*, 40, 525–538.

Caplin, A. and J. Leahy (2001): "Psychological Expected Utility Theory and Anticipatory Feelings," *The Quarterly Journal of Economics*, 116, 55–79.

Chen, D. L., M. Schonger, and C. Wickens (2016): "oTree-An open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 9, 88–97.

Chew, S. H. and J. L. Ho (1994): "Hope: An empirical study of attitude toward the timing of uncertainty resolution," *Journal of Risk and Uncertainty*, 8, 267–288.

Conlisk, J. (1993): "The utility of gambling," *Journal of Risk and Uncertainty*, 6, 255–275.

Diecidue, E., U. Schmidt, and P. P. Wakker (2004): "The Utility of Gambling," *The Journal of Risk and Uncertainty*, 29, 241–259.

Dillenberger, D. (2010): "Preferences for One-Shot Resolution of Uncertainty and Allais-Type Behavior," *Econometrica*, 78, 1973–2004.

Ely, J., A. Frankel, and E. Kamenica (2015): "Suspense and Surprise," *Journal of Political Economy*, 123, 215–260.

Falk, A. and F. Zimmermann (2017): "Beliefs and Utility: Experimental Evidence on Preferences for Information," Working paper.

Frederick, S. (2005): "Cognitive reflection and decision making," *Journal of Economic Perspectives*, 19, 25–42.

Gneezy, U., E. Haruvy, and A. E. Roth (2003): "Bargaining under a deadline: Evidence from the reverse ultimatum game," *Games and Economic Behavior*, 45, 347–368.

Gneezy, U. and J. Potters (1997): "An Experiment on Risk Taking and Evaluation Periods," *The Quarterly Journal of Economics*, 112, 631–645.

Golman, R., D. Hagmann, and G. Loewenstein (2017): "Information avoidance," *Journal of Economic Literature*, 55, 96–135.

Gul, F., P. Natenzon, and W. Pesendorfer (2021): "Random Evolving Lotteries and intrinsic Preference for Information," *Econometrica*, 89, 2225–2259.

Haigh, M. S. and J. A. List (2005): "Do Professional Traders Exhibit Myopic Loss Aversion? An Experimental Analysis," *The Journal of Finance*, 60, 523–534.

Huettel, S. A., A. W. Song, and G. McCarthy (2005): "Decisions under Uncertainty: Probabilistic ContextInfluences Activation of Prefrontal and Parietal Cortices," *Journal of Neuroscience*, 25, 3304–3311.

Johnen, A. K. and N. R. Harrison (2020): "Level of uncertainty about the affective nature of a pictorial stimulus influences anticipatory neural processes: An event-related potential (ERP) study," *Neuropsychologia*, 146.

Karlsson, N., G. Loewenstein, and D. Seppi (2009): "The ostrich effect: Selective attention to information," *Journal of Risk and Uncertainty*, 38, 95–115.

Kocher, M. G., M. Krawczyk, and F. van Winden (2014): "'Let me dream on!' Anticipatory emotions and preference for timing in lotteries," *Journal of Economic Behavior and Organization*, 98, 29–40.

Koszegi, B. and M. Rabin (2009): "Reference-Dependent Consumption Plans," *American Economic Review*, 99, 909–936.

Kreps, D. M. and E. L. Porteus (1978): "Temporal Resolution of Uncertainty and Dynamic Choice Theory," *Econometrica*, 46, 185–200.

Loewenstein, G. (1987): "Anticipation and the Valuation of Delayed Consumption," *The Economic Journal*, 97, 666–684.

Masatlioglu, Y., A. Y. Orhun, and C. Raymond (2017): "Intrinsic Information Preferences and Skewness," Working paper.

Menestrel, M. L. (2001): "A Process Approach to the Utility for Gambling," *Theory and Decision*, 50, 249–262.

Mononen, L. (2022): "On Preference for Simplicity and Probability Weighting," Working paper.

Nielsen, K. (2020): "Preferences for the resolution of uncertainty and the timing of information," *Journal of Economic Theory*, 189.

O'Donoghue, T. and C. Sprenger (2018): "Reference-Dependent Preferences," Book section.

Palacios-Huerta, I. (1999): "The Aversion to the Sequential Resolution of Uncertainty," *Journal of Risk and Uncertainty*, 18, 249–269.

Payne, J. W. (1973): "Alternative approaches to decision making under risk: Moments versus risk dimensions," *Psychological Bulletin*, 80, 439–453.

Puri, I. (2022): "Simplicity and Risk," Working paper.

Rubinstein, A. (1988): "Similarity and Decision-making under Risk (Is There a Utility Theory Resolution to the Allais Paradox?)," *Journal of Economic Theory*, 46, 145–153.

Segal, U. (1993): "Order indifference and rank-dependent probabilities," *Journal of Mathematical Economics*, 22, 373–397.

Sicherman, N., G. Loewenstein, D. J. Seppi, and S. P. Utkus (2016): "Financial Attention," *Review of Financial Studies*, 29, 863–897.

Slovic, P. and S. Lichtenstein (1968): "Relative Importance of Probabilities and Payoffs in Risk Taking," *Journal of Experimental Psychology Monograph*, 78, 1–18.

Strzalecki, T. (2013): "Temporal Resolution of Uncertainty and Recursive Models of Ambiguity Aversion," *Econometrica*, 81, 1039–1074.

Thaler, R. H., A. Tversky, D. Kahneman, and A. Schwartz (1997): "The Effect of Myopia and Loss Aversion on Risk Taking: An Experimental Test," *The Quarterly Journal of Economics*, 112, 647–661.

Tversky, A. and I. Simonson (1993): "Context-Dependent Preferences," *Management Science*, 39, 1179–1189.

Weil, P. (1990): "Nonexpected Utility in Macroeconomics," *The Quarterly Journal of Economics*, 105, 29–42.

Zimmermann, F. (2015): "Clumped or Piecewise? Evidence on Preferences for Information," *Management Science*, 61, 740–753.