Modeling attitudes towards uncertainty across attributes, sources and time

Emmanuel Kemel

Modeling attitudes towards uncertainty across attributes, sources and time
Research summary of period 2014-2020, presented for application to "Habilitation à Diriger des Recherches" at University of Rennes

Emmanuel Kemel, GREGHEC (CNRS & HEC Paris)

Soutenue le 12 février 2021

Jury

- Gildas Appere, Université d'Angers, Président
- Olivier L'Haridon, Université de Rennes, Membre
- Meglena Jeleva, Université Paris Nanterre, Rapporteur avant soutenance
- Kirsten Rohde, Erasmus School of Economics, Rapporteur avant soutenance
- Sebastian Ebert, Frankfurt School of Finance and Management, Rapporteur avant soutenance

Abstract

This document presents a summary of my research since my PhD Defense in March 2014. My work mainly consists in empirical investigations of preferences in decisions involving uncertainty and/or time. Most empirical studies on preferences under uncertainty focus on a restricted context where probabilities are known and outcomes are immediately-received monetary outcomes. My research extends the scope of decision contexts by exploring the impact of the type of consequence (the attribute), the source of uncertainty, and the timing of resolution of uncertainty and/or reception of outcomes on attitudes towards uncertainty. The document summarizes my main papers contributing to this research direction. It then presents a critical discussion of this work and proposes directions for future research.

Key words: risk attitudes, time preferences, ambiguity, probability weighting, reference dependence, discrete-choice econometrics, real incentives

Contents

Introduction

Most of the research efforts for studying decision under uncertainty have been dedicated to the analysis of preferences in decisions under risk, and several anomalies in the model of rational choice, expected utility, have been highlighted. So far, the main answers to these anomalies can be summarized into two main insights: probability weighting (Starmer, 2000) and reference dependence (O'Donoghue and Sprenger, 2018), two insights are combined in Prospect Theory (Tversky and Kahneman, 1992). Under genuine risk, the objects of choices are immediately-received monetary payoffs conditional on known and immediately-resolved probabilities. In contrast, real-life decisions feature cases where each of these aspects can be modified: the outcomes can be non-monetary, probabilities can be unknown, and risk resolution and payment of outcomes can be delayed. This manuscript presents a series of empirical investigations that explore the role of attributes, sources and time on preferences towards uncertainty. Even though, for all these contexts, preferences can be, in theory, modeled by Expected Utility (EU) or extensions, empirical evidence suggests that behavioral models accounting for reference dependence and probability weighting better account for empirical patterns. At first glance, this may not be surprising because reference-dependent models allowing for probability weighting are more flexible and are therefore more likely to fit empirical patterns. However, it also appears from the studies that the improved goodness of fit is not due to over-fitting. First, more modern statistical techniques inspired from machine learning suggest that behavioral model also feature better out-of-sample predictions (e.g Kothiyal et al., 2014). Second and more importantly, the behavioral patterns captured by these models can be rationalized, thereby offering not only better predictive or explanatory power, but also better "understanding power" than EU. The studies summarized in this manuscript show that behavioral models are better able to fit, but also understand, the impact of attributes, sources and time on attitudes towards uncertainty, than EU.

The next section reports studies investigating the stability of risk preferences over different types (i.e attributes) of outcomes. Intuition suggests that the impact of attribute type should be mainly captured by the utility function that explicitly measures the perception of outcomes. Studies show that this is indeed the case, however, probability weighting is also impacted by the nature of outcomes. This suggests that the psychological perception of risk itself is *attribute dependent*: it is impacted by the type of outcome at stake. This pattern questions the generally-assumed separability between states (or probabilities) and outcomes.

Section 2 explores the rich domain of uncertainty, as it investigates preferences over situations where outcomes are conditional on events for which no objective probabilities are available. In such situations, events are generated by sources of uncertainty. Under the model of rational choice, events should matter only through their perceived likelihood. In fact, however, empirical patterns reveal that the source generating the events can itself impact preferences. A phenomenon referred to as *source dependence*. The chapter shows how to disentangle beliefs from attitudes in such contexts, and reports evidence of such source dependence.

Section 3 investigates the impact of time on risk preferences. The first two studies investigate the interaction between risk and time, that are often considered as two separate topics in the literature. They implement the discounted expected utility model, the rational model of decisions involving risk and time, and reveal some of its limitations. In Section 3, we also study attitudes towards the timing of resolution of uncertainty. In this context, rational preferences are modeled by a recursive version of EU, proposed by Kreps and Porteus (1978). Building on experimental data, the study shows that a model capturing attitudes towards risk resolution though probability weighting offers a better explanatory and also predictive power.

The last section offers a critical review of the presented studies and draws directions for further research. In particular, it draws direction of theoretical developments for connecting the different probability-weighting functions observed over different attributes, sources and time periods.

Notation, vocabulary and theoretical background

- Modeling preferences towards uncertainty

In decisions under uncertainty, outcomes are dependent on events. The decision maker has no impact on which event will realize, but she can choose which distribution of outcomes over events is preferred. The objects of choices, called *acts* by Savage (1954), but now mainly called *prospects*, are thus distributions of outcomes over events. Like in most empirical studies, those reported in this manuscript focus on binary prospects. We denote $(x, E; y)$ with $x \geq y$ the prospect that gives x if event E obtains and y otherwise. When the probabilities of events are objectively known, a situations called *risk*, events become immaterial and the carriers of uncertainty are probabilities. A risky prospect, also called *lottery*, is denoted $(x, p; y)$, and gives x with a probability p and y with probability $1 - p$. In this document, we consider risk as a specific type of uncertainty, where probabilities are known and the term ambiguity will be used to compare situations of known versus situations of unknown probabilities.

I use notation $\succ$ and $\sim$ to denote strict preference and indifference in the preferences of the decision make (DM). The certainty equivalent CE of a prospect $(x, E; y)$ is the outcome c such that $c \sim (x, E; y)$. The matching probability of a prospect $(x, E; y)$ is the known probability m such that $(x, m; y) \sim (x, E; y)$.

Expected utility (EU) is the benchmark model of rational choice for decisions under uncertainty (Savage, 1954). Under this model, preferences are captured by two components: a strictly increasing utility function u and a probability distribution μ over events. The value assigned to a binary prospect (x, E, y) is

$$\mu(E)u(x) + (1 - \mu(E))u(y). \tag{1}$$

In the case of risk, the value of (x, p, y) is

$$pu(x) + (1 - p)u(y). \tag{2}$$

Despite its normative appeal, this model fails to capture major psychological aspects of decision under uncertainty: probability weighting and (non-neutral) ambiguity attitudes, and reference dependence. Probability weighting refers to the observation that decision-makers do not treat probabilities linearly (Kahneman and Tversky, 1979). Under risk, this bias can be accommodated by a strictly increasing probability-weighting function w mapping $[0, 1]$ to $[0, 1]$ and by assuming that a prospect $(x, p; y)$ is evaluated by

$$w(p)u(x) + (1 - w(p))u(y). \tag{3}$$

This model has also been applied to uncertainty by replacing objectives probabilities p with subjective probabilities $\mu(E)$ (Fox and Tversky, 1998).

Non-neutral ambiguity attitudes, the other well-documented deviation from EU, refers to the observation that decision-makers may exhibit a preference between known and unknown probability distributions over events; in other words, they behave as if they do not assign the same weight to objective and subjective probabilities. In a famous illustration of this behavior, Ellsberg (1961) intuited that people would prefer to bet on an urn with a known composition (i.e. with objective probabilities) than on an urn with an unknown composition (i.e., with subjective probabilities), even if there is no reason to believe that one composition would be more favorable than the other. This behavior can be accommodated by the introduction of a specific weighting function w_a, and by assuming that an ambiguous prospect $(x, E; y)$ is evaluated by

$$w_a(\mu(E))u(x) + (1 - w_a(\mu(E)))u(y). \tag{4}$$

Under this model, ambiguity attitudes are captured by the difference between the weighting functions of subjective probabilities w_a and the weighting function of objective probabilities w. This model accounts for ambiguity aversion while assuming the existence of a unique distribution of subjective probabilities μ. This probabilities are called a-neutral (Baillon et al., 2018) inasmuch as their correspond to the beliefs that would be measured using matching probabilities if the DM was ambiguity neutral.

While improving descriptive power, the weighting function w_a does not allow for ambiguity attitudes to vary depending on the (non-risky) source generating events. To overcome this limitation, Chew and Sagi (2008), followed by Abdellaoui et al. (2011a), developed an approach assuming that the weighting function is different for each source, and hence called this function a *source function*. Using the source function w_S, we evaluate an ambiguous prospect $(x, E; y)$ with event E generated by a source S by

$$w_S(\mu(E))u(x) + (1 - w_S(\mu(E)))u(y). \tag{5}$$

Comparing w_S to w characterizes the ambiguity attitude towards a given source S. The difference between source functions of two distinct sources characterizes *source dependence*: i.e., the fact that ambiguity attitudes differ across sources.[1] The source model preserves the notion of a-neutral probabilities for measuring beliefs, despite source dependence. In this context, probabilistic sophistication is said to be "local" as it holds within sources, but not between sources. The assumption that probabilistic sophistication holds within source, entails that the source is homogenous: all the events from the source feature the same level of ambiguity, and are waited by a same source function w_s.

Other types of behavioral extensions of EU, called *multiple-prior* models (e.g. Gilboa and Schmeidler, 1989, Klibanoff et al., 2005), account for ambiguity attitudes by relaxing the unicity of the distribution of subjective probabilities. However, the flexibility of these multiple priors make these models difficult to elicit, except in specific settings like in the study reported in Section 2.3.

The third well-documented deviation from EU is reference dependence. Whereas EU assumes that the utility applies to final levels of wealth, empirical research has however shown that decision makers behave as if they edited consequences in terms of gains and losses, i.e recode them as positive or negative deviations from a reference point, and exhibit different preferences towards each *domain* (Kahneman and Tversky 1979). In order to account for this behavior, consequences are modeled in terms of deviations from a reference point. In studies involving money, the reference point is the current wealth level of the decision maker (that includes the fixed payoff and endowment offered in the experiment). For studies involving other attributes, such as time consequences, a specific reference point has to be set up. A prospect is a gain (loss) prospect if it involves positive (negative) deviations, i.e. gains (losses), only; it is a mixed prospect if it involves both gains and losses. In order

[1] Several authors have proposed to consider risk as a specific source of uncertainty. Under this convention, ambiguity aversion ($w_a \neq w$) is a specific case of source dependence: a preference between a source with known probabilities over source(s) with unknown probabilities.

to account for reference dependence, the components of the model (e.g. utility and/or probability weighting) are allowed to differ across domains. Reference dependence also applies to inter-temporal preferences, whose models are presented in the next paragraph.

- Modeling inter-temporal preferences

The objects of choice in inter-temporal decisions are streams of outcomes. Like in most empirical studies, the papers mentioned in this manuscript focus on streams involving a maximum of two outcomes, including one received now. A time prospect denoted $(x, t; y)$ gives outcome $y \geq 0$ now and $x \geq 0$ at a future date t. This notation is simplified to (x, t) when $y = 0$.

The benchmark model of rational choice over time is the discounted utility model Samuelson (1937), where a time prospect is valued

$$D(t)u(x) + u(y)$$

where u is a strictly increasing utility function, and $D(t)$ is a discount function that follows $D(t) = e^{-\delta t}$ where $\delta > 0$ is the discount rate.

This model fails to account several patterns of inter-temporal choices. The most documented of them is decreasing impatience. It is indeed common to find patterns of choices such that $(x, t_1) \prec (y, t_2)$ but $(x, t_1 + \tau) \succ (y, t_2 + \tau)$, with $\tau > 0$. When this type of preference is observed only in cases with $t_1 = 0$, it can be accommodated by a the quasi-hyperbolic discount function $D(t) = \gamma e^{-\delta t}$ with $\gamma < 1$ and $D(0) = 1$. Parameter γ captures the present bias, i.e. the fact that decision makers may over-weight the present w.r.t to the future.

When this type of preference holds also in cases where $t_1 > 0$, functions featuring decreasing impatience must be used. Popular specifications with this characteristic include the hyperbolic function $D(t) = (1 + \gamma t)^{-\delta/\gamma}$ (proposed by Loewenstein and Prelec, 1992), and the constant-sensitivity function $D(t) = e^{-(\delta t)^{\gamma}}$ proposed by Ebert and Prelec (2007).

- Modeling preferences in decisions involving both risk and time

We also consider situations that combine risk and time in the following way: probability-contingent outcomes are replaced by streams of outcomes. In addition, the simplest types of streams of outcomes are considered: single future consequences.

Therefore, objects of choices are denoted $(x_t, p, ; y)$ with $x \geq y$ and $t \geq 0$. The rational model for evaluating these objects is discounted expected utility (used by Andersen et al., 2008), that evaluates these objects

$$e^{-\delta t}[pu(x) + (1 - p)u(y)].$$

This model that combines EU and DU, also combines their limitations. It therefore seems natural to augment this model with the elements that have be developed for fixing these limitations, namely probability weighting and non-exponential discounting.

Such a model, that could be called hyperbolic rank-dependent discounted utility would evaluate a prospect with the functional

$$D(t)[w(p)u(x) + (1 - w(p))u(y)].$$

This model is elicited in the study reported in Section 3.2. In Section 3.3 we investigate the stationarity of function u and w.

A recent stream of research has investigated if the utility function revealed from risky choices is the same as the utility function revealed from inter-temporal choices (e.g. Andreoni and Sprenger 2012). If the two functions a different, then it is not obvious which when should be used when risk and time are combined. I am involved in a research study in progress on this topic, but it is not reported in this manuscript.

Considering choices involving risk and time opens the question of the timing of resolution of risk, and its impact on risk preferences. This aspect is addressed in Sect. 4.4. There, we consider choices between objects $(x, p^t; y)$ where outcomes are received at a future time T and the risk, described by a probability p is solved either at $t = 0$ ("now") or at $t = T$ ("later"). The rational-choice model accounting for preferences towards the timing of resolution of risk (Kreps and Porteus 1978) is an extension of EU that allows for different utility functions, depending on the resolution timing. Prospect (x, p^t, y) is evaluated by

$$pu^t(x) + (1 - p)u^t(y)$$

Here again, I explored the empirical performance of an extension of this model that accounts for cause of descriptive limitations of EU: probability weighting.

1 The role of attributes in preferences towards risk and towards time

1.1 Introduction

The theory of decision under risk has been highly influenced by games of chance involving monetary outcomes. The St Petersburg game is probably one of the most famous, as it is told to be a motivation for Bernoulli to propose the EU functional. Other influential games such as the choice between bets involving a common consequence or a common ratio, proposed by Allais (1953), or the choice between bets involving a known or an unknown urn proposed by Ellsberg (1961) also involve monetary consequences. Even though decision science does not restrict its scope to situations involving monetary consequences, the large majority of choice-data collected in experiments that involve money. The main reason for this is certainly convenience as money allows for an easy implementation of real incentives in the lab. Another reason is that, for economists, money is the most general source of utility. Economist are indeed used to converting various types of consequences on the monetary scale. If any type of consequence can be converted into money, then any risky choice can be converted into an equivalent monetary lottery. Under this assumption, studying risk attitude towards money would suffice to explain risk attitudes towards any type of consequence.

This section reports a series of experimental studies that investigate the impact of the type of consequence, the *attribute*, on preferences. The results show that the type of attributes impacts the perception of probabilities and time, thereby questioning the separability of outcomes, probabilities and time periods, assumed by models rational decisions. Behavioral studies have already highlighted two types of violations of this separability: the gain-loss asymmetry, and the magnitude effect, that I now briefly introduce.

It has been shown that preferences depend on the domain of consequences (gains vs losses). Prospect theory (Tversky and Kahneman 1992) accounts for the fact that decision makers exhibit different attitudes towards gains and losses by allowing the gain and loss components of prospects to be evaluated using different functionals. The utility function and the probability weighting function are domain dependent: they differ between gains and losses. Similarly, in inter-temporal choice, empirical investigations also suggest that the discount function can be domain dependent. These patterns suggest that the separation between perception of risk or time on the one hand side, and perception of consequences is not complete. Indeed, the (perceived) domain of the consequence can impact the perception of risk or time. Another example of non-separation between outcomes and risk or time is the magnitude effect. It has been shown that probability weighting is less elevated for high gains than for low gains (Fehr-Duda et al. 2010). The probability weighting function is therefore impacted by the magnitude of consequences. Evidence for a magnitude effect has also been observed in the context of inter-temporal choice. Impatience is generally stronger when stakes are small (e.g. Frederick et al. 2002). Intuition would suggest that the magnitude effect "should" be captured by the utility function. Indeed, the very objective of the utility function is to capture the impact of variations in the magnitude of consequences on preferences. However, empirical studies suggest that magnitude effects cannot be captured solely by the utility function, the magnitude of the consequences also impacts the perception of time or probabilities.

This section contributes to the empirical investigation of these non-separabilities. Considering consequences related to different attributes (e.g. money, time, human lives), it shows that the type of attribute at stakes influences risk and time preferences. This influence is only partially captured by the utility function, and the perception of probabilities or time is also impacted by the type of attribute.

The first study compares risk attitudes towards money and time. The second study takes the perspective of public choice and compares risk attitudes towards saved or lost human lives to attitudes towards their monetary equivalent. The third study deviates from decision under risk and studies inter-temporal discounting. Comparing discounting of money versus discounting of time, the study reveals that the type of consequences also impacts inter-temporal preferences.

1.2 Risking money versus time

This section summarizes the paper: *Abdellaoui, M., & Kemel, E. (2014). Eliciting prospect theory when consequences are measured in time units: "Time is not money". Management Science, 60(7), 1844-1859.*

1.2.1 Motivation

Many real-life decisions involve time monetary risk, i.e. risk whose consequences are expressed in time units. This is particularly the case in health, where the consequences of diseases or treatments can be expressed in terms of duration of a current health state, or life expectancy, in transport where travel time and its reliability is a key dimension of travel or mode choice, or management where the time dedicated to repeated to particular tasks can be impactful for the productivity of organizations. For these reasons, studying attitudes towards time risk, and its difference with genuine (i.e. monetary) risk, is of interest for economic application.

In decision science, time is generally considered in terms of delay for the reception of consequences, in the context of inter-temporal choice. In this regard, considering time as consequence opens a new research direction. For example, recent studies have considered attitudes towards risky delays:

the time a decision maker has to wait until reception/payment of an outcome is conditional on probabilities (Li et al., 2017, Ebert, 2020).

Here, we consider time consequences in terms of saved/free time that can be allocated to any activity, or wasted/lost time that cannot be allocated to any activity.

A priori, time features several specific characteristics that may suggest that attitudes towards time risk differ from attitudes towards monetary risk. Unlike money, time does not need to have a monotonic subjective value or utility. A reader enjoying a 300-page book would not necessarily enjoy it more if the book were longer. For the sake of simplicity however, this study focuses on situations in which the value of time is monotonic. For example, a business traveler who is heading to the airport by car, and who has to answer important e-mails before boarding, will monotonically appreciate any time saved in relation to the expected arrival time. Another key feature of money is fungibility. The fungibility of money may encourage people to take risks more easily, because a loss resulting from a particular decision can be compensated by a gain in another activity, which is less true in the case of time. Leclerc et al. (1995) suggested that the difference in fungibility between time and money can result in more risk aversion when individuals face time losses than when they face monetary losses.

A third difference between time and money lies in the fact that time aggregation is difficult. Unlike money, time cannot be easily saved or stored. Zeckhauser (1973, p. 670) remarked that many activities require a degree of preparation that implies time indivisibility: "the process of preparation may be enjoyable, but surely the payoff to a half-completed painting, manuscript, education is not proportional to the payoff for the whole". This specificity may impact the value of time as "indivisibilities create one source of increasing returns for time allocations".

A last difference between time and money related to the mental accounting of time amounts: durations. A behavioral pattern referred to a duration neglect suggests people they take the end and the peak of the experience into consideration when considering a time laps rather than its duration (Fredrickson and Kahneman, 1993).

From a methodological perspective, time offers the possibility to implement real losses in a laboratory setting, something that is difficult with money (e.g. Etchart-Vincent and l'Haridon, 2011). The objective of the study is to elicit preference towards time risk, under a flexible model that can account for behavioral (and only rational) choice patterns. To this aim, we use Prospect Theory (PT), arguably the most flexible model of risk preferences.

1.2.2 Method

We implemented the method used by Abdellaoui et al. (2011c) and adapted from Abdellaoui et al. (2008) to measure PT components for time and money, in a laboratory experiment. For time, the consequences ranged from -60min to +60min. We also elicited attitudes towards monetary risk as a benchmark on the range [-1200€, 1200€], a range of outcome similar to the values used in Abdellaoui et al. (2008).

We used 5 certainty equivalents to elicit the utility function and 5 additional ones to capture the decision weights assigned to probabilities $\{\frac{1}{6}, \ldots, \frac{5}{6}\}$ in each domain (gains and losses) and attribute (time versus money). For each attribute, two additional indifferences were used to elicit loss aversion. Overall each subject completed a total of 44 certainty equivalents.

A key innovation of the study is that real incentives were implemented for time consequences. To this aim, a specific scenario was considered: subjects where invited to participate to a two-hour experiment for a given fixed payoff. The two-hour duration materialized a reference point from which subject would gain or lose time, by leaving the lab (up to one hour) sooner or later. The first hour of the experiment was dedicated to the measurement of preferences. The second hour (that could be shorten in case of time gain) was dedicated to a filler task. 70 subjects participated to the experiment, and real incentives was implemented on half of them. For the others, choices were hypothetical but used the same framing. Figure 1 illustrates the experimental procedure.

1.2.3 Results

The data are first analyzed in a model-free fashion, focusing on differences of risk attitudes across domain (gains vs losses) and attributes (time versus money). When addressing the impact of attribute on risk attitudes, we observed more risk aversion for monetary prospects than for time prospects in the gain domain. In contrast, when dealing with losses, the subjects exhibited more risk aversion for time than for money.

We then elicited the PT components and compared them. For gains, the utility function exhibited less concavity for time consequences than for monetary ones. The same applied for losses, when the curvature of the utility for time was milder. The most striking differences regard the decision weights assigned to probabilities. Although probabilities are a-dimensional and "should" not be impacted by the type of consequences, we observed that probability distortion was markedly different for time and money. In particular probability for time exhibited more elevation (which is interpreted as more optimism in gains and more pessimism in losses) and less sensitivity to changes in probabilities.

These patterns are illustrated in Figure 2. They appear from a parameter-free assessment of probability weighting (decision weight, DW) and are confirmed when two-parameter specifications

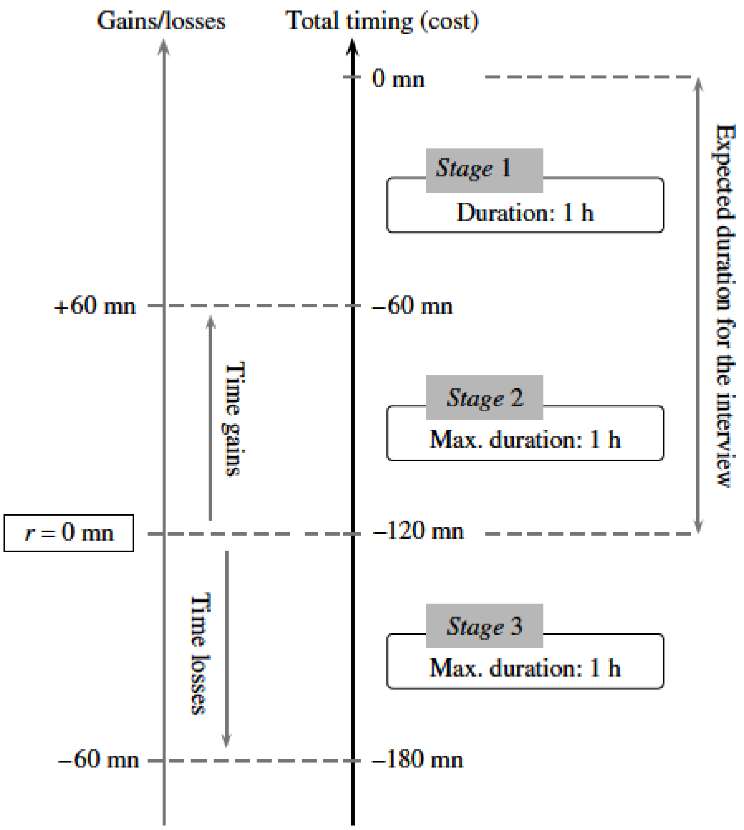

Figure 1: Experimental set up

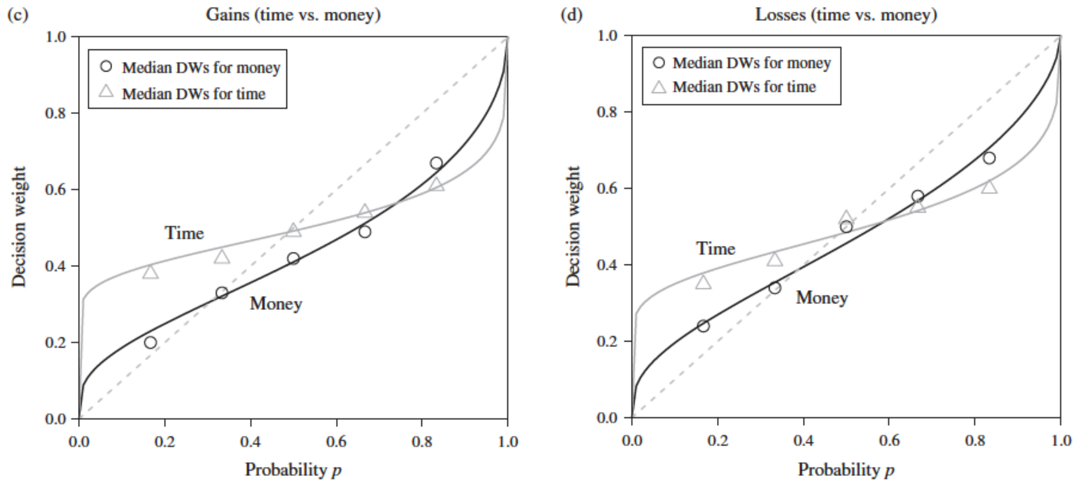

Figure 2: Probability weighting: time versus money

are considered. We also observed that the loss version parameter was larger for money than for time.

The results show that attitudes towards time risk differ from attitudes towards monetary risk. This suggests that specific measurement should be considered for applications to health, transport and management. The patterns generally observed for money cannot be directly used. From a theoretical perspective, a key insight from the analysis is that the type of attribute at stake does not only impact the utility function, but also the perception of probabilities.

1.3 Risking money versus human lives

This section summarizes the paper: Kemel, E., & Paraschiv, C. (2018). Deciding about human lives: an experimental measure of risk attitudes under prospect theory. Social Choice and Welfare, 51(1), 163-192.

1.3.1 Motivation

Many health, security or safety policies aim to reduce mortality and save human lives, and cost-benefit analysis is generally used to select the most efficient policies. However, at the time of the ex ante evaluation, the consequences of the different policies in terms of saved or lost human lives are often uncertain. Policy makers therefore have to make risky decisions involving human lives. In classic cost-benefit evaluations, gains and losses in terms of human lives are generally monetized and aggregated with other aspects of the decision. This approach is based on the implicit assumption that attitudes towards uncertainty are the same for human lives and for money. In this study we question empirically whether people would take the same risks when consequences are expressed in terms of saved/lost human lives, or in monetary terms.

This empirical investigation completes previous studies assessing attitudes towards risk when human lives are involved (e.g. Keeney, 1980, Abrahamsson and Johansson, 2006 and Rheinberger, 2010). Regarding the behavioral-science literature, such decisions echo the famous Asian Disease example proposed by Tversky and Kahneman (1981) to illustrate the impact of framing effects on decisions. Was the context of humans lives decisions used by the authors because the reflexion effect (the difference between gains and losses) is stronger in this context? Among the numerous replications of the Asian Disease problem, Fischhoff (1983) reported that 37% of subjects make different choices with respect to the money version and the human-lives version. Wang (1996) reported more risk seeking in situations involving human lives than in situations involving public good or personal money. Similarly, Fagley and Miller (1997) reported more risk seeking with human lives than with money for both gains and losses. The present study contributes to this literature by comparing risk attitudes towards money and human lives under PT.

The study also contributes to the analysis of the impact of the type of consequences on risk attitudes, and completes the study reported in Section 1.2. Indeed, the previously mentioned study compared two attributes, money and time, that were expressed on different numeric scales and had different values. Here, we measured risk preferences in the range of either [-100,100] human lives, or [-100,100] million euros of public money. We chose this monetary unit because the reference value for a statistical human life in France is around one million euro (Boiteux and Baumstark, 2001). The two attributes were thus compared on intervals of the same numeric range and similar economic values.

1.3.2 Method

Like in the study reported in Section 1.2, we used certainty equivalents to measure preferences. Each subject completed 8 CEs for each domain (gains versus money) and attribute (money versus lives). Two indifferences were also measure for each attribute in order to capture loss aversion.

Two key methodological aspects differ from the experiment of Section 1.2. First, regarding the experimental set, choices were (obviously) hypothetical: 56 subjects were invited to play the role of a social decision maker and make a series of binary choices between public measures defined in terms of probability distributions over possible consequences. The consequences were either gains or losses for society (and not for themselves), expressed in either monetary term or human lives. Subjects were clearly informed that the consequences of the choices did not impact them directly, but impacted the public good. In particular, when the consequences were monetary, decision concerned "public money". Several examples of public decisions involving gains or losses of money or human lives where given in the instructions in order to ensure that subjects understood the meaning of such consequences for society.

The second methodological difference regards the statistical treatment of the data. The raw data analysis also allowed to directly compare risk attitudes across attributes in a model-free approach. Regarding the PT parameters however, a different estimation method was used. We employ a fully parametric version of PT (with parametric specifications of the utility and the probability weighting functions) and estimated the parameters using likelihood maximization of a structural-equation model. In order to account for heterogeneity in preferences, PT parameters where allowed to be randomly distributed across subjects, and the moments the distributions were estimated.

Here is an explanation of why this approach is preferable to individual estimations. Consider individuals $i \in \{1, \ldots, n\}$ with preferences captured by model parameters $\beta_1, \ldots, \beta_n$. A popular approach in the literature consists in running estimations on each individual considered separately. Individual estimates provide values $\hat{\beta}_1, \ldots, \hat{\beta}_n$ with $\hat{\beta}_i = \beta_i + \nu_i$ where $E(\nu_i) = 0$ and $\sqrt{var(\nu_i)}$ are the mean and standard error of individual estimates. The sample variance of individual estimates is therefore $var(\hat{\beta}) = var(\beta + \nu) \neq var(\beta)$. Estimation imprecision of each individual estimate can bias the estimation of parameter variance. Because the sample variance impacts the precision of the sample average estimator, estimation imprecision of individual parameters can also impact inference on parameter means. Random-coefficient models avoid this difficulty by directly estimating sample mean and variance, and their standard errors. These econometric models assume that individual parameters follow a given distribution (e.g. $\beta_i \sim N(\bar{\beta}, \sigma_\beta)$) and estimate parameters (e.g. moments $\bar{\beta}$, and σ_β,) of the distribution. In other words, they estimate moments, whereas usual individual-level analyses rely on moments of estimates. Random-coefficient models can be estimated using simulated likelihood, or Hierarchical Bayes methods, the two approaches being asymptotically equivalent (Train, 2009). In this study, the random coefficient model is estimated using simulated likelihood.

1.3.3 Results

The raw data analysis shows that the fourfold pattern of risk attitudes, the behavioral pattern motivating PT, holds for the two attributes. In the gain domain, no systematic difference in risk attitudes appears between human lives and money. For losses, however, the mean CEs are systematically higher for human lives than for money, suggesting a higher propensity for risk seeking. The indifferences used for measuring loss aversion reveal that that the maximum loss accepted to

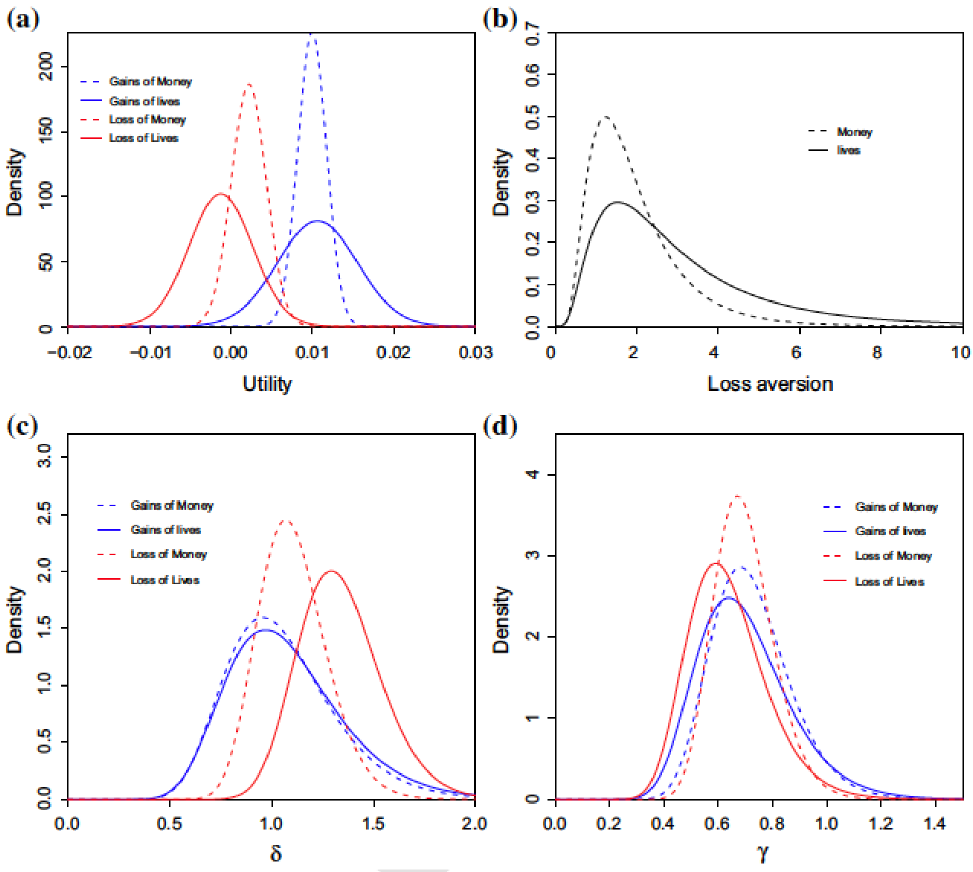

Figure 3: Results of the random coefficient estimations

compensate for a potential gain is much lower for lives than it is for money. This suggests a higher loss aversion towards human lives than money.

Regarding PT parameters, the usual patterns are observed for the two attributes: the utility, modeled by an exponential function, is concave for gains, and linear or close to linear in losses, and probability weighting (modeled by a Prelec specification) is inverse S shaped for gains and losses. The sensitivity γ has a similar distribution in each treatment. Nonetheless, the comparison of the estimated means of others parameters reveals clear differences regarding the treatment of losses. For choices involving only losses, the anti-elevation parameter δ of the probability-weighting function is significantly larger for human lives than for money. This means that the subjects under-weighted worst-case probabilities when human lives were involved. This pattern suggests that the subjects exhibited greater risk seeking for choices involving losses of lives than losses of money, as observed in the raw data analysis. This also refines the raw data result by showing that the difference is driven by the probability weighting function, rather than by the utility function. The second and probably most important difference concerns loss aversion, which is significantly greater for lives than for money. While the average loss aversion for money is close to 2, it is close to 3 for human lives.

Overall, the results show that our subjects did not treat losses of human lives like losses of money: probability weighting in the loss domain is less elevated and loss aversion is greater. Worth noting is that these differences have radically different effects on risk attitudes depending on the situation. If the situation involves losses only, the lower elevation of the weighting function induces greater risk seeking. For decisions involving possible live savings, but also the possibility of losses, the large degree of loss aversion results in a higher degree of risk aversion. Hence, when consequences are expressed in terms of human lives, individuals take more risks in decisions involving losses, but fewer risks in decisions involving both gains and losses. The patterns highlight the value of using a flexible model such as PT to compare these two attributes, since this type of refinement would not be observable with a model that does not distinguish between attitudes towards gains, losses, and mixed situations.

1.4 Discounting money versus time

This section summarizes the paper: Abdellaoui, M., Gutierrez, C., & Kemel, E. (2018). Temporal discounting of gains and losses of time: An experimental investigation. Journal of Risk and Uncertainty, 57(1), 1-28.

1.4.1 Motivation

The studies presented in the previous sections focused on the "attribute dependence" of risk attitudes, i.e. the fact that preferences are impacted by the type of outcomes at stake. In this study, we investigate the attribute-dependence of preferences in inter-temporal choice. In this context, decision makers choose between streams of outcomes. The outcomes are received for sure at a given date. Most models of inter-temporal choice assume a separability between time and outcomes, and a stream $(x_, t; y)$ that gives outcome y now and outcome x at time t is valued

$$D(t)u(x) + u(y)$$

where u is a strictly increasing utility function. A concave utility characterize a preference for smoothing consumption over time and fosters patience. D is a monotonic, generally assumed decreasing, function that captures the perception of time. When decreasing (increasing), decision maker exhibit preference for present (future) consumption which foster patience (impatience). The formula expresses a weighted utility model, whose form reminds models used for preference under risk, except that the decision weights do not sum to 1. Our investigation of attribute dependence of inter-temporal preference therefore extends the scope of this research direction to another type of weighted utility model. Here again, intuition may suggest that the nature of consequences should impact the utility only.

Using time a consequence of inter-temporal choice allows to contribute to a recent stream in the literature: the search for a "good" proxy for consumption. Most empirical research on inter-temporal choice use monetary payoffs, mainly for a matter of convenience. Yet, the model of inter-temporal choice is supposed to apply to consumption, and money may not be the best proxy for consumption. Indeed, the easier transferability of money over time, due to its intrinsic fungibility and the access to credit that decouples money from consumption, may generate specific findings regarding discounting and utility (see Bleichrodt et al., 2016; Cubitt and Read, 2007). A few research works have explored the idea that the time dedicated to a specific task or activity may be a better proxy for consumption than money (e.g., Meissner and Pfeiffer, 2015). Similarly, Augenblick et al. (2015) measured temporal preferences and dynamic inconsistencies for decisions involving the effort dedicated to a specific task. Unlike money, time can neither be stored nor exchanged on a market: an hour of spare (or working) time at a given period has to be enjoyed (or endured) at the same period. The present paper contributes to this research direction.

Regarding possible applications, decisions involving time allocation in a more less distant future are not uncommon in every-day or managerial life. For example, an employee with a standard working contract of 40 hours per week would consider any decrease (increase) of weekly working time as a gain (loss) of time. Regarding time gains, any company or State policy of working time reduction raises the question of how employees prefer to allocate their spare time. Similarly, the decision of when to edit the minutes of a meeting is another example of inter-temporal choice involving losses of time that most professionals have experienced. The task can be completed right after the meeting when the memory is still "fresh", or may be postponed and take longer.

Given the specificities of time, we may expect preference patterns that are non-standard for money. In particular, we may expect negative discounting i.e. a preference for postponing positive consequences. This pattern is considered as irrational for money that should always be received sooner and saved for future consumption. For consumption however, several economist consider such preference as rational: "the pleasurable deferral of a vacation, the speeding up of a dental appointment, the prolonged storage of a bottle of expensive champagne are all instances of this phenomenon." (Loewenstein, 1987). Regarding time gains, the behavior of a person attempting to accomplish a time-consuming task as soon as possible and to "save" free time for the future is also considered rational by Adam Smith: "He is enabled gradually to relax, both in the rigor of his parsimony and in the severity of his application; and he feels with double satisfaction this gradual increase of ease and enjoyment, from having felt before the hardship which attended the want of them".

1.4.2 Method

We developed an experimental set up where subjects could gain or lose time, at different time periods. To this aim, the following scenario was considered. A concrete research-assistantship contract was implemented in order to create a reference point, in the form of a reference duration, against which subjects could evaluate gains and losses of time. In the present study, prospects involved time outcomes received at two different time periods now and at t. Therefore, a reference point had to be set up for each of these periods. Each subject was asked to imagine that a professor had offered her the possibility of receiving a given salary in exchange of several hours of research assistantship. The research assistantship contract required the subject to attend 2 working sessions: the first at now and the second at t. Each session was initially planned to last 4 hours. This reference duration defined the reference point. Gains (losses) of time received at t were defined as the possibility to shorten (extend) the duration of a session scheduled at t. We do not make any assumption about the way the subjects enjoy the time gained: it could be dedicated to leisure, homework or any other activity.

We recruited 101 bachelor and master's students from a Business School. In order to avoid a confounding effect due to a major increase in wealth or to a change in availability, we considered only participants who would still be students a year after the day of the experiment.

We measured preferences using present equivalents, a standard procedure in inter-temporal choice. A specificity of our approach is that we allowed for preference for the future in our measurement. The considered consequences ranging from a loss of 3 hours to a gain of 3 hours. We also measured preferences towards gains of money as a benchmark.

Flexible parametric specifications of the model were considered, with a power function for the utility and the specification proposed by Ebert and Prelec (2007) for the discount. With this specification, parameter δ measures impatience, and parameter γ captures deviations from exponential discounting. This specification can capture both decreasing ($\gamma < 1$) and decreasing ($\gamma > 1$) impatience. A random-coefficient model was estimated, in order to capture the moments of the distribution of preference parameters in our sample.

We compared these moments between gains and losses of time, and gains of money versus gains

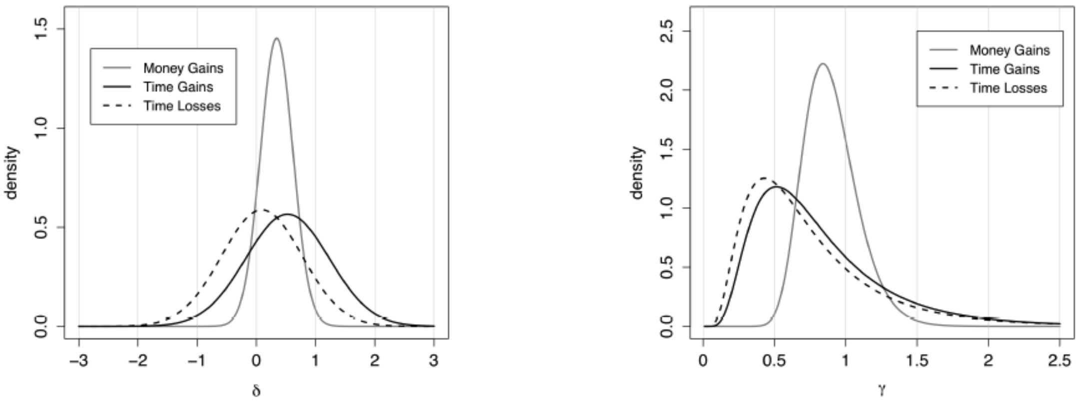

(a) Discount rate Parameter

(b) Delay-sensitivity Parameter

Figure 4: Results of the random coefficient estimations

of time.

1.4.3 Results

Our experiment shows that people do not discount time in the same way they discount money. The results of the random-coefficient estimations are illustrated in Figure 4. Overall, according to model patterns, both the level of impatience (measured by the discount rate parameter δ) and the degree of deviation from constant discounting (measured by the delay sensitivity parameter γ that takes value 1 for exponential discounting) are higher for gains of time than for gains of money. We also observe a much higher heterogeneity in discounting behavior for time than for money. Furthermore, our main finding is the strong asymmetry in discounting behavior between gains and losses of time: the level of impatience is much higher for gains than for losses of time. More specifically, in contrast to the classical view that people would prefer to experience losses further along in the future, we observe that the majority of our subjects prefers to expedite losses of time, i.e. exhibits negative discounting. On the other hand, the large majority of the subjects exhibits positive discounting for gains of time, i.e. they prefer to gain time now than in the future. In other words, the subjects were at the same time eager to work and impatient to rest.

The results show that the attribute dependence of preferences not only concerns risk, but also inter-temporal discounting. In this domain also, using behavioral models accounting for reference dependence allows to fully capture the attribute dependence, that, for example differs between gains and losses.

1.5 Concluding remarks

Rational models of decision under risk and over time, as well as their behavioral extensions assume a separability between attitudes towards outcomes and attitudes towards their carriers (states, probabilities, or time periods). The studies reported in this section question this separability, as the perception of time or probabilities is also impacted by the nature of consequences. The perception of probabilities and time periods is not only domain dependent, and magnitude dependent; it is also attribute dependent. This result suggests that the type of attribute at stake influences the perception of the source of uncertainty: risk is not the same when relating to money or when relating to human lives. Therefore, the notion of source of uncertainty may need to be characterized not only by a random process, but as the combination of a random process and the related consequences. Investigating the impact of sources on preferences revealed the "rich domain of uncertainty" (see the next Section); investigation of attribute dependence reveals that this domain is yet richer: as well as there may be as many attitudes as sources, there may be as many attitudes as types of outcomes. This challenges the parsimony of risk-attitude and time-preference models and questions the external validity of attitudes measured from monetary choices in laboratory experiments.

2 The role of sources of uncertainty on preferences

2.1 Introduction

The large majority of empirical investigations of choice under uncertainty focuses on specific situations where outcomes are conditional of events whose probabilities are objectively known, a situation called risk. Except for games of chance, real-life decisions mainly involve cases where probabilities are unknown. The rational decision model, expected utility, offers a simple bridge between situations of known and unknown probabilities. When no objective probabilities are available decision makers use their own subjective probabilities. This elegant and tractable model been challenged by the decision example proposed by Ellsberg (1961). People prefer betting on a “known” urn containing 50 black balls and 50 red balls to betting on an “unknown” urn containing 100 black and red balls in an unknown proportion. Such preference entails sub-additive probabilities under EU and is interpreted as ambiguity aversion. In terms of behavior, ambiguity attitudes are therefore an additional aspect of preferences that must be accounted form. More technically, (non-neutral) ambiguity attitudes question the possibility to measure beliefs using probabilities, and introduces a new challenge: disentangling beliefs from ambiguity attitudes.

A first solution, Choquet Expected Utility (Schmeidler, 1989), accommodates Ellsberg preferences replacing additive probabilities by sub-additive measures (called capacities) that satisfy monotonicity with respect to event inclusion. While “technically satisfying” these capacities capture both beliefs and attitude with no possibility to disentangle them. Another solution relaxes the unicity of the probability distribution. In these multiple priors models, decision makers are assumed to assign a set of priors to the events (e.g. Gilboa and Schmeidler, 1989, Klibanoff et al., 2005). The spread of priors can be interpreted as a measure of perceived ambiguity and the way multiple priors are integrated characterize ambiguity attitudes. This approach severely complicates the notion of beliefs and its measurement. A third approach assumes local probabilistic sophistication. Additive probabilities holds within source, but not between sources[2] (Chew and Sagi, 2008, Abdellaoui et al., 2011a). This approach separates decision weights into two components: an additive probability measuring beliefs and a source function capturing (ambiguity) attitudes towards the source. In this section, studies build on this model in order to develop a method for measuring beliefs in the form of additive probabilities. The method is then used for investigating the impact of learning on ambiguity attitudes, and for measuring preferences across natural sources.

The first study presents a method for measuring beliefs when decision makers exhibit non-neutral ambiguity attitudes. The second study builds on the observation that the two urns considered by

[2] In this set up a source a mechanism generating events with a homogenous degree of ambiguity.

Ellsberg represent extreme cases: more often than not, decision makers are not given the probabilities related to outcomes, but they are not totally ignorant of them either. They have had some opportunity for learning. It studies the impact of learning through sampling, on the evolution of beliefs and ambiguity attitudes.

The third study extends the investigation of ambiguity attitudes beyond artificial sources à la Ellsberg. Indeed, most studies on ambiguity focus on artificial sources and there are apparent limits to the insights of studying such artificial setting (see Camerer and Weber, 1992; Li et al., 2017 and Baillon et al., 2018). In particular the study show how to measure preferences between two natural sources.

2.2 Disentangling beliefs from attitudes in decision under uncertainty

This section summarizes the paper: Abdellaoui, M., Bleichrodt, H., Kemel, E., & l'Haridon, O. (2020). Measuring beliefs under ambiguity. Operations Research.

2.2.1 Motivation

Under expected utility, the standard model of rational choice under uncertainty, beliefs are represented by a probability measure that can be elicited by observing choices between uncertain bets (Ramsey, 1931). This position was challenged by Ellsberg (1961) famous examples, which suggest that preferences depend on a third dimension, in addition to risk attitude and subjective probability, which he called ambiguity and which reflects the reliability, credibility, and adequacy of the decision maker's information. It raises the question how beliefs can be measured when EU does not hold and decision makers have non-neutral ambiguity attitudes. The purpose of this paper is to present a simple method to achieve this measurement. Like the method proposed by Baillon (2008), our method builds on the construction of equally-likely events has the advantage not to involve chained measurements.

2.2.2 Method

We denote $[a, b]$ the support of a continuous random variable, for which we want to measure the subjective distribution. The method builds on equally likely events. In our context, events are intervals of values that can be taken by the random variable. Two intervals are equally likely if a decision maker is indifferent between betting a positive value x on either of these two events. The method consists in subdividing intervals into equally likely subintervals. Specifically, it proceeds in three tasks, eliciting such subdivisions. First, we measure the median of the distribution by subdividing $[a, b]$ into two equally-likely subintervals. In tasks 2 and 3 we choose a value $c \in [a, b]$

and we subdivide $[a; c]$ and $[c; b]$ into two equally-likely subintervals. The first measurement directly provides the median of the distribution. Measurements 2 and 3 can be combined to derive a statistic related to the dispersion of the distribution. Overall, the method allows to elicit the first two moments of a continuous distribution using three non-chained tasks.

We implement the method to a laboratory experiment were 82 subjects made incentivized choices. The two sources of uncertainty were the temperatures (in degrees Celsius) in Rotterdam and in New York City on January 15, 2013 at 2pm, i.e. one month after the experiment. For these sources, the values $a = -50$ and $b = 50$ were considered. For each source, three series of tasks were completed. The first series consisted in the 3 tasks basing our method. These tasks permit a "deterministic" elicitations of the subjective distributions. The second series consisted in additional measurements of equally-likely events. These additional measurement tasks were used to account for stochastic choice errors in the estimations of the subjective distributions. Concretely, these data were used to feed econometric estimations. The third series of measurement allowed to measure attitudes in addition to beliefs.

An innovation of the study consisted in measuring beliefs assuming a parametric specification for the distribution of subjective probabilities: the flexible Beta distribution was used, and beliefs were measured in terms of parameters of the Beta. The stochastic estimations of beliefs parameters (second series of tasks) and the estimations of attitudes (third series) were treated using individual or random-coefficient likelihood maximization. The objective was to compare the stability and calibration (w.r.t. historical temperature data) of the distributions deriving from the different tasks.

2.2.3 Results

The deterministic method provides subjective distributions that are plausible and well calibrated with historical data (see Figure 5).

At the individual level, a large diversity of patterns is revealed, highlighting the interest of efficient methods allowing for individual-level elicitations. The econometric estimations accounting to stochastic decision errors revealed no substantial differences with the results from the deterministic approach.

The results of the econometric estimations combining beliefs and attitudes produced beliefs that were also consistent with the deterministic ones. This check supports a key aspect of the deterministic method: it allows to measure beliefs that are robust to non-neutral attitudes. In other words, the beliefs that are measured using this method, that "neutralizes" the role of attitudes, are the same as

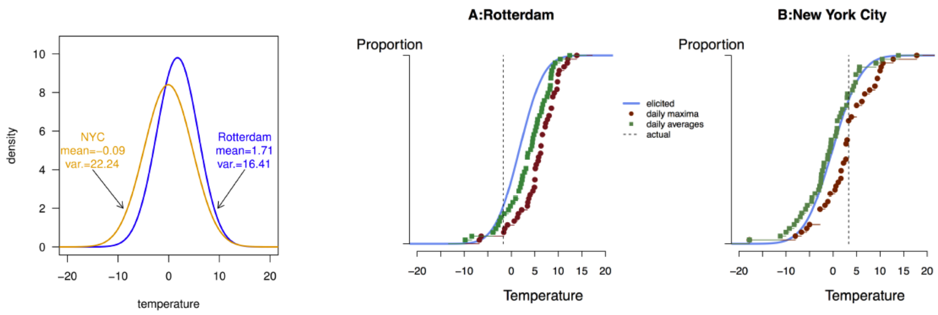

Figure 5: Elicited subjective distributions

the beliefs that are measured when attitudes are explicitly accounted for. Regarding the parameters characterizing attitudes, the expected patterns are observed. The utility function is concave and the median probability-weighting parameters suggest an inverse S-shape probability weighting function for the two sources. A slight difference is captured between the probability weighting of the two sources, that is consistent with source preferences. This aspect of decision under uncertainty is further developed in the following sections. These estimations allow for another important check of the method. The third series of choice-data allows to measure beliefs assuming Expected Utility, the rational model of choice under uncertainty. The analysis shows that the beliefs distributions derived from estimations under this assumption is less well-calibrated with historical data than beliefs that are robust to deviations from EU. The beliefs elicited by our method are therefore more plausible than the beliefs measured under EU. This suggests that omitting deviations from EU does not only entail conceptual considerations, it also has concrete implications for applications based on measured beliefs.

2.3 Learning and the evolution of ambiguity attitudes

This section summarizes the paper: Abdellaoui, M., Hill, B., Kemel, E., & Maafi, H. (2020). The Evolution of Ambiguity Attitudes and Perceptions through Learning. Under revision.

2.3.1 Motivation

In most economic decisions, consequences are conditional on events whose probabilities are not known. Likelihoods are not totally unknown either. Decision makers have some opportunity for learning. In the choice of real-estate, most buyers do not know the precise probability distribution of relevant characteristics, but they have had the opportunity to experience some houses – or data points. Diners, whilst not having complete knowledge of the quality of different restaurants, are not totally ignorant either, having received information from friends, restaurant guides, internet sites and previous experience. These cases of limited information are less ambiguous than Ellsberg's unknown urn, but more ambiguous than the known urn. Ellsberg (1961, p. 659) himself speculated that "If all the information about the events in a set of gambles were in the form of sample-distributions, then ambiguity might be closely related, inversely, to the size of the sample". This study proposes an experimental set up that tests this intuition. We considered a series of urns with unknown compositions, whose composition could be learned by subjects through random sampling. The number of draws available for learning offers an objective and "controllable" measure of available information and ambiguity. Ellsberg's unknown urn would correspond to a situation with no draw. According to the law of large numbers, an urn with an infinitely large sample could be drawn would correspond to the known urn. The set up therefore allows to scan attitudes towards various degrees of ambiguity, between Ellsberg's known and unknown urns. In the main experiment, both beliefs and attitudes towards each urn were measured, allowing to observe and compare the impact of learning on both beliefs and attitudes.

2.3.2 Method

A total of four experiments were run, labeled A1, A2, B1 and B2.

In A1, the experiment measured willingness to bet using matching probabilities for a series of Bernoulli distributions of unknown parameter. Subjects could learn about this parameter through sampling. Three conditions were considered: free sample where subjects could sample as much as they wanted till a maximum of 100, 8 draws and 4 draws. Within each condition, the parameter of the Bernoulli distribution varied in order to scan the probability interval. Risk attitudes were also measured using CEs, and the utility towards urns discovered through 8 draws was also estimated. The experiment involved 71 subjects.

A2 involved 40 subjects and was a robustness check of the patterns captured by A1. It involved matching probability (MP) measurements for treatments "no draws" (Ellsberg unknown urn), 6 draws and 12 draws.

In A1 and A2, the beliefs were not measured and were assumed to follow standard Bayesian updating. In order to assess the robustness of the findings when the assumption of Bayesian updating is relaxed, two additional experiments (B1 and B2) were run, in which beliefs were explicitly measured (using the method from Section 3.2) along with willingness to bet.

The first of these experiments, B1, considered treatments with 0, 5 and 10 draws. In this experiment, sources where represented by physical urns from which subjects made manual draws. B2 investigates attitudes towards large samples, with draws of 100 and 10, 000. Given these large

Studies	Treatments	Choice-based elicitations		Sampling	# Subjects
		Preferences	Beliefs		
A1	$n = 4, 8, \text{free}^{(*)}, \text{risk}$	CEs, MPs	No	Computer	$N = 71$
A2	$n = 0, 6, 12, \text{risk}$	CEs, MPs	No	Computer	$N = 40$
B1	$n = 0, 5, 10$	MPs	Yes	Physical urn	$N = 80$
B2	$n = 100, 10000$	MPs	Yes	Virtual urn	$N = 60$

$^{(*)}$: In the "free" treatment, the subject was allowed to perform as many draws as she wanted, up to a maximum of 100.

Table 1: Description of the four experiments

sample size, sampling was made on a computer. The characteristics of these four experiments are summarized in Table 1

The data are treated under the source model, whose components are estimated using likelihood maximization either at the aggregated or at the individual level. As a robustness check, data are also analyzed under the smooth ambiguity model proposed by Klibanoff et al. (2005), that uses multiple priors to model ambiguity.

2.3.3 Results

Our main finding is that ambiguity attitudes are impacted by learning, with a move towards ambiguity neutrality as sample size increases. Moreover, the impact of this effect on preferences is comparable, if not greater than that of the divergence from Bayesian update. These findings emerge from analysis under both the source model and the smooth ambiguity model. Our investigations also suggest that, though ambiguity attitude becomes more neutral as sample size increases, it does not vanish at very large sample sizes.

Figure 6 plots the regression lines of subjective posterior probabilities (i.e. measured after sampling) $\pi_{\tilde{p}}$ and matching probabilities $m_{\tilde{p}}$ against the corresponding Bayesian posterior probabilities $\pi_{\tilde{p}}^{BU}$ in study B1. and gives an indication of our general findings. Comparing $\pi_{\tilde{p}}$ to $\pi_{\tilde{p}}^{BU}$ measures deviations from Bayesian updating; comparing $m_{\tilde{p}}$ to $\pi_{\tilde{p}}^{BU}$ tells about deviations from Bayesian updating and Bayesian (i.e. ambiguity neutral) choice. The 45° line corresponds to the Bayesian benchmark concerning choice and update. The strong relationship we observe between $m_{\tilde{p}}$, $\pi_{\tilde{p}}$, and $\pi_{\tilde{p}}^{BU}$ confirms that, as one would expect, standard Bayesian updating has by far the strongest impact on both beliefs and preferences under ambiguity. However, $\pi_{\tilde{p}}$ and $\pi_{\tilde{p}}^{BU}$ differ over most of the probability range, indicating a small though significant divergence from Bayesian updating. Moreover, for both sample sizes, $m_{\tilde{p}}$ and $\pi_{\tilde{p}}$ differ considerably over most of the probability range, suggesting non-neutrality to ambiguity. More specifically, the ambiguity premium is negative for medium and likely events, suggesting ambiguity aversion, and positive for unlikely events, indicating ambiguity proneness.

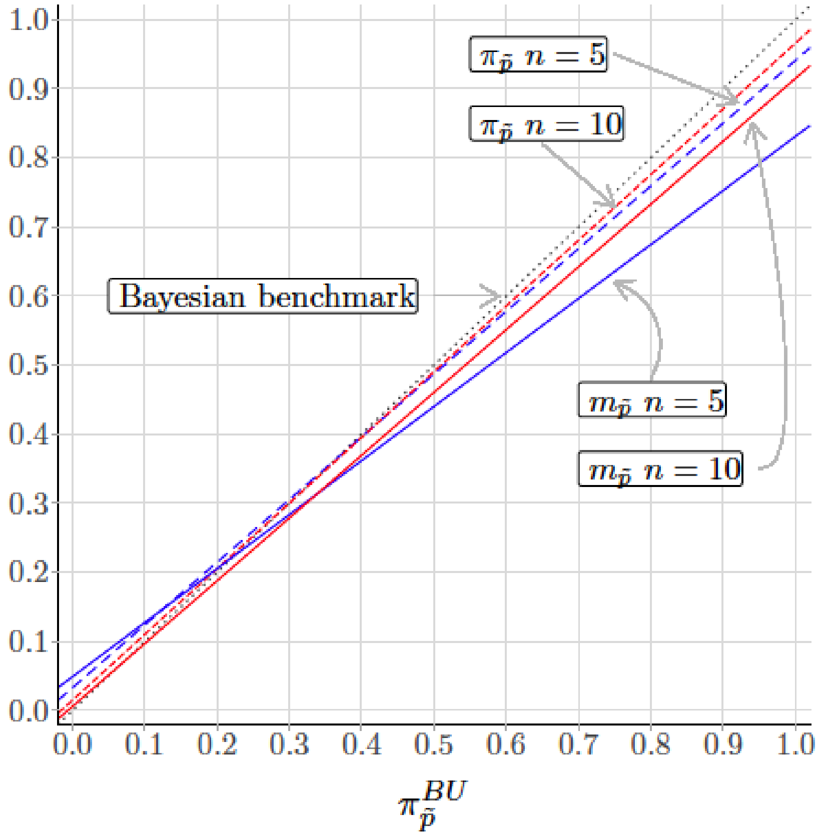

Figure 6: Results from Experiment B1

The Figure also provides an indication of the relative size of the impacts of the various factors that change on learning. Bayesian belief update clearly has the largest impact on matching probabilities. The non-Bayesian factors – divergence from Bayesian update and ambiguity related changes – constitute second-order impacts of learning, but how do they compare to each other? A cautious conclusion that can be drawn from the Figure is that the impact of changes in ambiguity attitude and perception (as indicated by the ambiguity premium) is no less important than that due to divergences from Bayesian update.

Overall the results show that besides beliefs, ambiguity attitudes also vary (though to a lower extent) with learning. Attitudes converge towards ambiguity neutrality, even though non-neutral ambiguity attitudes pertain even for large sample size. The results appear both under the source model and under the smooth ambiguity model, even though the two models build on radically different approaches.

2.4 Measuring natural source dependence

This section summarizes the paper: Gutierrez, C. & Kemel, E. (2020). Measuring Source Dependence.

2.4.1 Motivation

The previous study investigates the impact of learning on ambiguity attitudes, considering urns explored through different sample size. If one follows the literature and considers that a given quantity of information to determine a source of uncertainty (Tversky and Fox, 1995), then each sample size in the study can be thought of as a different source of uncertainty. However, all these sources were created form artificial setting, using urns. I now present a series of experiments that study natural sources and measure how preferences vary from one source to another. According to expected utility, rational decisions depend only on the perceived likelihood of events and the utility of the related consequences. However, the source generating events may also affect decisions: for equal levels of likelihood, decision makers may prefer to bet on one source than another, for instance when they feel more knowledgeable or familiar about it (Fox and Tversky, 1995; Chew et al., 2012). This behavioral pattern, called source dependence, has been investigated mostly using artificial sources, comparing behavior for urns with a known and unknown composition (for a review, see Trautmann and Van De Kuilen, 2015). However, there are apparent limits to the insights of studying such artificial settings, and it is therefore essential to understand how people make decisions for natural sources (see Camerer and Weber, 1992; Li et al., 2017 and Baillon et al., 2018). This study introduces a simple way to measure differences of attitudes across natural sources,

Study	N	Valuation method	Elicitation of beliefs	Sources
Study A	62	CE	EE	Temperature in Paris
				Temperature in a foreign city
Study B	95	MP	EE	Approval rating of D. Trump
				Approval rating of E. Macron
Study C	200	CE	BH	Temperature in Paris
				Temperature in Belgrade
Study D	200	MP/CE[3]	BH	NA: Simulation

Table 2: Summary of the four datasets

through a transformation functions that offers a cardinal measurement of source preferences and a straightforward interpretation. The method is then applied to a series of four datasets.

2.4.2 Method

Consider two natural sources A and B and their source functions w_A and w_B. I introduce the function ϕ_{AB}, such that $w_B = w_A \circ \phi_{AB}$ (i.e., $\phi_{AB} = w_A^{-1} \circ w_B$). The function ϕ_{AB} is strictly increasing, satisfies $\phi_{AB}(0) = 0$ and $\phi_{AB}(1) = 1$, and maps subjective probabilities of events generated by the source B to subjective probabilities of events generated by the source A. Deviations of ϕ_{AB} from identity directly characterize source dependence: A is strictly preferred to B if $\phi_{AB}(x) < x$. In turn, $x - \phi_{AB}(x)$ represents the source-dependence premium, i.e. the decrease in likelihood the decision maker is ready to accept in order to bet on source A instead of source B. Therefore, the transformation function ϕ_{AB} offers a direct measure of source preference of B over A. Inversely, the source preference of A over B is captured by $\phi_{BA} = \phi_{AB}^{-1}$. When A is a risky source (R), we have $w_B = w \circ \phi_{RB}$ and $\phi_{RB} = w^{-1} \circ w_B$. In this case, the transformation function ϕ_{RB} corresponds to the ambiguity function proposed by Dimmock et al. (2016) for capturing ambiguity attitudes. The function ϕ can therefore be considered as a generalization of their approach for capturing source dependence between natural sources.

The paper shows that the function ϕ can be measures easily using structural equation econometrics, from the types of choice data generally used in experiments: either certainty equivalents or matching probabilities. In particular, when using CE, the function can be estimated from a reduced number of choices as compared to previous approaches, since the utility function needs not to be measured.

This approach is illustrated on four data sets: one from a the experiment reported by Abdellaoui et al. (2011a) (study A), two from original experiments (studies B and C) and one from a data recovery study. The characteristics of the datasets used for our implementation are presented in Table 2. Parametric specifications of source dependence function ϕ are estimated using Hierarchical Bayes modelling.

[3]The simulation used in this study applies to both types of data for valuing uncertain prospects, see Section 3.1.4.

2.4.3 Results

In all of the datasets, one source is local (an arguably more familiar to the subjects than the other), and is considered as the reference source.

The main results of the analysis are given in Table 3. Parameter α measures the anti-elevation of function ϕ . It can be interpreted as twice the premium, in terms of winning likelihood, a DM is ready to pay/forgo to bet on the reference (local) source rather than on the other one, for an event of probability 0.5. More details about the parameters used in this study are given in the paper. Parameter β relates to the slope of ϕ and can be interpreted as an elasticity of this premium to changes of likelihood from 0.5. The case where $\alpha = \beta = 0$ refers to a linear ϕ which corresponds to an absence of source dependence.

Estimates $\bar{\alpha}$ and $\bar{\beta}$ (resp. σ_α, σ_β) refer to the mean (resp. standard deviation) of the parameters and capture the modal patterns (resp. heterogeneity).

	Study A (only real incentives)		Study B		Study C	
$\bar{\alpha}$	-0.205	[-0.286; -0.096]	0.353	[0.251; 0.449]	0.051	[0.012; 0.091]
$\bar{\beta}$	0.104	[0.030; 0.176]	0.277	[0.171; 0.377]	0.059	[0.032; 0.085]
σ_α	0.181	[0.115; 0.279]	0.319	[0.255; 0.396]	0.229	[0.200; 0.264]
σ_β	0.137	[0.086; 0.198]	0.360	[0.294; 0.436]	0.125	[0.101; 0.164]
LL	-982.404		-3313.293		-3449.682	

Table 3: Summary of HB estimations - Studies A, B and C

All the studies reveal significant average deviations of α or β from 0, thereby providing clear evidence of source dependence of attitudes towards natural sources. The average parameter β is larger than 0, entailing a preference for the local to the foreign source, a pattern that is constant with home bias. We note also that average parameter α differs from 0 is each study, although in various directions and magnitude. This suggests that source preferences are also likelihood dependent: the magnitude and even possible the sign of the source premium vary with the likelihood.

Parameters σ_α and σ_β are generally larger or equal of the mean value, suggesting over dispersion. This suggests a large degree of heterogeneity in source dependence. Beyond average patterns, effects of source dependence can be important, although in heterogenous magnitude and even direction. This reminds that only focusing on modal patterns may mask individual-level effects that cancel each other when aggregated.

2.5 Concluding remarks

The source model, that assumes local probabilistic sophistication allows to preserve the notion of "well behaved" subjective probabilities within source, and captures between-source violations of probabilistic sophistication by source-specific probability-weighting functions. This model offers two advantages for the analysis of choice under uncertainty. First it allows for new approaches for capturing beliefs. The study reported in Section 2.2 contributes to this direction. Measuring beliefs that are disentangled from attitudes can be useful per se, in order for application or for measuring how beliefs vary in light of new information. They also allows to identify source preferences. The identification and analysis of source preferences is the second advantage. Prior studies highlighted the rich domain of uncertainty, i.e. the heterogeneity of preferences across sources. However, little is known about the reasons why some source generate different attitudes than others. Efficient and tractable methods for capturing source preferences are necessary for going further and investigating the causes and consequences of these source preferences.

3 Preferences in decisions combining risk and time

3.1 Introduction

Risk and time are two key dimensions of economic decisions. We already covered several aspects of decision under risk in the previous sections. Regarding time, decisions involving the planning of consumption over time are frequent, and have important consequences for individuals and economic systems. In particular, for individuals, they relate to saving for retirement and consumption over life time; for governments, inter-temporal preferences are involved in the question of debt management.

Risk and time are generally considered as separate topics. The objects of choice under risk are lotteries, i.e. probability contingent outcomes, and preferences over these objects are modeled under expected utility or behavioral extensions (such at RDU or PT). The objects of choice over time are streams of outcomes, i.e. series of dated outcomes, and preferences over these objects are modeled by the discounted utility model (Samuelson, 1937) or extensions (involving non-exponential discounting for example). However, more often than not, risk and time interact. The future is intrinsically risky, and outcomes of risky decisions materialize in a more or less distant future. In contexts involving both risk and time, the objects of choices can be modeled as probability contingent streams of outcomes. Preferences over such objects are modeled by discounted expected utility. This simple model can be used for applications. For example, while the discounted utility model is generally used assuming a linear utility, Andersen et al. (2008) proposed to use the utility function derived from risky choices. They show that accounting for the curvature of the utility substantially changed the estimated discount rate.

A recent stream of research has questioned the separation between risk and time preference entailed by DEU, showing for example that the presence of risk impacts time preferences (Andreoni and Sprenger, 2012). In this section, we report a series of studies that question the empirical validity of DEU. We first report an experiment where time preferences are measured in a risky environment. It shows that commonly observed biases such as non-exponential discounting and probability weighting hold in this context. The second study reports an elicitation of risk preferences when all possible outcomes are received in a same future date. It shows that these preferences differ from genuine risk (where outcomes are not delayed), and that this pattern cannot be taken into account by DEU. The third subsection addresses an aspect of risk and time that is not covered by DEU: attitudes towards the timing of resolution of uncertainty. When outcomes are received in the future, the resolution of risk can also take place in (a closer) future. The timing of resolution of risk introduces a new aspect of preferences: attitudes towards information. Do people always prefer to receive information regarding the outcome of their risky choice as soon as possible? The study proposes an efficient method for eliciting attitudes towards the timing of resolution of risk, and compares several models in terms of goodness of fit and prediction accuracy.

3.2 Measuring inter-temporal discounting under risk

This section summarizes the paper: Abdellaoui, M., Kemel, E., Panin, A., & Vieider, F. M. (2019). Measuring time and risk preferences in an integrated framework. Games and Economic Behavior, 115, 459-469.

3.2.1 Motivation

The study investigates the properties of time discounting under risk, using an extension of the method popularized by Holt and Laury (2002). Doing so, it provides three contributions to the literature.

Firstly, it contributes to the investigation of time preferences under risk. It has been suggested that deviations from the standard model of inter-temporal decision making, discounted utility with an exponentially decreasing discount function (DU; Samuelson, 1937), may be largely or entirely due to elicitation methods positing certainty of future outcomes (Keren and Roelofsma, 1995;Weber and Chapman, 2005; Halevy, 2008; Gerber and Rohde, 2010; Epper et al., 2011). According to this account, (quasi-) hyperbolic preferences (Laibson, 1997; Rohde, 2010) are imputable to the absence of risk in the present, while risk is inherent in any future outcomes. A dislike of risk would then result in a preference for immediate outcomes over future ones, regardless of a respondent's true underlying discount rate. Under this assumption, deviations from exponential discounting should disappear when time preferences are measured under risk. The study tests this hypothesis.

Secondly, estimating probability weighting in addition to utility curvature further allows to examine the effect of the model adopted under risk on the estimated discount function. The study shows that accounting for nonlinear probability weighting indirectly impacts the measurement of discount rates, through the role of the utility. In the presence of pessimism in the probability weighting function, utility obtained under EU will be excessively concave (Wakker, 1994). Correcting for probability weighting will resolve this issue, thus resulting in reduced concavity in utility (Abdellaoui et al., 2008). This issue will also influence estimates of time discounting.

Thirdly, the study shows how to use the popular Holt-and-Laury method to elicit probability weighting jointly with utility curvature. The methodological approach provides a robustness check for the factors driving inverse-S-shaped probability weighting. Indeed, when elicited from the usual methods, such as certainty equivalents, this shape can be due to a tendency of subjects to shift in the middle of the choice list. Under the Holt-and-Laury method, such choice pattern would produce the opposite shape. Therefore, observing an inverse S-shaped probability weighting under this method suggests that this behavioral pattern is a key aspect of preferences and not an artifact from the measurement method.

3.2.2 Method

The usual Holt-and-Laury method fixes four outcomes $x_r > x_s > y_s > y_r > 0$. Then, for probabilities $p \in \{0.1, \dots, 0.9\}$, its observes choices between the "risky lottery" (x_r, p, y_r) and a "safe lottery" (x_s, p, y_s). Monotonicity imposes that there exists a single shifting probability $p^\star$such that $(x_r, p^\star, y_r) \prec (x_s, p^\star, y_s)$ and $(x_r, p^\star + 0.1, y_r) \succ (x_s, p^\star + 0.1, y_s)$. Under EU, the value $p^\star$ gives information about the curvature of the utility function. A single set of values x_r, x_s, y_s, y_r are usually considered, and they are chosen such that and expected value maximizer would have a shifting probability close to 0.3.

This approach is extended in two directions. First, several series of values x_r, x_s, y_s, y_r are considered, including cases where the shifting probability of an EV maximizer would be close to 0.1 or 0.8. This aspect of the design allows to elicit probability weighting jointly with utility curvature.

Second, the experiment introduces a risk versus time tradeoff in the choices by considering cases where the outcomes of the safe lottery are delayed. This aspect of the design allows to capture time preferences in a risky environment.

A total of 100 subjects were recruited at the laboratory of the Technical University in Berlin, Germany for a 45min individual experiment, run in 20 small group sessions. Outcomes ranged from 0 to 500 euros and subjects had the possibility to have one of their choice played for real in addition to their flat payment. Payment were processed by bank transfer and a particular attention was payed to the credibility of the real incentives mechanism. Each subject completed a total of 42 distinct choice lists.

The choice data are analyzed in a model free fashion that provides model-free evidence of non-constant discounting. They are then analyzed using discrete choice econometrics. Parametric specifications of several models are considered, and parameters are estimated either at the aggregated or at the individual level using likelihood maximization.

3.2.3 Results

The main results of the study are synthesized in Table 4. The benchmark is DEU. In Discounted RDU, probability weighting is allowed. The last two columns report the results of models that relax exponential discounting. In Quasi-Hypeorbolic RDU, present bias is allowed, and Hyperbolic RDU assumes and hyperbolic discount function.

parameter	DEU	DRDU	QHRDU	HRDU
ρ (utility curvature)	0.273 (0.268, 0.279)	0.512 (0.499, 0.526)	0.517 (0.503, 0.531)	0.514 (0.5, 0.528)
r (discount rate)	0.059 (0.056, 0.061)	0.141 (0.135, 0.148)	0.111 (0.103, 0.119)	0.239 (0.211, 0.268)
γ (prob. sensitivity)		0.675 (0.655, 0.694)	0.672 (0.652, 0.691)	0.674 (0.654, 0.693)
η (prob. pessimism)		1.405 (1.364, 1.447)	1.42 (1.378, 1.462)	1.411 (1.369, 1.452)
β (<1: present bias)			0.972 (0.967, 0.977)	
ζ (hyperbolicity)				1.788 (1.201, 2.376)
max LL	-37348.93	-36130.75	-36073.05	-36066.97

95% confidence intervals in parentheses below the estimates.

Table 4: Parameter estimates of structural models

We can see that relaxing the assumptions of probability weighting and exponential discounting both improve the goodness of fit. Significant probability weighting is observed, and the probability weighting function is inverse S-shaped. This shows that the commonly observed pattern of probability weighting can be recovered using the Holt-and-Laury measurement method. Omitting probability weighting distorts the measurement of the utility which, in turn impacts the estimations of the discount rate. Estimating a DEU model with constant discounting and linear probabilities, we estimate a low yearly discount rate of around 6%. Once we allow for nonlinear probability weighting, however, the estimated discount rate more than doubles to 14%. Eventually, models allowing for non-constant discount rate offer a better goodness of fit. This suggests that deviations from exponential discounting are not due to the fact that previous measurement mainly focuses in cases when future outcomes where obtained for sure. These anomalies for inter-temporal preferences are also observed in the context of risk.

3.3 Risk attitudes when consequences are received in the future

This section summarizes the paper: Kemel, E. & Paraschiv, C. Risking the Future? Measuring risk attitudes towards future consequences.

3.3.1 Motivation

The classical literature investigating decision making under risk generally assumes "immediate" outcomes: the decision making process ends by the decision maker experiencing the consequences arising from the choice. However, decision-making under risk in everyday life does not generally corresponds to this theoretical setting of immediate consequences. Instead, real situations of risky decisions often involve a delay between the moment when the decision is made and the moment when the outcomes are received by the decision maker. For example, during elections, citizen vote for a political program whose consequences will realize the future. It was for example the case of Brexit that was announced to take start a couple of years after the referendum if adopted. Would British citizen have made the same choice if they knew the program would be implemented right after the referendum? Delayed outcomes are also common in the health domain. Risky sexual behavior can result in diseases which, albeit contracted/diagnosed now, have consequences that will appear in the future (e.g. cancer). The delay between the (risk taking) decision and the reception of the consequences is also important for deterrence. Law offenders may have different perception of the risk of sanction if the fines are received long after the reckless behavior. Overall, the delay in the materialization of the consequences resulting from the decision process may have an important role in explaining attitudes toward risk in real-life decisions.

We are not aware of any study that investigates precisely the impact of delayed outcomes on the degree of risk aversion of the decision-maker. The closest study to our was run by Abdellaoui et al. (2011b). The authors considered lotteries that were either solved and payed now or solved and payed later, and indeed observed more risk tolerance in the second context. However, by manipulating the timing of the resolution **and** the type of payment of outcomes, their experimental design cannot isolate the sole effect postponing consequences. The present studies pursues this objectives by considering lotteries that are always solved now, but payed either now of later.

3.3.2 Method

We focus on choice objects of the type $(x_t, p; y_t)$ with $t \in \{0, 1\}$. The uncertainty is always solved at time 0 ("now") and the outcomes are both received either at time 0 ("now") or at time 1 ("one year from now"). For each treatment (i.e payment now or later) we measured preferences under rank-dependent utility (RDU), using certainty equivalents. The CEs was received at the same time as the risky consequences. With this, our design neutralizes the role of discounting. 70 subjects participated to a one-hour experiment in the form of individual computer-assisted interviews. Each subject completed 11 CEs for each treatment, as well as other choice tasked used in another project.

Like in the studies reported in Sections 1.3 and 3.2, parametric specification for the RDU model were considered, and where estimated using likelihood maximization at the aggregated or at the individual level. In this study the objective is to compare the RDU parameters across treatments.

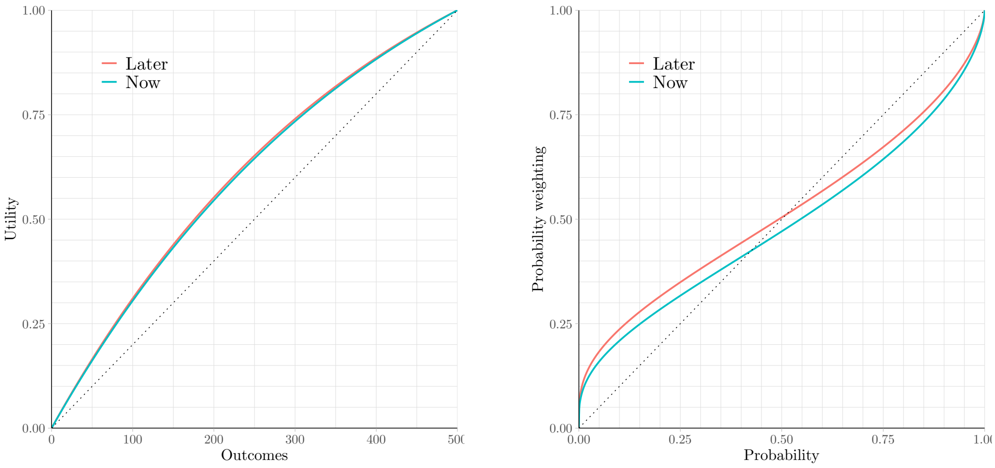

Figure 7: Estimated mean patterns

3.3.3 Results

Raw data analysis shows that the subjects exhibited higher risk tolerance in the treatment where outcomes where delayed. Under RDU risk attitudes derive from the combination of both utility and probability weighting. The econometric estimations are used to identify which of these components are impacted by the treatment. The main results of the study are reported in Table 5, which report aggregated-level and individual-level parameter estimates for each treatment.

	Aggregate		Individual	
	Now	Later	Now	Later
Utility	1.197	1.257	1.096	1.233
	(0.097)	(0.131)	[0.433, 2.190]	[0.631, 2.102]
Elevation	0.941	0.859	0.978	0.865
	(0.036)	(0.035)	[0.781, 1.144]	[0.719, 1.073]
Sensitivity	0.609	0.621	0.645	0.671
	(0.025)	(0.025)	[0.496, 0.836]	[0.546, 0.877]
LL	-5878.904		-4568.743	

Table 5: Estimations with Prelec specification

We can see that utility parameter are similar from one treatment to another, as well as the sensitivity parameter of the probability weighting function. The only difference regards the elevation parameter according to which the probability-weighting function is more elevated (contributing to higher risk seeking) when consequences are delayed. These patterns are confirmed by statistical tests and are illustrated in Figure 7.

Overall, the study replicates the findings reported by Abdellaoui et al. (2011b) with a control for the timing of resolution of uncertainty: subjects exhibit more risk tolerance when the reception of outcomes are delayed. According to the econometric analysis this effect is captured by the probability weighting function that is more elevated in this case.

3.4 Attitudes towards the timing of resolution of risk

3.4.1 Motivation

Many real-world choices involve the resolution of uncertainty over time. Examples include such economically important decisions as consumption, savings, investment, portfolio management, and production. Temporal resolution of uncertainty also plays a role in most medical decisions such as when patients undergo genetic tests to determine the likelihood of getting a disease in the future. In all of these cases, we expect that the decision maker is not indifferent to temporal resolution of uncertainty because (s)he assigns a value to informative signals about it. The value of information is instrumental when it can be used to take action; it is intrinsic when psychological in nature and associated with feelings of anxiety or hopefulness.

The benchmark model of attitudes towards resolution timing has been proposed by Kreps and Porteus (1978, hereafter KP). The model is a recursive extension of EU where the utility U_0 of a lottery solved now is allowed to differ from the utility of a lottery solved later U_T. The function φ mapping the two utility functions, such that $U_0 = \varphi \circ U_T$ captures attitudes towards resolutions timing. Preference for an early (late) resolution is associated with higher (lower) CEs for lotteries solved now, hence a utility function U_0 less (more) concave than U_T and a convex (concave) φ.

Prior elicitations of this model required the measurement of both U_0 and U_T in order to capture φ. We propose a more direct measurement methods for measuring φ with no need to measure U_0 nor U_T. The method thus requires fewer choices and avoids the risk of error propagation due to the measurement of utility functions. In this study, the choice task also allows for model-free characterization of attitudes towards resolution timing (as well as CEs allow for model-free characterization of risk attitudes). It is therefore possible to use the choice data to test alternative models, such as models that allows for probability weighting, a behavioral aspect that is omitted by EU and KP. In particular, the study considers a model proposed by Wu (1999) that introduces probability weighting functions that may differ depending on the resolution time. This model captures attitudes towards resolution timing by the differences of probability weighting. Eventually, it also considers the general set up proposed by Epstein (2008) that allows both utility and probability weighting to vary depending on resolution time.

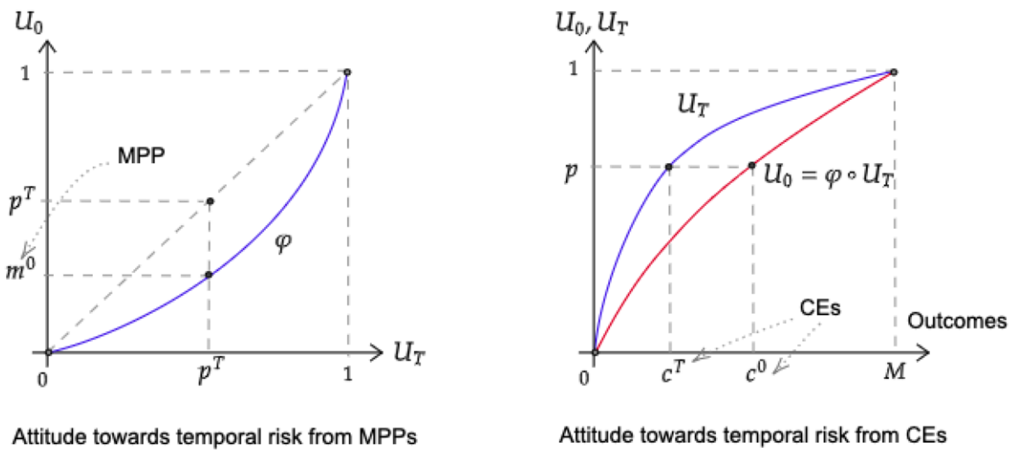

Figure 8: Illustration of the measurement method

3.4.2 Method

The method for characterizing attitudes towards resolution timing involves choices between objects of type $(X, p^t; x)$ with consequences $X \geq x$ received in the future, at time T and risk, described by probability p^t solved either now $(t = 0)$ or in the future $t = T$. The matching present probability m^0 such that $(X, m^0; x) \sim (X, p^T; x)$ is measured. The DM exhibits preference for early (late) resolution of risk if $m^0 < p^t$ $(m^0 > p^t)$ and neutrality otherwise. Under KP, $m^0 = \varphi(p^T)$. The method allows for a direct measurement of φ. When a one parameter specification of φ is considered, the measurement of a single indifference is required. The method is illustrated in Figure 8. In particular it compares the elicitation of φ from MPPs (left hand side panel) and from CEs (right hand side panel).

Two experiments were run in order to implement the method and, more generally investigate attitudes towards resolution timing. In experiment A, 70 students participated to one-hour computer-assisted interviews. A real incentive system was implemented for half of the group. The experiment measured attitudes towards risk solved now with consequences payed at T, one year from now, using CEs. It also implemented our method, measuring matching present probabilities for various levels of probabilities, solved in either 3, 6, 9 or 12 months from now.

In experiment B, we measured attitudes towards risk with consequences payed in one year from now $(t = T)$ but solved either at $t = 0$ or at $t = T$, using CEs. These choices allow for the assessment of attitudes towards resolution timing through the comparison of CEs, like in prior studies. They also allow to estimate RDU in each context, allowing for the comparison of probability weighting and utility across resolution timings. The experiment also measured matching present probabilities for different levels of probabilities. These data were used for two complementary analysis. First I could compare the φ estimated from CEs, to the φ estimated from MPPs. Second, I used these data as "out-of-sample" data for the comparison of prediction accuracy of the models estimated using CEs. This approach, inspired from machine learning, aims at avoiding over-fitting in model comparison/selection.

3.4.3 Results

Results from Experiments A and B show that preference for early resolution clearly prevails among our subjects. However, this preference is less pronounced for small probabilities, suggesting that hopefulness may partially counterbalance anxiety for rare wining events. Analyses under KP, these preference are materialized by a convex function φ.

Data from experiment B allows further comparisons of estimations methods and models. First, it shows that preferences measured using CEs and preferences MPPs are consistent. Since, the later method is more direct, it may be preferred. Regarding model comparisons, models allowing for probability weighting perform best. This is not surprising. Nevertheless, we observe that the bias induced by the omission of probability weighting impacts the φ when measured using CEs. A gap is indeed observed between the φ measured from CE and the φ measured from MPPs. This gap is reduced when probability weighting is accounted for. The model proposed by Wu (1999) captures attitudes towards timing of resolution through a probability weighting function that varies with the timing of resolution. Under this model the weighting function is less elevated when risk is solved later. Preference of early resolution of risk is then captured by pessimism towards probabilities solved in the future. Testing the flexible model proposed by Epstein (2008) we observe that the weighting function varies with resolution timing whereas utility remains almost the same. This suggests that the largest part of attitudes towards resolution timing is captured by probability weighting. These patterns are illustrated in Figure 9.

Eventually, when comparing models capturing resolution timing by different utility function, to models capturing resolution timing by different probability weighting functions, we observe that the later offer both a better goodness of fit, and a better prediction accuracy on out-of-sample data.

3.5 Concluding remarks

The studies reported in this section highlight the richness of interactions between risk and time. The first study shows that probability weighting matters when modeling time preferences under risk. The second study shows that, besides discounting, the timing of reception of outcomes impacts risk

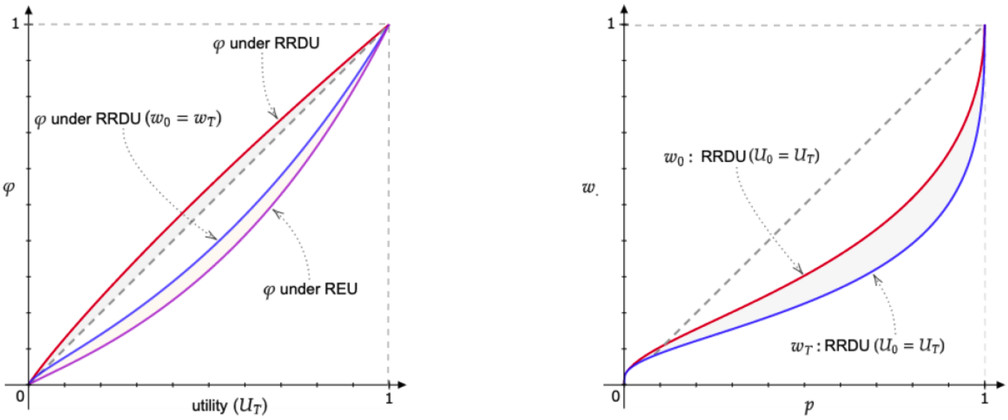

(a) Transformation function (b) Probability weighting

Figure 9: Illustration of the measurement method

attitudes, that this effect is captured by probability weighting. The third study investigates attitudes towards the timing of resolution of risk and, among other things, shows that these attitudes are better captured by probability-weighting functions that by utility functions. These results suggest that the probability-weighting function is a flexible component of risk and time models, that can account for many aspects of time preferences. It seems that probability weighting functions can capture the diversity of situations involving time, as well as they can capture the diversity of situations involving uncertainty (cf Section 2). Nevertheless, like in the case of uncertainty, models using probability weighting offer the flexibility to fit data ex-post, but they suffer from a lack of parsimony: are there as many probability weighting functions as time periods?

4 General discussion

This section proposes a critical discussion of the studies presented in this manuscript, then, it formulates directions for further research. In particular, the first part will address the limitations of the previously-mentioned studies. These limitations mainly consist in restraining assumptions, related to the objects of choice that are considered, the data collection method or modeling approach. The discussion explains the reasons why these assumptions have been made, what are the associated drawbacks and what can be alternative assumptions. Exploring alternative assumptions will open a first series directions for methodological developments. The second part presents a series of directions for further research, mainly in terms of model development for connecting the components of models across different attributes, sources and time periods.

4.1 Limitations

4.1.1 Observing binary choices between simple alternatives

All the choices considered on this manuscript are arguably the simplest choices possible. Decision makers have to express a preference between two alternatives. Additionally the alternatives are also designed to be as simple as possible. Prospects are binary, and involve a non-zero minimum outcome only when needed for model identification[4]. Similarly, for inter-temporal choices, the simplest streams of outcomes are considered. They involve only one future outcome, sometimes in addition to a present outcome, for the needs of model identification[5]. Of course, real-life decisions involve much more general and complex choice situations. The decision maker must choose between more than two alternatives and each alternative may be complex. Uncertainty often involves more than two states, and inter-temporal decisions such as investments involve long series of future payments or cash flows.

Therefore, not only my studies clearly focus on restricted choice sets, that do not apprehend the complexity of real-life decisions, but a specific effort has been made in order to measure preferences (under the models considered) from the simplest choice possible. Here are the reasons why.

The main reason is that our objective is explore the limitations of rational models, and in particular, situations where they are systematically violated. Given the cognitive limitations of human decision makers, it is likely violation of the rules of rationality (such as transitivity or monotonicity) will arise when dealing with complex objects. However, if violations are observed on simple choices,

[4]choices between binary prospects with only one non-zero outcome do not allow to separately identify the utility function, and the weighting function.

[5]choices between sooner-smaller vs larger-later outcomes do not allow to separately identify the utility function, and the discount function.

they cannot be interpreted as deriving from the treatment of complexity. For example, if a DM violates first-order stochastic dominance when considering two complex lotteries, the analyst may consider that a possible reason of the violation is that complex lotteries have been been correctly processes or understood by the decision maker. If a violation of first order stochastic dominance is observed from simple binary lotteries, then the violation will be explained by the misunderstanding of risk itself, which is more specific.

A second and more pragmatic reason is that elicitation of the models considered in this manuscript require a large quantity of data. Concretely, subjects participating to laboratory experiments have to make a large number of choices. The assumption that is made is that subjects have limited attention resources and that considering cognitively-simple choices allows them to make more careful choices.

In this respect, I followed in these studies the objectives that is often pursued in survey design, and that consists in developing questions that are as simple and explicit as possible.

A more fundamental motivation for considering simple choices relates the application of decision science: decision analysis. One of the objectives of decision analysis is to help decision makers coping with complex decisions. To this aim, a general procedure consists in separating the complex decision into several simple decisions that are less cognitively difficult (see examples in Keeney et al., 1993). Simple choices therefore play a key role in decision analysis. They are present both during the process steps of the decision analysis (for measuring preferences and replacing elements of the decision tree by certainty equivalents), and during the final stage. Indeed, the process consists in simplifying the decision context, the final decision may take the form of a simple choice. Overall, from the perspective of decision analysis, simple choices are not "artifactual simplifications", they are a pillar of the discipline. It is therefore of interest to observe preferences over such choices. In particular, a key aspect of decision analysis is to help decision-makers follow the rules of rationality in complex choices. But does it make sense to guide the decision maker in this direction if (s)he seems to deny rules of rationality even in simple choices? This point is made by Baucells and Katsikopoulos (2010) as a motivation for the development of descriptive models of decision making.

Eventually, a theoretical motivation for considering binary prospects in the case of decision under uncertainty is that, in this context, most of the available models take a similar functional, called biseparable preferences (Ghirardato and Marinacci, 2001). Measuring preferences in this context allows to remain general and agnostic about which specific model to employ. Inversely, capturing deviations from the biseparable preference allows to capture limitations that concern many models.

Obviously, the decision to work on simple choices has a cost that I now discuss. I can see three main types of limitations. The first one is that we intentionally ovoid several aspects of the behaviors that are specific to complex environments. For example, I did not try to understand why people sometimes violate transitivity or monotonicity. Instead, I tried to develop choice environment that, by their design and simplicity permit to reduce these violations. My research consists in developing simple choices for "extended" decision contexts (involving different attributes, sources and time periods). A "dual" and complementary approach would consist in studying complex choices in standard contexts (e.g. Birnbaum and Schmidt, 2008).

Another limitation of simple choices is that they to not allow to investigate contexts where models make specific predictions. In the context of decision under risk, an illustration is provided by the notion of rank dependence. Rank dependence entails that the decision weight assigned to probabilities depends on the rank of the related consequence. Rank dependence derives from the fact that the transformation of probabilities applies to cumulative distributions rather than to densities, thereby preserving first order statistical dominance despite probability weighting. It is sometimes considered as a "technical trick" that avoids that the model predicts violations of FOSD. It has also been considered as a plausible behavioral, and psychologically grounded, phenomenon (Diecidue and Wakker, 2001). The role of rank dependence only arises when prospects feature more than two outcomes. Bernheim and Sprenger (2019) did not find evidence for rank dependence among choices between three-outcome lotteries with small and similar monetary outcomes. More research is needed to finely capture the specificities of preferences in choices involving more than two consequences.

In the context of inter-temporal choice, considering "minimal" streams of outcomes also prevents from discriminating between competing models. For example, Blavatskyy (2016) proposed a model of inter-temporal preferences inspired from rank-dependent utility. This model requires streams of at least two outcomes in order to best tested. Considering streams of at least two outcomes is also relevant in the context of risk and time. As explained by Rohde and Yu (2020), binary lotteries with outcome streams introduce a new dimension of the choice objects: correlation of outcomes over time. This aspect of preferences, attitudes towards correlation, can be of interest in itself. It can also be used for comparing models. In particular, models assuming that situations involving risk and time are treated by integrating over risk then over time cannot account for non-neutral attitudes towards correlation. Therefore, observing such attitudes allows to falsify these models.

Overall, considering simple choice is a way to isolate the aspects of preferences that are studied. The complexity of the choices under study that be progressively increased depending on the research question and the available models.

4.1.2 Observing choices in laboratory experiments

Economists have access to two main types of choice data. Observed choices and declared choices. The former offer high validity but little control, the latter offer high control but arguably little validity. Economic experiments are supposed to combine both control and validity: the experimentalist controls the experimental environment, and, thanks to the implementation of real incentives, subjects are assumed to provide valid responses. A limitation is this approach is that experiments create artificial environments, which question the external validity of the results.

The studies presented in this manuscript all rely on data collected through experiments. They nevertheless illustrate the effort that I have been made, in the course of my research, in order to improve the external validity of experiments. Regarding the experimental set ups, the standard context of games involving money has, when possible, been extended to cases that are arguably closer to real-life decisions. This is particularly the case for the experiment of Section 1.4 that created a real research-assistantship contract that allowed to implement real gains and losses of time. My work in progress also explores different types of populations, beyond convenience subjects. For example, in a project with Antoine Nebout (INRAE) and Bruno Ventelou (CNRS & INSERM) we measured risk attitudes of French General Practitioners (GPs) and matched them with observed prescribing behavior. A survey was run on a large sample of French GPs (final sample size of 939) in order to measure their risk attitudes using a choice-based methods, similar to those employed in studies of Sections 1.2 and 1.3. An originality of this study is that respondents could be matched with administrative records, related to their observed prescribing behavior. Specifically, the annual volume of prescribed lab tests was available. Intuitively, the decision to prescribe or not a lab tests may relate to risk attitudes. Not prescribing is a risky option, where the lab-test reduces the risk with a cost for the health system. The decision involves many types of consequences: health of the patient, public health (in the case of viruses for example), public money and also reputation or juridic risk for the practitioner. If risk attitudes are attribute dependent, then it is not clear wether commonly measured risk attitudes towards money correlate with prescribing. The data analysis shows that risk attitudes are indeed correlated with prescription. The effect size is small but has the expected sign: more risk averse GP tend to prescribe lab tests. The effect is also robust to various model specifications.

In Another study in progress, in collaboration with Antoine Nebout (INRAE) and Noémi Berlin (CNRS & University Paris Nanterre), I investigate the correlates of risk and time preferences through a general population survey. Choice-based measures of risk and time preferences towards money were collected through a specific module in the ELIPSS survey [6]. The data where matched with a core

[6] https://www.elipss.fr/fr/

survey that contains many questions related to self-reported real life behavior in the domains of food consumption, health and finance. Preliminary results suggests that the choice-based preferences are significantly correlated, with the expected sign, to various types of behaviors, such as saving or investment, sport activity, alcohol consumption and smoking. These results suggest that despite attribute dependence, choice-based risk attitudes towards money, can have explanatory power for real life behavior involving non-monetary attributes.

All the studies of this manuscript are based on the measurement of indifferences. Using indifferences raises two possible limitations. First, the method used for capturing the indifference may distort preferences. Second, subjects may have incomplete preferences when choices are close to indifference. Regarding the measurement of indifferences, the two popular methods available are the bisection method, and the choice list. The former is based on a series of binary choices, that depend on one-another. Thereby, it is vulnerable to error propagation and is not directly incentive compatible. The latter stacks choices in a list. The presence of the list and the fact that all the choices to be made appear instantaneously may create framing effects that distort preferences. For example, subjects may be reluctant to express a shifting point away from the middle of the list, that constitute a natural anchor. In order to cope with the respective limitations of these methods, the most recent studies (e.g. Section 2.3) combine the two methods. The procedure started with a bisection, that was used to complete a choice list. After the bisection, the pre-filled choices list appeared and the subject had the possibility change the choices from the list before confirming the whole list. All the choices from the list were eligible for the real incentives. This procedure is incentive compatible, allows for error correction and avoids the framing effect of the choice list where preferences are collected. A limitation of this approach is that indifference still required subjects to complete a series of choices for each indifference. This increases the lengths of experiments. Also, the responses given to a choice may depend on the previous choices, that may create anchoring effects or modify the reference point. In order to avoid these possible distortions, subjects should answer only one choice, or a very limited number of choices in a randomized order. Because binary choices are less informative than indifferences, the loss of information should be compensated by stronger assumptions regarding the parametric specification of the model to be used, or the distribution of parameters in the sample of respondents (cf Section 4.1.3).

The second main limitation of experimental procedures is that generally assumed complete preferences. Concretely, respondents must express preferences and do not have the possibility not to answer questions. The main reason for this is that, to my knowledge, no procedure has been found to make incomplete responses incentive compatible. An intuitive solution consists in indicating to subjects that a non-completed choice would be completed using a random device. However, using this procedure, the "no choice" option would be selected by subjects exhibiting a preference for randomization as well as by subjects with incomplete preferences. Given that incomplete preferences are a vast area of empirical research, precious information could be collected from experiments that are not incentivized (e.g. Cubitt et al., 2015).

Most of the experiments reported here implemented real incentives. However, when the design involved a subgroup with hypothetical choices, no differences were observed in terms of responses between the real incentive and the hypothetical choices group (e.g. Sections 1.2 and 3.3). Of course, this does not mean that there is no risk of hypothetical bias. However, what is clear is that the implementation of real incentives imposes sever constraints on investigations. For example, decision theory has been developed for "important" or at least impacting decisions. Nevertheless, because of real incentives, laboratory decisions generally involve very small monetary amounts. It is not clear to me if one should prefer incentivized choices on virtually meaning-less contexts, or hypothetical choices on real-life decisions. Because the latter are less costly to collect, they could be used more frequently to complement incentivized experiments.

4.1.3 Statistical modeling: the role of the assumptions about the specification of functions, parameter distributions and errors

All the studies reported in the manuscript employ at some point econometric estimations. This approach has the advantage to offer estimates of the key parameters of the models and to provide inference tools that allows to comparisons across treatments. The drawback for this approach is that is relies on parametric assumptions about (1) the components of the model (e.g. utility, discount, probability weighting, or beliefs functions), (2) the distribution of errors. In the case of random-coefficient models that capture heterogeneity of preferences, (3) the shape of the distribution of coefficients in the sample must be assumed. Generally, a (log)normal distribution is taken (for non-zero parameters).

Regarding (1), the components of decision models can be measured either in a parameter-free fashion, or using parametric specifications. The former are particularly welcome in exploratory analysis, when the shape of the component is unknown and needs to be investigated. For example, pioneer studies on probability weighting elicited the decisions weights assigned to each probability in order the reveal the inverse S-shape of this function. Then, suitable parametric specifications have been considered to capture key aspects of the function. Parametric specifications have two advantages. First, they save degrees of freedom and allow to capture the global shape of functions using a limited number of data. For example, assuming a one-parameter utility function allows to capture the shape of the utility function using one indifference only (a CE and an indifference probability a la Holt and Laury). A second advantage of parametric specification is that the parameter may have an interpretable value. For example, the parameter of a one parameter exponential utility function can be interpreted in terms of the Arrow-Pratt index of absolute risk aversion. Parameters of probability-weighting functions are generally interpreted in terms of elevation and sensitivity, that have receive psychological interpretations. The former relates to optimism; the latter refers to the cognitive ability to discriminate probabilities, an aspect that relates to perceived ambiguity when probabilities are unknown.

Using parametric specifications for model components has the following drawbacks. First, not all the specifications feature axiomatic foundations. For example, the specification proposed by Prelec et al. (1998) has been axiomatized, whereas the specification proposed by Tversky and Kahneman (1992) has not. Therefore, some specifications are used in a purely descriptive way, with little or no behavioral foundation. Another limitation arises from the combination of several specifications. Not only the theoretical properties of combined specifications can be unknown, but the parameters of each components can become difficult to disentangle, and results can vary depending on the chosen specifications. For example, I often observed that the estimator of the elevation parameter of the probability weighting function is correlated with the one of the utility function, making the identification of each parameter unstable when insufficient data are used. Regarding the impact of the chosen specification on the estimated parameters, Köbberling and Wakker (2005) show how the choice of parametric specifications for the utility function can impact the definition (and therefore, the measurement) of loss aversion. Combinations of parametric specifications are often compared and selected based on goodness of fit (e.g. Stott, 2006). A limitation of this approach is that the goodness of fit of these combinations considered may depend on the set of stimuli. For these reason, I generally try to include in the experiments, some stimuli that allow for parameter-free tests of the axioms underlying the specification utility functions. For example, the assumption of constant absolute risk aversion or constant relative risk aversion is explicitly tested. This allows to base the choice of the utility specification on a test of its property, rather than on goodness of fit. Similarly, studies on inter-temporal choices generally include stimuli allowing to test the assumption of constant discount rate or quasi-hyperbolic discounting using simple statistics. This approach could be extended with the development of stimuli that are explicitly designed to test the specificities of parametric specifications. Steps in this direction include the development of adaptive surveys (e.g. Toubia et al., 2013). In a more descriptive perspective, methods inspired from machine learning can be used to fit very flexible specification that are assumption-free and will not distort measurements because of non-suitable assumptions (e.g. Bertani et al., 2020).

Regarding point (3), the shape distribution of parameters in a given sample can be observed looking at histograms of individual-level estimations. Assumptions related to these distributions may therefore be the easy to test. The most difficult point, in my opinion is point (2): how to models decision errors? There is a large heterogeneity in the way errors are modeled. Generally, an error term is added to the function that values prospects. This assumption could itself be questioned: are errors additive or multiplicative? Second, for the function valuing the prospects is sometimes the utility function (e.g. Stott, 2006, Holt and Laury, 2002), and sometimes, prospects are valued on the scale of outcomes (e.g. Bruhin et al., 2010). Apesteguia and Ballester (2018) showed that the former approach is problematic as it violates a monotonicity property, and the later is preferable. Following this argument, the studies reported in the manuscript express the error on non-utility scales (outcomes, probabilities or events, depending on the study). There are also variations regarding the way data are treated when collected through choice lists that capture indifferences. In some studies, data from a choice list are considered as independent discrete choices (e.g. Holt and Laury, 2002). Given that the discrete choices are grouped in a choice lists, it is unlikely that choice errors are independent within a list. Other studies consider that the choices from a choice list provide only one information: the indifference. They therefore model the data from a given choice list as a single observation of a continuous random variable (e.g. Bruhin et al., 2010). This approach accounts for the fact that choices from a list are not independent, but it does not account for the precision with which the indifference is measured. For example, a matching probability can captured through steps of 0.1 or 0.05 which can have a major effect on the precision of the measurements. Following Beauchamp et al. (2019), I generally model observations from a choice list as an interval containing the observation of a continuous variable, in order to explicitly account for the precision of the measurement in the error specification. With this approach sometimes called interval regression, the size of the interval captures the precision of the measurement.

Despite these variations, all these approaches model errors in terms of random utility. The idea is that the decision maker has known and constant preference but makes errors (that are non-biased and often symmetric) when deciding. These assumptions are questionable. The decision maker may have incomplete or changing preferences. Also, errors may not always be non-biased: the decision maker may sometimes decide totally at random. Changing preferences may be modeled using a random preference model. In such model, the decision make does not make decision errors, but has non-constant preference and draws from a distribution of preference parameters for each choice (e.g. Eliashberg and Hauser, 1985). This approach has been implement under EU and may be harder to implement under more complicated models. In addition, it assumes that each of the random preference is well behaved, the approach cannot account for violations of stochastic dominance (unlike random utility models), that are not unfrequent in choice data. Regarding the possibility of totally random errors, such pattern can be modeled using a mixture model where observations can derive either from a totally random distribution, or from a distribution centered on well behaved preferences (e.g. von Gaudecker et al., 2011). Overall many different and sometime complementary, sometimes mutually exclusive possibilities are available to model decision errors, and a systematic investigation of their theoretical or empirical implications may be deserved.

4.1.4 Interactions between time and uncertainty: beyond risky gains

All the studies presented so far in this section contribute the recent stream in the empirical literature that questions the impacts of interactions between risk and time on preferences. All of them imply only positive outcomes (gains) and known probabilities (risk). Straightforward extensions of these studies would consist in measuring attitudes towards the timing of consequences or risk resolution in the context of losses. This would be particularly interesting in the context of attitudes towards the timing of risk resolution and realization of consequences. Indeed, in the gain domain, rationality suggests that people should prefer to receive consequences as soon as possible, and receive information with instrumental value as soon as possible. In the loss domain, information with instrumental value should still be preferred to be received as soon as possible, but consequences should be preferred to be realized as late as possible.

If is well known that the implementation of monetary losses in experiments is difficult. Except a few exceptions (e.g. Etchart-Vincent and l'Haridon, 2011), hypothetical choices are generally considered in the loss domain. Another approach consists in endowing subjects with a monetary amount that can be lost in the course of the experiment. This approach relies on the assumption that the endowment is well integrating in the wealth of the subject. In the context of choices involving time, this approach creates another complication: how to make subjects come back to the experimenter in the future and pay their loss? A solution, inspired from the experimental set ups in the studies of Section 2 would consist in using time consequences (e.g. Section 1.2) or real effort (Augenblick et al., 2015). For example, the study reported in Section 1.4 implemented losses in the future, in the form of extra time spend on a task under a contract with fixed payment. This procedure, albeit expensive and complex in terms of implementation, could be used for investigating attitudes towards risky inter-temporal choice, towards delayed outcomes or delayed resolution of risk.

A second direction for follow-up studies would be to replace, in these studies, risk by ambiguity. Replacing risk, materialized by Ellsberg's known urn by an unknown urn would be straightforward in terms of experimental implementation. Given that most real-life decisions involve natural source, interactions between uncertainty and time should rather be measured using natural sources. These sources require the measurement of beliefs. To this end, the method for measuring beliefs about natural source, presented in Section 2.2 could be used. In the case of delayed resolution of uncertainty, using natural sources brings however a difficulty: the value of the source realized at time period t' may follow a different distribution than the value of the same source at time t. The studies should therefore focus on stationary sources, thereby restraining the scope of possible natural sources considered. Another approach would be to consider the value at time t and disclose it later. This option implies to consider sources that are privately observed by the experimentalist, which also restrains the scope of possible sources. The most general option would then consist in measuring beliefs towards the source at different time periods. This approach is also the one requiring the largest amount of observations.

4.2 Directions for further research

I now draw directions for further theoretical and empirical research, in order to address questions raised by the results of the studies presented in the manuscript. The studies reported so far highlight the complexity of preferences towards uncertainty and time. In particular, attitudes vary depending on the type of consequences, the source of uncertainty and the payment time, and these aspects also impact the probability weighting. These results suggest that the probability-weighting function is a flexible component of uncertainty and time models, that can account for many aspects of preferences. Nevertheless, models using probability weighting offer the flexibility to fit data ex-post, but they suffer from a lack of parsimony. If the weighting function is attribute/source/time dependent, is there an infinite number of weighting functions, for a single decision maker, as there is an infinite number of attributes, source or time periods? Besides the problem of parsimony, these models also lack of prediction power: knowing the probability weighting function for a given attribute, source or time period, does not allow to make any informed prediction about the probability weighting function in another context. Additional structure is needed in order to link these functions one to another. I now present suggestions for addressing this issue.

4.2.1 Linking probability weighting functions across attributes

As we saw with the studies reported in Section 1, behavioral model have the flexibility to account for attribute dependence. A further step may consist in the development of models that could also explain or predict it. An approach for explicitly modeling attribute dependence may consist in a two-stage process where consequences are first converted to their monetary equivalent, and a usual risk model applies to the risk equivalents. This model would simply explain how the utility function varies from one attribute to another, but would not explain why decision weights are also attribute dependent. A key insight from this model would be to allow for imprecise monetary valuations of outcomes. The imprecision may derived either from incomplete preferences in the tradeoffs between money and the attribute, from a lack of experience with decisions involving this attribute, or because the value of the non-monetary consequence can be risky. Consider for example a gain of one hour. Decision makers may not have a precise idea of the value of such consequence. The value itself may also be uncertain. For example the utility derived from one hour of free time can depend on the weather or the health or mental state of the decision maker, that are unknown ex ante. For these reasons, the non-monetary consequences are converted to monetary lotteries, and the two stage model then involved compound lotteries. If preferences do not satisfy reduction of compound lotteries, preferences will be impacted by the imprecision on the evaluation of non-monetary consequences.

Here is simple example. Consider the choice between saving 1h for sure and 2h or noting with probability 0.5. Suppose that the monetary equivalent of 1 hours (2 hours) is 10 euros (20 euros) if the weather is good (an event to which the DM assigns a probability ρ) and 0 if it rains. If we assume RDU preference and denote u the utility of money and w the probability weighting for time, the sure outcome will be chosen if

$$
\begin{aligned}
w(\rho)u(10) &> w(0.5\rho)u(20) \\
u(10) &> \frac{w(0.5\rho)}{w(\rho)}u(20)
\end{aligned}
$$

The decision maker will behave as if she used a specific weighting function $w^\star(x) = \frac{w(\rho x)}{w(\rho)}$ for non-monetary consequences. This model builds on the idea proposed by Epper and Fehr-Duda (2018) for accounting for anomalies in decisions involving risky and inter-temporal consequences. Applied to attribute dependence, the model can explain and predict why/to what extent, the probability weighting function is impacted by the degree of imprecision in the evaluation of the utility of non-monetary consequences.

In this very simple example, the weighting function $w^\star$ observed for time risk would be more elevated than the weighting function for money. This pattern is consistent with the empirical findings reported in the study of Section 1.2. The model deserves to be further developed. For example, the monetary evaluation could take the form of intervals of possible values. This assumption would allow to cross-fertilize research on attribute dependence and research on imprecise consequences (e.g Liu et al., 2020).

4.2.2 Linking probability weighting functions across time periods

In Section 3.4 we considered tradeoffs between prospects $(X, p^0; x)$ and prospects $(X, p^t; x)$. In such situation, the DM must choose between a probability p^0 solved now, and a probability p^t , possibly larger but solved sooner. This types of tradeoffs remind choices between smaller-sooner and larger-later outcomes made in inter-temporal choice, and modeled through discounting. Following this analogy, the matching present probability m^0 such that $(X, m^0; x) \sim (X, p^t; x)$ could be modeled as $m^0 = D_r(t)p$ where $D_r(t)$ is a discount function that captures attitudes towards resolution timing (hence the subscript r). The discount function $D_r(t)$ could then become the link between probability weighting functions of risk solved at different time periods: $w_t(p) = w(D_r(t)p)$, where w is the weighting function for genuine (i.e. atemporal) risk. Imposing restrictions in $D_r(t)$, like in inter-temporal discounting, would impose a particular structure linking temporal weighting functions one to another.

The same idea can be applied to situations where risk is solved now but outcomes are delayed. Consider the choice between $(X_0, p; 0)$ and $(X_t, q; 0)$. Preference for present implies that a DM is ready to give up on winning probabilities in order to receive the consequence sooner than later. The matching present probability m such that $(X_0, m; 0) \sim (X_t, p; 0)$ could be modeled as $m = D_o(t)p$ where $D_o(t)$ is a discount function that captures attitudes towards outcome-payment timing (hence the subscript o). Such a model has been proposed by Baucells and Heukamp (2012). The authors propose a condition of probability-time tradeoff consistency that is necessary and sufficient such that $D_o(t) = exp(-\rho t)$ where ρ is a discount rate.

Combining the two dimensions, preferences between objects of type $(X_t, p^\tau; 0)$ with $p < 1$ and $\tau < t$, could be modeled using the following functional

$$w_{t,\tau}(p)u(X) \text{ with } w(D_o(t)D_r(\tau)p)$$

Adding condition of consistency to the time probability tradeoffs for outcome delays and resolution delays may lead to the specification

$$w_{t,\tau}(p) = w(e^{-\rho t - \mu\tau}p)$$

where ρ is a discount rate capturing attitudes towards delayed outcomes and μ is a discount rate capturing attitudes towards delayed resolution of risk. Theoretical investigations of the necessary and sufficient properties for this specification to hold are in progress in collaboration with Manel Baucells (University of Virginia) and Veronica Cappelli (Stanford University).

The project is completed by an experiment that aims at measuring and comparing discount rates

ρ and μ.

Whatever the goodness of fit of this model to experimental data, an obvious limitation of this model is that it applies only to binary lotteries with one non-zero outcome. Further research is thus also needed to investigate if similar ideas can be applied to more general lotteries.

Section 3.3 shows that DM are more risk tolerant when outcomes are received in the future. The future is intrinsically uncertain. Even if an outcome is announced to be obtained "for sure" in the future, there are risks that the decision maker will be be able to receive it. A first type of risk relate to the trust in the contract that announces the outcomes in the future. The contractor may default. Even if the contract is reliable there may be reasons why the decision maker may not be able to receive the consequence. Death is one of them, that precludes the notion of absolute certainty. The fact that future is risky have been used to explain present bias. Indeed, the risk induced by (even small) delays in the reception of outcomes may introduce a discontinuity or a gap in the discount function. Epper and Fehr-Duda (2018) show that this aspect can account for other anomalies in preferences towards risk and time, including those reported in the study of Section 3.3. Consider a time lottery (x_t, p, y_t) that pays at time t, x with probability p (solved now) or y with probability $1 - p$. Assume that the DM considers that there is a survival probability $\rho_t < 1$ that either of these consequences will be indeed received. The DM may recode the prospect as a three-outcome lottery $(x_t, \rho p, y_t, \rho(1-p), 0)$. A certain outcome c_t received at t would be recoded as $(c_t, \rho, 0)$.

Therefore, under a discounted RDU model as the one considered in Section 3.3, an indifference $c_t \sim (x_t, p, y_t)$ would write

$$w(\rho_t)u(c)D(t) = w(\rho p)D(t)u(x) + [w(\rho) - w(\rho p)]u(y)D(t)$$

$$\begin{aligned} w(\rho_t)u(c)D(t) &= w(\rho_t p)D(t)u(x) + [w(\rho_t) - w(\rho_t p)]u(y)D(t) \\ w(\rho)u(c) &= w(\rho_t p)[u(x) - u(y)] + w(\rho_t)u(y) \\ u(c) &= \frac{w(\rho_t p)}{w(\rho_t)}[u(x) - u(y)] + u(y) \end{aligned}$$

The formula correspond to the time dependent RDU formula use in Section 3.3 with $w_t(p) = \frac{w(\rho_t p)}{w(\rho_t)}$. Further assuming discrete time periods with a constant survival probability ρ from one period to another, and also assuming that DM compound the survival probability rationally over periods, we would have $w_t(p) = \frac{w(\rho^t p)}{w(\rho^t)}$.

This model explains why DM exhibit more elevated probability weighting when consequences are received in the future. The last version of the model is very parsimonious at the probability weighting functions for all time periods are connected to the standard one through a single parameter ρ that has a direct behavioral interpretation.

This model is tractable and can be estimated with the design presented in Section 3.3. It could be used for in further empirical investigations in order to see how the survival probability is compounded across periods. Manipulations in the experimental set up could also allow to observe which aspect of the experimental environment or which characteristics of the subject (e.g. age) impact the revealed survival probability.

Besides its tractability, a limitation of this model is that (1) it assumes that subjects rationally edit the future lotteries, compounding the survival probability with the lottery probability (2) weight this compounded probability by the risky probability weighting function. Indeed, the survival probability is not objective and DM may exhibit non-neutral ambiguity attitudes towards it. The model may thus deserve to be refined in order to account for non-neutral ambiguity aversion, even though the impact of ambiguity attitudes may be of a second order of magnitude, in comparison to the impact of time.

4.2.3 Linking probability weighting functions across degrees of ambiguity

Ambiguity is generally understood as imprecision about probabilities. A natural way to express imprecision is to replace values by intervals. In the case of binary prospect with an unknown probability $\tilde{p}$ to receive the winning event, ambiguity about $\tilde{p}$ can be represented by an interval $[p^-, p^+]$. Here, the size of the interval, $p^+ - p^-$ measures the imprecision about $\tilde{p}$, and provides a measure of "objective ambiguity". Chew et al. (2017) recently studied attitudes towards such situations, referred to as partial ambiguity, in their paper. Intervals of size 1 correspond to Ellsberg's unknown urns, and intervals of size 0 correspond to Ellsberg known urn. Intermediate cases allow to scan the continuous range of ambiguity degrees between the two extreme case materialized by Ellsberg's urns. If there exists an infinite number of ambiguity levels $a = p^+ - p^-$, under the source model, each of these levels of ambiguity should be captured by a specific probability weighting function w_a. This implies an infinite number of source function, with no a priori relationship linking them. This creates an obvious lack of parsimony for the model. Furthermore, assume that attitudes are known for intervals of size $a = 0.5$ and for $a = 0.1$ the model cannot tell anything about attitudes towards intervals of size $a = 0.3$.

In a project in progress with M. Abdellaoui (CNRS & HEC Paris), T. Astebro (HEC Paris) and C. Paraschiv (Université Paris Descartes), we model propose to measure attitudes towards objective ambiguity under a restrained specification of the source model that imposes that:

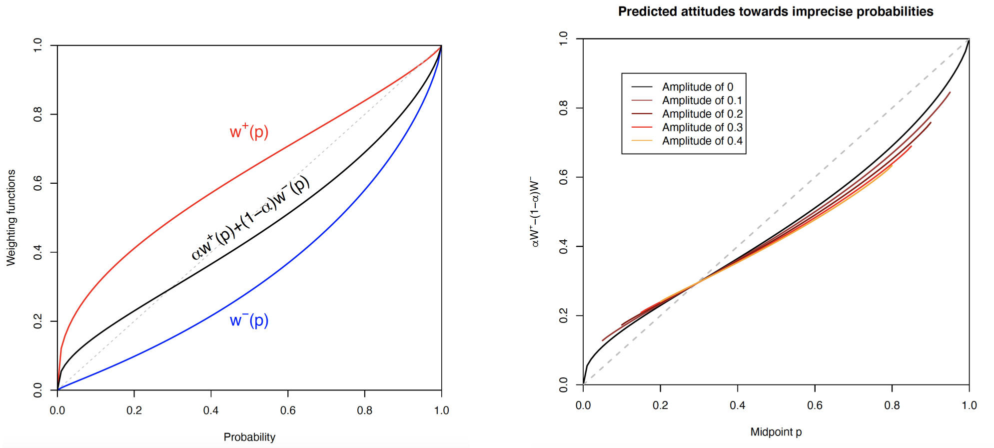

Figure 10: Estimated mean patterns

$$w_a(p) = \alpha w^+(p + \frac{a}{2}) + (1-\alpha)w^-(p - \frac{a}{2})$$

for p in $[\frac{a}{2}, 1 - \frac{a}{2}]$, where w^+ and w^- are strictly increasing probability weighting functions and α lies between 0 and 1.

This model has been elicited in a lab experiment, and plausible shapes for w^+ and w^- were obtained: they are illustrated in Figure 10 (left hand side panel).

In particular, our mean parameter estimates allow to predict the source function w_a for any level of ambiguity, as illustrated in the right-hand side panel of the figure. The model captures the pattern revealed by the study in Section 2.3: the weighting function becomes more sensitive and more elevated by ambiguity diminishes. The model therefore offers a parsimonious and plausible "restrained" version of the source model for objective ambiguity.

Considering the limit case where $a = 0$, the model also applies to risk and provides a specification of the risk weighting function. Given our empirical finding that w^+ is globally concave and and w^- is globally convex, the model shed a new light on the interpretation of the probability weighting of risk and its inverse-S shape. The model indeed suggests that this shape originates from the convex combination of two components: the possibility effect, captured by w^+ , and the certainty effect, captured by w^-. This interpretation is a cognitively plausible explanation of the reason why the risky weighting function is inverse S-shape. It also offers new directions parametric specifications for this function. For example, w^+ and w^- could be modeled by power functions with different parameters, allowing to isolated the measurement of the possibility effect (or more generally, the shape of the function for likely events) from the possibility effect (or more generally, the shape of the function for unlikely events).

Concluding remarks

This documents gives the opportunity to take an overview on my research activity in the last 6 years. This section presents comments on key aspects of the development of my research over this period, and its connection with other aspects of academic life such as teaching and publicizing.

- From decision theory to behavioral decision analysis

All the studies reported in this manuscript belong to a field that can be called "behavioral decision analysis" or "behavioral decision science". I propose a rough definition and personal view of this field. Behavioral decision analysis consists in analyzing observed decisions while accounting for aspects that are considered as irrational or irrelevant by decision science. Decision science formalizes decision contexts and builds rules of rational choice, thereby taking a normative perspective. Decision analysis uses the normative models of decision science for making recommendations in applied decision making. In crude words, the decision analyst would help a decision maker saying: "given the preferences that I measured assuming that you are rational, you should do choose this option". This approach can be considered as paradoxical. If the decision maker has indeed complete and rational preferences, why would she need the help of a decision analyst? We could argue in this case, that either (1) the contribution of the decision analyst is simply to reduce the cognitive cost of the decision by separating the decision process into simpler steps, or that (2) the contribution of the decision analyst is to help the decision maker building her (rational) preferences, thereby taking the risk to distort her (possibly irrational) preferences which would lead to suboptimal well being. The opposite perspective would consist in taking a purely descriptive perspective. To this aim, the best approach would be to employ data science. Data science, and in particular machine learning explicitly aims at recovering choice patterns. The best performing methods, such as random forests, boosted trees or neural networks (that are the basis of deep learning), are also the ones that are the less interpretable. Data science can be used for developing decision-support tools. In such cases, the analyst would help the decision maker saying "given the decisions that you (or other people) made in the past, you should choose this". In this case, the analyst would be of little help. The decision help would again reduce the cognitive cost of the decision, and its main contribution would be to help the decision maker to be consistent with previous choices. However, if the decision maker is not rational, being consistent is not necessarily an objective, and the decision maker can consistently make suboptimal decisions.

Behavioral decision analysis offers a compromise between these two approaches. All the behavioral models considered in this manuscript are extensions of normative models: they contain the ingredients of rational models, and augment them with "behavioral components" that capture deviations from rationality. This features allow to account for the two aspects of preferences: rational ones and irrational ones. Therefore, the analyst can keep both a normative **and** a descriptive analysis. In particular, the components of the model that capture deviations from rationality have behavioral interpretation. This characteristic has several advantages. It allows not only to describe preferences, but also to understand them. This understanding gives room for policy intervention. In the context of decision analysis, the analyst can leave the choice to the decision maker saying: "according to your previous choices, you take into account aspects that can be considered as irrelevant, if you want to keep taking them into account, you should choose this, otherwise, omitting these aspects, you should choose that". In some case, the decision maker may realize that (s)he is indeed sensitive to aspects that should be ignored, and decide to follow the rational part of the elicited preferences. This is likely to be the case for framing effects, for example. In other cases, (s)he may prefer to keep the behavioral aspects into account. This is likely to be the case with ambiguity attitudes. for example. Indeed, students generally stick to their preference even after being explained that the preference can be considered as irrational.

Because the "irrational" components of behavioral models have behavioral interpretation, the analysis may reveal that these aspects are, eventually, not so irrational. The model accounting for a survival probability presented in Section 4.5 offers an illustration. The fact that risk attitudes change depending on the timing of reception of consequences may appear as irrational. However, when this pattern is explained by the fact that the future is perceived as uncertain, it appears decision makers may have good reasons to follow this pattern. In this example, the appearing bias may come from the fact that the decision context was not modeled properly: the analyst considered the possibility of future outcomes obtained with certainty, although the future is necessarily uncertain. The survival probability should be accounted for by rational models. This example illustrates that behavioral analysis can also enrich rational decision models.

Eventually, behavioral analysis can be used to measure preferences that can be compatible with rational models, and also robust to behavioral biases. This is the case of the method presented in Section 3.2 for measuring beliefs in the form of subjective probabilities. These probabilities are consistent with EU probabilities, and are also robust to non-neutral ambiguity attitudes. In this respect, this type of approach dominates measurement methods that are valid only under rational models.

The knowledge on behavioral decision analysis that I developed for and from my research was used for teaching and publicizing this approach to a broader audience. Regarding teaching, I built a course of Behavioral Economics for HEC Paris L3/M1 students, and taught this course four years in a row, between 2015 and 2019. For publicizing, I wrote a few papers commenting real-life decisions in the light of behavioral science, in media dedicated to general audience.[7].

- Development of digital tools for research and teaching

The validity of the choice data can also be impacted to the number of choice and/or duration of the experiment. When subjects are exposed to many choices, it is likely that fatigue effects, or use of heuristics appear and distort observed choices.

In this regards also, efforts have been made in order to use econometric methods that optimize available data, and allow to estimate models from fewer choices. The knowledge acquired in econometrics allowed me to develop a course in advanced econometric choice modeling, in addition to the introductory course that I give to HEC Master and PhD Students.

In terms of technical developments, efforts have also been invested in the development of digital tools for collecting choices. While my first experiments were coded in Python, in the last years, I developed web applications that can be used for collecting choices. These applications can be deployed more easily (in large scale surveys for example) and can be highly interactive. Interactivity allows for example to include procedures that check the attention or the understanding of the respondents before or in the course of the experiment. The skills acquired in the development these apps was also used for pedagogical purposes in my teaching activity[8].

- Collaborations and supervising

Since the defense of my PhD, I mainly focused my research activity into 3 types of projects. The first consisted in valorization of projects initiated during my PhD (e.g. Section 3.2), or in follow-up projects. The second consisted in working on the research project proposed when applying to CNRS. The project was about interactions between risk and time, and studies presented in Section 3 contribute to this direction. The third type of projects have been initiated through interactions with the research community. Working in the decision science team at GREGHEC gave me the opportunity to collaborate with senior researchers, and also to develop projects with students HEC PhD students (e.g. Cedric Gutierrez, Fan Wang, Veronica Cappelli). I also developed collaborations with senior or senior researchers from other institutions. Table 6 lists the co-authors with whom projects have been initiated or papers have been published since my PhD. Collaborating with junior and senior researchers from national and international institutions clearly stimulates the development of research ideas and skills. It also allows to vary the type of role played in the projects. Teaching

[7] see https://www.hec.edu/en/knowledge/articles/it-rational-stockpile-times-crisis , or https://www.forbes.fr/technologie/nudges-et-intelligence-artificielle-unis-pour-le-meilleur-ou-pour-le-pire/?cn-reloaded=1 , or https://www.ladn.eu/entreprises-innovantes/parole-expert/nudge-technique-marketing-influence-nos-comportements/

[8] https://www.hec.edu/en/knowledge/instants/how-hec-professors-enhance-their-research-and-courses-using-data-visualization-app

Name	Institution	Paper(s) published	Working papers	Projects in Progress
Abdellaoui Mohammed	CNRS & HEC Paris	3	2	2
Astebro Tom	HEC Paris			1
Aydogan Ilke	IESEG		1	
Baillon Aurélien	Erasmus University		1	
Baucells Manel	University of Virginia			1
Bertani Nicolo	INSEAD			1
Bleichrodt Han	Erasmus University	1		
Cappelli Veronica	Stanford University			1
Gutierrez Cedric	University of Bocconi	1	1	1
Hill Brian	CNRS & HEC Paris		1	
L'Haridon Olivier	University of Rennes	1		1
Li Chen	Erasmus University		1	1
Maafi Hela	University Paris 8		2	
Mun Sofiia	Paris School of Economics			1
Nebout Antoine	INRAE and Paris-Saclay		1	2
Panin Amma	World Bank	1	1	1
Paraschiv Corina	University Paris Descartes	1	1	
Ventelou Bruno	INSERM & Aix Marseille		1	
Vieider Ferdinand	University of Gent	1	2	1
Wang Fan	ESSEC Singapore			1

Table 6: List of co-authors since the PhD defense

at HEC Paris also gave to advise 9 master students in the last three years. I am also currently collaborating with a PhD student, Sofiia Mun, from Paris School of Economics. These collaborations gave me an idea of the investments and skills required for advising PhD students.

References

Abdellaoui, M., Baillon, A., Placido, L., and Wakker, P. P. (2011a). The rich domain of uncertainty: Source functions and their experimental implementation. *American Economic Review*, 101(2):695–723.

Abdellaoui, M., Bleichrodt, H., and l'Haridon, O. (2008). A tractable method to measure utility and loss aversion under prospect theory. *Journal of Risk and uncertainty*, 36(3):245.

Abdellaoui, M., Diecidue, E., and Öncüler, A. (2011b). Risk preferences at different time periods: An experimental investigation. *Management Science*, 57(5):975–987.

Abdellaoui, M., l'Haridon, O., and Paraschiv, C. (2011c). Experienced vs. described uncertainty: Do we need two prospect theory specifications? *Management Science*, 57(10):1879–1895.

Abrahamsson, M. and Johansson, H. (2006). Risk preferences regarding multiple fatalities and some implications for societal risk decision making-an empirical study. *Journal of risk research*, 9(7):703–715.

Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine. *Econometrica: Journal of the Econometric Society*, pages 503–546.

Andersen, S., Harrison, G. W., Lau, M. I., and Rutström, E. E. (2008). Eliciting risk and time preferences. *Econometrica*, 76(3):583–618.

Andreoni, J. and Sprenger, C. (2012). Risk preferences are not time preferences. *American Economic Review*, 102(7):3357–76.

Apesteguia, J. and Ballester, M. A. (2018). Monotone stochastic choice models: The case of risk and time preferences. *Journal of Political Economy*, 126(1):74–106.

Augenblick, N., Niederle, M., and Sprenger, C. (2015). Working Over Time: Dynamic Inconsistency in Real Effort Tasks. *Quarterly Journal of Economics*, 130(3):1067–1115.

Baillon, A. (2008). Eliciting subjective probabilities through exchangeable events: An advantage and a limitation. *Decision Analysis*, 5(2):76–87.

Baillon, A., Huang, Z., Selim, A., and Wakker, P. P. (2018). Measuring ambiguity attitudes for all (natural) events. *Econometrica*, 86(5):1839–1858.

Baucells, M. and Heukamp, F. H. (2012). Probability and time trade-off. *Management Science*, 58(4):831–842.

Baucells, M. and Katsikopoulos, K. V. (2010). Descriptive models of decision making. *Wiley Encyclopedia of Operations Research and Management Science.*

Beauchamp, J. P., Benjamin, D. J., Laibson, D. I., and Chabris, C. F. (2019). Measuring and controlling for the compromise effect when estimating risk preference parameters. *Experimental Economics*, pages 1–31.

Bernheim, B. D. and Sprenger, C. (2019). Direct tests of cumulative prospect theory. *Available at SSRN 3350196.*

Bertani, N., Boukhatem, A., Diecidue, E., Perny, P., and Viappiani, P. (2020). Fast and simple adaptive elicitations: Experimental test for probability weighting. *Available at SSRN 3569625.*

Birnbaum, M. H. and Schmidt, U. (2008). An experimental investigation of violations of transitivity in choice under uncertainty. *Journal of Risk and Uncertainty*, 37(1):77.

Blavatskyy, P. R. (2016). A monotone model of intertemporal choice. *Economic Theory*, 62(4):785–812.

Bleichrodt, H., Gao, Y., and Rohde, K. I. (2016). A measurement of decreasing impatience for health and money. *Journal of Risk and Uncertainty*, 52(3):213–231.

Boiteux, M. and Baumstark, L. (2001). Transports: choix des investissements et coût des nuisances. *La documentation française*, 325.

Bruhin, A., Fehr-Duda, H., and Epper, T. (2010). Risk and rationality: Uncovering heterogeneity in probability distortion. *Econometrica*, 78(4):1375–1412.

Camerer, C. and Weber, M. (1992). Recent developments in modeling preferences: Uncertainty and ambiguity. *Journal of Risk and Uncertainty*, 5:325–370.

Chew, S. H., Ebstein, R. P., and Zhong, S. (2012). Ambiguity aversion and familiarity bias: Evidence from behavioral and gene association studies. *Journal of Risk and Uncertainty*, 44(1):1–18.

Chew, S. H., Miao, B., and Zhong, S. (2017). Partial ambiguity. *Econometrica*, 85(4):1239–1260.

Chew, S. H. and Sagi, J. S. (2008). Small worlds: Modeling attitudes toward sources of uncertainty. *Journal of Economic Theory*, 139(1):1–24.

Cubitt, R. P., Navarro-Martinez, D., and Starmer, C. (2015). On preference imprecision. *Journal of Risk and Uncertainty*, 50(1):1–34.

Cubitt, R. P. and Read, D. (2007). Can intertemporal choice experiments elicit time preferences for consumption? *Experimental Economics*, 10(4):369–389.

Diecidue, E. and Wakker, P. P. (2001). On the intuition of rank-dependent utility. *Journal of Risk and Uncertainty*, 23(3):281–298.

Dimmock, S. G., Kouwenberg, R., and Wakker, P. P. (2016). Ambiguity attitudes in a large representative sample. *Management Science*, 62(5):1363–1380.

Ebert, J. E. J. and Prelec, D. (2007). The Fragility of Time: Time-Insensitivity and Valuation of the Near and Far Future. *Management Science*, 53(9):1423–1438.

Ebert, S. (2020). Decision making when things are only a matter of time. *Operations Research*.

Eliashberg, J. and Hauser, J. R. (1985). A measurement error approach for modeling consumer risk preference. *Management Science*, 31(1):1–25.

Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The quarterly journal of economics*, pages 643–669.

Epper, T. and Fehr-Duda, H. (2018). The missing link: Unifying risk taking and time discounting. *University of Zurich Department of Economics Working Paper*, (96).

Epper, T., Fehr-Duda, H., and Bruhin, A. (2011). Viewing the future through a warped lens: Why uncertainty generates hyperbolic discounting. *Journal of risk and uncertainty*, 43(3):169–203.

Epstein, L. G. (2008). Living with risk. *The Review of Economic Studies*, 75(4):1121–1141.

Etchart-Vincent, N. and l'Haridon, O. (2011). Monetary incentives in the loss domain and behavior toward risk: An experimental comparison of three reward schemes including real losses. *Journal of risk and uncertainty*, 42(1):61–83.

Fagley, N. S. and Miller, P. M. (1997). Framing effects and arenas of choice: Your money or your life? *Organizational behavior and human decision processes*, 71(3):355–373.

Fehr-Duda, H., Bruhin, A., Epper, T., and Schubert, R. (2010). Rationality on the rise: Why relative risk aversion increases with stake size. *Journal of Risk and Uncertainty*, 40(2):147–180.

Fischhoff, B. (1983). Predicting frames. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 9(1):103.

Fox, C. R. and Tversky, A. (1995). Ambiguity aversion and comparative ignorance. *The quarterly journal of economics*, 110(3):585–603.

Fox, C. R. and Tversky, A. (1998). A belief-based account of decision under uncertainty. *Management science*, 44(7):879–895.

Frederick, S., Loewenstein, G., and O'Donoghue, T. (2002). Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature*, 40(2):351–401.

Fredrickson, B. L. and Kahneman, D. (1993). Duration neglect in retrospective evaluations of affective episodes. *Journal of personality and social psychology*, 65(1):45.

Gerber, A. and Rohde, K. I. (2010). Risk and preference reversals in intertemporal choice. *Journal of Economic Behavior & Organization*, 76(3):654–668.

Ghirardato, P. and Marinacci, M. (2001). Risk, ambiguity, and the separation of utility and beliefs. *Mathematics of operations research*, 26(4):864–890.

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with a non-unique prior.

Halevy, Y. (2008). Strotz meets allais: Diminishing impatience and the certainty effect. *American Economic Review*, 98(3):1145–62.

Holt, C. A. and Laury, S. K. (2002). Risk aversion and incentive effects. *American economic review*, 92(5):1644–1655.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society*, pages 263–291.

Keeney, R. L. (1980). Evaluating alternatives involving potential fatalities. *Operations Research*, 28(1):188–205.

Keeney, R. L., Raiffa, H., et al. (1993). *Decisions with multiple objectives: preferences and value trade-offs*. Cambridge university press.

Keren, G. and Roelofsma, P. (1995). Immediacy and certainty in intertemporal choice. *Organizational Behavior and Human Decision Processes*, 63(3):287–297.

Klibanoff, P., Marinacci, M., and Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892.

Köbberling, V. and Wakker, P. P. (2005). An index of loss aversion. *Journal of Economic Theory*, 122(1):119–131.

Kothiyal, A., Spinu, V., and Wakker, P. P. (2014). An experimental test of prospect theory for predicting choice under ambiguity. *Journal of risk and uncertainty*, 48(1):1–17.

Kreps, D. M. and Porteus, E. L. (1978). Temporal resolution of uncertainty and dynamic choice theory. *Econometrica: journal of the Econometric Society*, pages 185–200.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2):443–478.

Leclerc, F., Schmitt, B. H., and Dube, L. (1995). Waiting time and decision making: Is time like money? *Journal of Consumer Research*, 22(1):110–119.

Li, Z., Müller, J., Wakker, P. P., and Wang, T. V. (2017). The rich domain of ambiguity explored. *Management Science*, 64(7):3227–3240.

Liu, Y., Heath, T. B., and Onculer, A. (2020). The future ambiguity effect: How narrow payoff ranges increase future payoff appeal. *Management Science*.

Loewenstein, G. (1987). Anticipation and the Valuation of Delayed Consumption. *Economic Journal*, 97(387):666–684.

Loewenstein, G. and Prelec, D. (1992). Anomalies in Intertemporal Choice: Evidence and an Interpretation. *The Quarterly Journal of Economics*, 107(2):573–597.

Meissner, T. and Pfeiffer, P. (2015). I want to know it now: Measuring preferences over the temporal resolution of consumption uncertainty. Unpublished Manuscript. Available at SSRN: https://ssrn.com/abstract=2654668.

O'Donoghue, T. and Sprenger, C. (2018). Reference-dependent preferences. In *Handbook of Behavioral Economics: Applications and Foundations 1*, volume 1, pages 1–77. Elsevier.

Prelec, D. et al. (1998). The probability weighting function. *ECONOMETRICA-EVANSTON ILL-*, 66:497–528.

Ramsey, F. P. (1931). *The foundations of mathematics and other logical essays.* Number 214. K. Paul, Trench, Trubner & Company, Limited.

Rheinberger, C. M. (2010). Experimental evidence against the paradigm of mortality risk aversion. *Risk analysis*, 30(4):590–604.

Rohde, K. I. (2010). The hyperbolic factor: A measure of time inconsistency. *Journal of Risk and Uncertainty*, 41(2):125–140.

Rohde, K. I. and Yu, X. (2020). Aggregating over risk and time in risky intertemporal choice: which order? *Working paper.*

Samuelson, P. A. (1937). A note on measurement of utility. *The review of economic studies*, 4(2):155–161.

Savage, L. J. (1954). *The foundations of statistics*.

Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica: Journal of the Econometric Society*, pages 571–587.

Starmer, C. (2000). Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. *Journal of economic literature*, 38(2):332–382.

Stott, H. P. (2006). Cumulative prospect theory's functional menagerie. *Journal of Risk and uncertainty*, 32(2):101–130.

Toubia, O., Johnson, E., Evgeniou, T., and Delquié, P. (2013). Dynamic experiments for estimating preferences: An adaptive method of eliciting time and risk parameters. *Management Science*, 59(3):613–640.

Train, K. E. (2009). *Discrete choice methods with simulation*. Cambridge university press.

Trautmann, S. T. and Van De Kuilen, G. (2015). Ambiguity attitudes. *The Wiley Blackwell handbook of judgment and decision making*, 1:89–116.

Tversky, A. and Fox, C. R. (1995). Weighing risk and uncertainty. *Psychological review*, 102(2):269.

Tversky, A. and Kahneman, D. (1981). The framing of decisions and the psychology of choice. *SCIENCE*, 211:30.

Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4):297–323.

von Gaudecker, H.-M., van Soest, A., and Wengström, E. (2011). Heterogeneity in risky choice behavior in a broad population. *The American Economic Review*, 101(2):664–694.

Wakker, P. (1994). Separating marginal utility and probabilistic risk aversion. *Theory and decision*, 36(1):1–44.

Wang, X. T. (1996). Framing effects: Dynamics and task domains. *Organizational Behavior and Human Decision Processes*, 68(2):145–157.

Weber, B. J. and Chapman, G. B. (2005). The combined effects of risk and time on choice: Does uncertainty eliminate the immediacy effect? does delay eliminate the certainty effect? *Organizational Behavior and Human Decision Processes*, 96(2):104–118.

Wu, G. (1999). Anxiety and decision making with delayed resolution of uncertainty. *Theory and Decision*, 46(2):159–199.

Zeckhauser, R. (1973). Time as the ultimate source of utility. *The Quarterly Journal of Economics*, pages 668–675.