European University Institute

Max Weber Programme

EUI Working Papers

MWP 2007/27

Decision Makers Facing Uncertainty:

Theory versus Evidence

Paolo E. Giordani   Karl H. Schlag   Sanne Zwart

EUROPEAN UNIVERSITY INSTITUTE

MAX WEBER PROGRAMME

Decision Makers Facing Uncertainty:
Theory versus Evidence

PAOLO E. GIORDANI KARL H. SCHLAG SANNE ZWART

EUI Working Paper MWP No. 2007/27

ISSN 1830-7728

Printed in Italy
European University Institute
Badia Fiesolana
I – 50014 San Domenico di Fiesole (FI)
Italy

http://www.eui.eu/
http://cadmus.eui.eu/

Abstract

We consider three competing normative theories of how to make choices when facing uncertainty: subjective expected utility, maximin utility and minimax regret. In simple decision problems, we compare how decision makers under each of these theories value safe options, freedom of choice and information. We then use these models to predict answers to questions in the European Values Survey and use these predictions via a latent class analysis to estimate the distribution of these behaviors across Europe. We find a larger proportion of Bayesians in the Northern countries than in Southern countries. The opposite is true for maximin utility behavior. Only a few are consistent with minimax regret behavior.

Keywords

Uncertainty, maximin, minimax regret, Bayesianism, European Values Survey, latent class analysis.

JEL Classification: D81, C4.

Decision Makers Facing Uncertainty: Theory versus Evidence*

Paolo E. Giordani† **Karl H. Schlag**‡ **Sanne Zwart**§

1 Introduction

In the decision-theoretic literature several alternative approaches have been developed to deal with choice under uncertainty. We consider three different axiomatic approaches to uncertainty where uncertainty refers to not knowing the distribution of the underlying states. In the Bayesian or subjective expected utility model (Savage, 1954) uncertainty is transformed into risk by assessing a probability distribution over the set of possible states. The maximin utility model (Wald, 1950, Milnor, 1954) is only concerned with the worst outcome possible for each choice. Minimax regret (Savage, 1951, Milnor, 1954) captures aversion to lost opportunities.

Our aim is to investigate the implications of these alternative behavioral models. At first we identify three contexts where attitudes towards uncertainty play a role. For each context we set up an elementary decision problem and then compare predictions of the three models. Later we use the resulting classifications of behavior in the three different contexts to uncover similarities and differences in human behavior across Europe. To achieve this we identify questions from the European Values Survey (EVS) that can be studied within the contexts we modeled, and use a latent class analysis to estimate the distribution of behavioral types across Europe.

*The authors would like to thank participants at the 2006 EEA conference in Vienna for their comments. Paolo Giordani thankfully acknowlegdes financial support from the Max Weber Programme, EUI. The usual disclaimers apply.

†Via delle Fontanelle 10, 50016 Florence (Italy). Tel.: +39 0554685698, email: paolo.giordani@eui.eu.

‡Via della Piazzuola 43, 50133 Florence (Italy). Tel.: +39 0554685951, email: karl.schlag@eui.eu.

§Via della Piazzuola 43, 50133 Florence (Italy). Tel.: +39 0554685927, email: sanne.zwart@eui.eu.

Each decision problem is constructed to uncover a specific feature of uncertainty, which we can also find in the questions in the EVS. The decision problems have in common the fact that payoffs depend on some unknown state. In 'safe versus uncertain' we are interested in the general attitude towards uncertainty. Here we consider a decision maker (DM) who has to choose between a safe option and an option yielding an uncertain outcome. The DM can avoid uncertainty by choosing the safe option at the expense of possibly foregoing better outcomes. In 'freedom of choice' we wish to investigate how confrontation with uncertainty influences incentives to use initiative. The idea is that preferences for an active or for a passive position can be influenced by how good one thinks one is in dealing with uncertainty. We confront the DM with a decision problem and ask her whether she wants to make the choice herself or whether she wants to leave the choice up to someone else instead. In 'value of information' we investigate how the attitude towards uncertainty influences the willingness to invest in reducing uncertainty by acquiring more information. Information is costly but possibly reveals which state will occur.

In all three decision problems we find that choice varies across the different behavioral models. Combining our findings across these three problems, we identify the following behavioral patterns. Bayesians may or may not like safe choices, like to exert their freedom of choice and give relatively low value to information. Maximin utility DMs like safe options, like freedom of choice and value information highly. Finally, minimax regret DMs are willing to partially engage in uncertain scenarios, moderately like freedom of choice and like to have more information.

In the second step we select the questions asked in the EVS that can be associated to the aspects or contexts surrounding uncertainty we have been modeling. To give an example, the following question from the EVS questionnaire is linked to the 'freedom of choice' problem: "Would you mention the 'opportunity to use initiative' as an important aspect in a job?". We find five relevant questions. Job security and long-term relationships are interpreted as safe options. Choosing to use initiative as well as emphasizing the value of teaching independence to one's own children is associated to the preference for freedom of choice. Preference for a job with not too much pressure is considered as a signal that information is valuable. By establishing a link between the questions and our choice problems we can predict the responses of our three behavioral types to the EVS questions. This allows us to categorize individuals by comparing the arrays of their answers and to thereby empirically gain insights into how Europeans deal with uncertainty.

In the empirical analysis we focus on EVS data from 1999 and consider the fifteen countries belonging to European Union at that date (EU15). We wish to classify how many individuals provide answers that are consistent with each of the behavioral models and to compare the resulting frequencies across Europe. To do this we perform a latent class analysis, where a class is defined by a DM type. The three behavioral patterns identified in the theoretical analysis guide us when imposing the constraints necessary to define the types. For example, if theory suggests that the maximin type likes safe options, we impose its preference for job security and long-term relationship. One additional unspecified (or 'free') type is added to allow for alternative behavior. For each type an array of probabilities of each answer across all questions is estimated. This estimation is done for the entire data set so that types are comparable across countries. Goodness of fit is assessed. In particular, we find support for explaining the data with four types, one behavioral type for each normative model together with the addition free type. The evidence is thus favorable to the restrictions identified in the theoretical analysis.

Across Europe we find around 69% of the sample population to be classified into one of the three behavioral types identified by the theory. In particular, the proportions of Bayesian, maximin utility and minimax regret types average around 39%, 23% and 6% respectively. The remaining 31% fall within the fourth type.

The fourth type is unconstrained, in the sense that no restriction is imposed on its choices. The estimation of the answer probabilities across all questions gives us its behavioral characterization. This type neither likes nor dislikes safe options, strongly dislikes freedom of choice and does not suffer pressure. The essentially 'passive' nature of this type, which prefers that choices are made by others, seems to be its most distinctive feature.

As types are not country-specific but defined uniformly across Europe we can compare frequencies across countries. Overall we find remarkable differences in the cultural attitude towards uncertainty across European countries. Proportions of Bayesians are typically higher than the average across both continental countries (such as Austria, Germany and the Netherlands) and Scandinavian countries (Denmark, Finland and Sweden), while they are typically lower than the average in Mediterranean countries (Greece, Italy, Spain and Portugal) and, to some extent, in Ireland. Roughly the opposite pattern holds for the maximin utility type, whose fractions are generally high in southern European countries and low in most continental and Nordic European countries. In light of the distinctive features of these two types identified above, we can thus essentially confirm the "north-south" interpretation (Hofstede, 2001 among others; see below for a comparison), which is the idea that individuals in southern European countries have a more conservative attitude towards uncertainty than the ones in the northern and continental countries. We also find an east-west divide in the sense that the majority of types in France, Portugal, Spain and Belgium fall within the fourth free type. Western European countries are less clearly captured by the predictions of our theoretical models. Finally, the proportions of minimax regret types are instead lower and generally more homogeneous across European countries albeit with a slight tendency to more behavior consistency with minimax regret in the north.

The empirical part of this paper can be compared to the cross-country comparative psychological literature, which is aimed at assessing essential differences in cultural traits across countries. In particular, Hofstede (2001) has built the *uncertainty avoidance index* (UAI), which measures "the extent to which the members of a culture feel threatened by uncertain or unknown situations"(Hofstede, 2001, p. 263). Our approach is however methodologically distant from Hofstede's one. Hofstede equally weighs the answers to three questions where answers are connected to behavior under uncertainty via analogy, and then builds an index ranking all countries from the least to the most uncertainty avoiding. We instead use decision theory to formulate predictions on the choices of different types along a series of behavioral dimensions (safe vs. uncertain, freedom, value of information). In a latent class analysis we then check whether our theoretical predictions can be accepted, and if so, evaluate the relative proportions of the different DM types across EU15 countries. Among other things, we also find evidence in favor of the "north-south" interpretation.

This paper can also be related to the literature on *ambiguity*, where by ambiguity it is meant the lack of a (single) prior over the possible states of nature or, equivalently, the tendency not to reduce uncertainty to risk (Ellsberg, 1961, Schmeidler, 1989, Gilboa-Schmeidler, 1989). In light of this definition Bayesians are unaffected by ambiguity, as they form priors, while maximin and minimax regret behaviors perceive the choice scenario as ambiguous as they do not form priors. Moreover, we can go further in our interpretation of the free type, and consider its aversion towards freedom of choice as contradicting the essential spirit of Bayesianism. We can then interpret our findings as providing evidence in favor of DMs in northern and continental countries perceiving a relatively less ambiguous choice environment and, hence, being more prone to transform uncertainty into risk than DMs in southern countries.

The rest of the paper is organized as follows. In the next section we introduce and compare the three alternative choice models. In Section 3 we construct the three decision problems and analyze the choices under each model. In Section 4 we link these problems to EVS questions, formulate and test our predictions. In Section 5 we present and comment on the findings.

2 Axiomatic Choice under Uncertainty

According to Knight (1921), *uncertainty* - as opposed to *risk* - is associated with decision problems where no objective probability distribution over the states of the world is given. There are several different axiomatic models of decision making under uncertainty. We restrict our attention to the following three approaches: the subjective expected utility criterion (Savage, 1954, Anscombe and Aumann, 1963), the maximin utility decision rule (Wald, 1950) and the minimax regret criterion (Savage, 1951).

Subjective expected utility (SEU) theory is based on axioms that extend Von Neumann-Morgenstern's expected utility principle, originally developed for risk, to the case of uncertainty. Accordingly, it is as if the decision maker (DM) *subjectively* assesses a probability distribution (or *prior*) over the states of the world, and then selects the action that yields the highest (subjective) expected utility. Discomfort with this approach of treating uncertainty as risk has led to a recent resurgence of alternative theories. Notice also that evidence in neuroscience (Rustichini et al., 2002) shows that different parts of the brain are used depending on whether there is risk or uncertainty.

The most common alternative is the *maximin utility criterion*. It was introduced by Wald (1950), axiomatized by Milnor (1954) and recently by Stoye (2006), and has received increasing attention since Gilboa and Schmeidler (1989). It is as if the following scenario takes place. The DM expects to be punished by a 'malevolent' Nature. She fears for any choice that the worst outcome possible under this choice will occur. Her sole concern is therefore to defend herself by choosing the action that maximizes this worst outcome. Randomization can be useful as the worst outcome is defined in expected terms once outcomes have been transformed into von Neumann-Morgenstern utilities. Note that the maximin utility decision rule can be interpreted as resulting from extreme pessimism. However, it cannot be associated to extreme risk aversion, since the degree of risk aversion is already captured in the measurement of utilities or payoffs.

An alternative approach that has recently attracted attention is the *minimax regret* *criterion* (Savage, 1951), axiomatized by Milnor (1954) and also by Stoye (2006). Here the DM does not care about the outcome per se but about lost opportunities. She is worried about not correctly anticipating which state of the world will occur and hence about not making the best choice that can be made ex post. Regret measures the loss due to not making the best choice ex post. Analogous to the maximin utility criterion, the DM fears that she is facing a malevolent Nature which is trying to maximize her regret, and consequently the DM defends herself by choosing the (mixed) action that minimizes maximal regret. Randomization is typically beneficial in this defence. Note that, while regret is defined in terms of hindsight, one should not interpret this as a model of a DM who lives in the past since this DM is assumed to anticipate possible future regret when making choices. Anticipation of aversion to lost opportunities finds support in experiments by Zeelenberg (1999)[1].

Both the maximin utility criterion and the minimax regret criterion have been given behavioral interpretations. However their foundations are purely axiomatic. In particular, the application of minimax regret does not depend on whether or not there is information ex post about which state occurred. Both criteria should be seen in light of their axiomatic foundations. The key departure from SEU is founded in the Symmetry Axiom. SEU allows via the prior for different states to be treated differently. The Symmetry Axiom rules out this possibility as it postulates that choice may not depend on labels. The underlying idea is that the definition of the decision problem must include all relevant aspects. If the decision maker nevertheless would like to make a choice that is not invariant to the relabelling of states and actions, then this would contradict the postulate that the definition of the decision problem captures all relevant aspects. Both the maximin utility criterion and the minimax regret criterion satisfy the Symmetry Axiom together with an additional convexity axiom which is associated to ambiguity aversion. According to the latter, when indifferent between two actions, the DM prefers to randomize between them in order to better protect against uncertainty. The maximin utility criterion and the minimax regret criterion differ in terms of which axioms of SEU are relaxed. Recall that Independence of Irrelevant Alternatives (IIA) postulates that preferences are not allowed to change if new actions are added. This should not be confused with the Independence Axiom that is used to enable rearranging mathematical terms and is associated to time consistent choice. The maximin utility decision rule satisfies IIA but only a weaker version of the Independence Axiom. The

[1] Minimax regret should not be confused with other forms of regret to incorporate lost opportunities when facing risk such as Bell (1982) and Loomes and Sugden (1982).

added thrust of the Symmetry Axiom together with ambiguity aversion embedded in the convexity axiom focuses attention on the worst outcome generated by each action. The DM seems to be extremely pessimistic. However, once IIA is relaxed, the outcome is less extreme despite Symmetry and ambiguity aversion. The minimax regret criterion satisfies Symmetry, ambiguity aversion, Independence and the following weaker version of IIA. Preferences are now allowed to depend on the set of actions available, a property called *menu dependence*. In order to achieve a form of consistent behavior across different sets or menus, the Independence to Never Best Alternatives (INA) axiom is postulated. Actions may be added without changing preferences as long as they do not change the outcome of an omniscient decision maker who knows which state will occur. In other words, the best outcome in each state cannot be changed. This invariance to situations that do not affect the well being of such an omniscient decision maker turns the focus to the best outcome in each state and thus leads to concern for regret.

We will focus on these three ways to deal with uncertainty: reversion to risk, extreme pessimism and aversion to lost opportunities.

Let us sum up the philosophy of choice behind each of the three decision criteria presented above. The Bayesian (or 'cool') decision maker always addresses uncertain settings by forming priors and making the necessary expected utility calculations. The maximin utility (or 'gloomy') DM has a highly conservative attitude towards uncertainty and acts as if the worst state of nature were certain to occur, while the minimax regret (or 'anxious') DM is concerned with not missing advantageous opportunities.

We proceed by comparing the choices of these three different decision makers in three specific settings, each focusing on a particular aspect of uncertainty. We investigate the benefits of safe actions, measure the freedom of choice and capture the value of information. Each of the settings points to a most basic behavioral trait. In a later section we will use our analysis to uncover underlying motives by looking at the answers to some of the questions posed in the European Values Survey.

3 Uncovering Models of Choice under Uncertainty

3.1 Safe versus Uncertain

At the heart of choice under uncertainty is the question of how one trades off making a safe choice against making a choice that yields an uncertain outcome. We investigate how the evaluation of uncertainty influences this tradeoff. Consider the following simple

payoff	α	β
safe	c	c
uncertain	l	h

Table 1: The 'safe vs. uncertain' decision problem.

decision problem with two actions labelled 'safe' and 'uncertain' and two states labelled α and β shown in Table 1 where $h > c > l$. Note that this can also be interpreted as a tradeoff between risk and uncertainty. One can imagine that c is the expected value of a risky outcome to be contrasted with the choice of 'uncertain' that has a truly uncertain outcome. We derive the most preferred choice for each of the three different decision makers.

The Bayesian DM assesses a subjective probability for the occurrence of each of the two states. If the bad state for the uncertain choice, state α, is sufficiently likely then the Bayesian chooses the safe action. More specifically, if μ denotes the probability that state α is believed to occur then the Bayesian chooses the safe action if $\mu > (h-c)/(h-l)$ and chooses the risky if $\mu < (h-c)/(h-l)$. In particular, if $\mu \neq (h-c)/(h-l)$ then the Bayesian will not randomize. Typically one does not even consider the case where $\mu = (h-c)/(h-l)$ as this prior is considered degenerate. However, even if this knife-edge case is considered, one would not predict that the Bayesian randomizes albeit with no specific formal reasoning.

A maximin utility DM always chooses 'safe'. This is immediate when considering only pure actions since $c > l$. However it is also true when one includes all mixed actions. This can be seen by considering the fictitious zero sum game between the DM and the malevolent Nature, in which the objective of DM is to maximize utility while that of Nature is to minimize it. It is well known (von Neumann and Morgenstern, 1947) that the equilibrium strategy of the DM in this fictitious game solves the maximin utility criterion. Now note that the pair (safe, α) is an equilibrium of this zero sum game. Hence 'safe' attains maximin utility. Here we see the pessimism of the maximin utility criterion at work. Regardless of how large l or h, as long as $l < c$, this DM always chooses 'safe'.

To analyze the choice of the minimax regret DM we transform utility into regret (see Table 2). The unique mixed action that attains minimax regret involves choosing 'safe' with probability $(c-l)/(h-l)$ and 'uncertain' with probability $(h-c)/(h-l)$.

regret	α	β
safe	0	$h-c$
uncertain	$c-l$	0

Table 2: The 'safe vs. uncertain' decision problem in terms of regret.

In particular, there is always a strictly positive probability of choosing 'uncertain'. The minimax regret criterion trades off the magnitudes of possible loss $(c-l)$ against possible gain $(h-l)$. A small probability is put on 'safe' if and only if the ratio of possible loss to possible gain is small. The value of minimax regret is equal to $(h-c)(c-l)/(h-l)$. Note here the advantage of mixing, which guarantees regret to be strictly below the maximal regret of 'safe' equal to $(h-c)$ and of 'uncertain' equal to $(c-l)$.

To summarize, the tradeoff of the Bayesian is embedded in the prior. While some Bayesians will choose 'safe' others will choose 'uncertain'. The conservatism of the maximin utility DM leads her to always choose 'safe'. Fear of missing advantageous opportunities causes the minimax regret DM neither to choose 'safe' nor 'uncertain' but instead to randomize between these two actions.

3.2 Freedom of Choice

We investigate the tradeoff between 'freedom of choice' and giving the responsibility for choice to others. We wish to understand how the model of choice for facing uncertainty influences the value of being allowed to choose in contrast with letting someone else choose and following their instructions. The latter situation can arise through delegation or by entering a relationship in which one no longer makes the choice. We build our model around the following simple decision problem. There are two actions and two states with a unique best action in each state and where different states have different best actions. Actions are labelled A and B, states labelled α and β, and payoffs shown in Table 3. We impose $y > x$ and $z > w$ to create a different best action in each state and assume further $x < z$ and $w < y$ in order to ensure that not all outcomes in one state are smaller than all outcomes in the other state.

Let us embed this choice between two actions in the following larger decision problem that has two stages. In the first stage, the decision maker has to decide whether she wishes to make the choice herself or whether instead she prefers to have the choice

payoff	α	β
A	y	w
B	x	z

Table 3: The decision problem faced in stage 2 when choosing 'free' in stage 1.

payoff	α	β
'free'	$\frac{yz-xw}{y-x+z-w}$	$\frac{yz-xw}{y-x+z-w}$
'third'	$\lambda y+(1-\lambda)\,x$	$\lambda w+(1-\lambda)\,z$

Table 4: The decision problem faced in stage 1 by a maximin utility agent.

made by a third party. In the second stage the actual choice between A and B occurs. At the time of the first stage the DM believes that the third party will choose A with probability λ and B with probability $1-\lambda$ for some specified $\lambda \in [0,1]$. Let 'free' denote the choice of the DM at stage one to retain the role of choosing an action at stage two. Let 'third' denote the choice at stage one to let the third party choose at stage two.

The analysis for a Bayesian DM is straightforward. Generically this DM will not be indifferent between the two actions, and, as the third party randomizes, she will strictly prefer 'free' and thus to retain the power to choose the action.

Consider now maximin utility. Here one has to specify how decision making takes place in this sequential setting. A natural approach (for an axiomatization see Siniscalchi, 2006) is to solve via backwards induction. If the decision maker chooses 'free' in stage one, then she faces the decision problem in Table 3 in stage two where she will then choose A with probability $(z-x)\,/\,(y-x+z-w)$ and B otherwise to then guarantee a minimal utility of $(yz-xw)\,/\,(y-x+z-w)$. Notice that the choice of this mixed action yields the same expected utility in both states. Anticipating the outcome obtained in stage 2 when choosing 'free' in stage 1 yields the reduced decision problem in stage 1 shown in Table 4. It then follows easily along the same line of argument as in the safe versus risky model that this decision maker chooses 'free'.

Consider now the minimax regret DM. Looking again first at the decision in stage 2 after having chosen 'free' in stage 1 we find that the DM chooses A with probability $(y-x)\,/\,(y-x+z-w)$. Anticipating this choice in stage 2 we obtain the reduced form for stage 1 shown in Table 5. For whichever value of λ in $[0,1]$ we find that the

payoff	α	β
'free'	$\frac{(y-x)y+(z-w)x}{z-w+y-x}$	$\frac{(y-x)w+(z-w)z}{z-w+y-x}$
'third'	$\lambda y+(1-\lambda)\,x$	$\lambda w+(1-\lambda)\,z$

Table 5: The decision problem faced in stage 1 by a minimax regret agent.

minimax regret DM will mix between 'free' and 'third'.

To sum up, both a Bayesian DM and a maximin utility DM have a preference for exerting their freedom of choice, while a minimax regret DM, in randomizing across the two choice options, exhibits a weaker preference towards freedom of choice.

3.3 Value of Information

Information plays an important role when making choices under uncertainty as more information can reduce this uncertainty. In the following we investigate the incentives for the gathering information when this is costly.

We build on the decision problem with two actions and two states presented in the previous subsection. To simplify notation we normalize payoffs so that $x = 0$ and $y = 1$, which can be done without loss of generality. We add the restriction that $w < 1$ and $z > 0$ in order to rule out that all outcomes in one state are larger than all outcomes in the other state. We add the possibility to learn more about the true state as follows. By incurring a cost c the DM learns the true state with probability γ and does not learn anything new with probability $1-\gamma$, where $c > 0$ and $\gamma \in (0, 1)$ are given. Hence, even when the DM pays c she may not learn anything. The strategy of the DM who decides not to purchase information and to choose action C is denoted by Cn. The strategy to buy information, to choose the best action whenever the true state is revealed and to choose action C when no new information is revealed is denoted by Cb. Payoffs of the enlarged decision problem are given in Table 6.

Consider a Bayesian decision maker who puts prior probability μ on state α occurring. It is easily shown that this DM has the highest willingness to pay for information when the prior $(\mu, (1-\mu))$ is such that she is indifferent between An and Bn. This indifference holds when $\mu = \mu_0 := (z-w)\,/\,(1+z-w)$. Note that when $\mu = \mu_0$ then the Bayesian is also indifferent between Ab and Bb. So the expected payoff to not buying information is equal to

$$\frac{z}{1+z-w},$$

payoff	state α	state β
An	1	w
Bn	0	z
Ab	$1-c$	$\gamma z+(1-\gamma)w-c$
Bb	$\gamma-c$	$z-c$

Table 6: The choice setting revealing the value of information for the decision maker. The DM is allowed to acquire at cost c a probability γ of learning the true state.

while the payoff to buying information is equal to

$$\frac{z+\gamma(z-w)-c(1+z-w)}{1+z-w}.$$

For this particular prior it follows that the Bayesian will buy information if $c < \gamma(z-w)/(1+z-w)$. Thus, for any prior μ on state α there exists $c_0(\mu)$ such that the Bayesian DM will buy information if $c < c_0(\mu)$ and she will not buy information if $c > c_0(\mu)$ where

$$c_0(\mu) \leq \frac{\gamma(z-w)}{1+z-w}.$$

In fact it is easily shown that $c_0(\mu)$ is strictly below this threshold whenever $\mu \neq \mu_0$, where $c_0(\mu)$ can be arbitrarily small if μ is either sufficiently large or sufficiently small.

Let us now turn to a maximin utility DM. Assume that this DM decides not to buy information. Since $w<1$ and $z>0$ it follows that the DM will mix between An and Bn in the same way she would mix between A and B in the original decision problem. Both An and Bn will maximize expected payoffs given the mixed action of malevolent Nature. Nature will then choose state α with probability μ_0, as only then will the DM be indifferent between An and Bn. Given this strategy of Nature, we determined above that An and Bn are only best responses if and only if $c \geq \gamma(z-w)/(1+z-w)$. Thus, the maximin utility DM will buy information if

$$c < \frac{\gamma(z-w)}{1+z-w},$$

and will not buy information if instead the above holds with "$>$".

Finally consider the minimax regret DM. In Table 7 we have transformed utility into regret.

We argue analogously to the case of maximin utility. If the minimax regret DM does not buy information then she mixes between An and Bn. Note that Nature again

regret	states α	β
An	0	$z-w$
Bn	1	0
Ab	c	$(1-\gamma)(z-w)+c$
Bb	$1+c-\gamma$	c

Table 7: The 'value of information' choice setting in terms of regret.

ensures that she will do so by assigning probability μ_0 to state α. Thus we conclude as in the case of maximin utility that the minimax regret DM buys information if

$$c < \frac{\gamma(z-w)}{1+z-w},$$

and will not buy information if instead the above holds with ">".

To summarize, the Bayesian is always less willing to buy information than either the maximin utility or the minimax regret type, where the latter two have the same threshold on costs below which they start buying information. The prior of the Bayesian makes her more confident about the situation and hence less willing to pay for more information.

3.4 Summary

Combining the different attitudes across these three decision problems we identify the following behavioral patterns. Bayesians may or may not like safe choices, enjoy being unconstrained in their choices and give relatively low value to acquiring information. Maximin DMs like safe options, like freedom and value information highly. Finally, minimax regret DMs are willing to partially engage in uncertain scenarios, moderately like freedom and give high value to new information.

4 Estimating Models of Choice under Uncertainty

We now use our models and their predictions across the three choice problems developed above to investigate the different ways in which Europeans deal with uncertainty. The data we use is based on answers to some questions posed in the European Values Survey (EVS). The hypothesis is that there are at least three behavioral types, the Bayesian, the maximin utility DM and the minimax regret DM. We proceed by establishing a link between the decision problems of the previous section and some of the questions in the survey, which allows us to predict the answers of each of the three DM types across these questions. This link is then brought to the data using a latent class analysis in which the proportions of each type are estimated.

4.1 Linking Choice Problems to Questions

In the following we identify five questions from the European Values Survey (EVS) that are related to each of our choice problems. In particular two questions are associated with the 'safe vs. uncertain' model, two to the 'freedom of choice' model, and one to the 'value of info' model. For clarity only summaries of the questions are presented here. We have included all the details in Appendix A.

Consider first the conflict between safe and uncertain. The first question in this group, coded C013 in the EVS, is taken from the section of the questionnaire devoted to 'work'. It asks: "Would you mention 'job security' as an important aspect in a job?". The possible answers are 'mention' and 'not mention'. We identify job security as being a safe option that is implicitly contrasted in this question to job insecurity, an option with an uncertain outcome. Mentioning job security can be directly interpreted as not wanting to lose one's job. However job security also indirectly means to forego opportunities of getting a better job. In light of the analysis of the safe versus uncertain model, we predict how each of the three types will answer each question. Some Bayesians will mention good job security while others will not. Maximin DMs will always mention job security. Minimax regret DMs will always hedge against uncertainty by randomizing, so some will mention and some will not mention job security. However, since we cannot observe this randomization we obtain the same prediction for the minimax regret DM as we have for the Bayesian. Not all will answer 'mention' but not all will answer 'not mention'.

The second question associated with the conflict between safe and uncertain, coded D026 in the EVS, is selected from the group of questions devoted to 'family'. It reads: "A marriage or a long-term stable relationship is necessary to be happy". We group the possible answers into 'agree', 'not agree' (see Appendix A for details). Paralleling the explanation given for the previous question, we identify the long-term stable relationship as the safe option, to be implicitly compared to the more uncertain life scenarios unfolding when opting for short-term and/or non stable relationships. Bayesians may agree or disagree with the statement, maximin utility DMs always agree, while minimax regret DMs randomize across the answer options and, hence, are associated to the same prediction as that for Bayesians.

Freedom of choice is addressed in two questions, respectively selected from the 'work' section of the questionnaire and from the one devoted to 'perceptions of life'. The former, coded C016, asks: "Would you mention the 'opportunity to use your own initiative' as an important aspect in a job?". The possible answers to this question are 'mention' and 'not mention'. Following our three choice models we predict that both Bayesians and maximin utility DMs like freedom of choice and hence will choose 'mention', while for minimax regret DMs some will choose 'mention' and others will choose 'not mention'.

The other question in the survey capturing a taste for freedom is the one coded A029 and asking: "Would you mention 'independence' as an especially important quality that children should be encouraged to learn at home?". The possible answers are 'mention' and 'not mention'. Once again, given our three choice models, we predict that both Bayesians and maximin utility DMs will answer 'mention', while for minimax regret DMs some will choose 'mention' and others will choose 'not mention'.

Finally consider the following question, selected from the 'work' section of the questionnaire and coded C012: "Would you mention 'not too much pressure' as an important aspect in a job?". The possible answers are 'mention' and 'not mention'. Pressure can be interpreted as the feeling of stressful urgency (Oxford Dictionary) associated to the lack of time to be able to deliberate and to gather more information. Lower value of information can mean that the decision maker is less concerned about pressure. We focus on the value of information aspect of this question and, according to our model, predict that a Bayesian does not feel much pressure and is likely to choose 'not mention' as information has less value. Both a maximin utility DM and a minimax regret DM care more about information, so we predict that they feel more pressure and hence are likely to choose 'mention'.

Table 8 summarizes the predicted answers from our three decision types across the five questions.

4.2 Empirical Analysis

We now wish to estimate how many individuals behave consistently with each of our behavioral models.

Group	Question	Content	Bayesian	Maximin	M. regret
safe	C013	Job security	some mention, some not	mention	randomize
safe	D026	Long-term relationship	some agree, some not	agree	randomize
free	C016	Initiative	mention	mention	randomize
free	A029	Independence	mention	mention	randomize
info	C012	No pressure	not mention	mention	mention

Table 8: Theoretical predictions

4.2.1 Selecting the Method

We need to select a method for estimating the existence of such underlying behavioral models in the EVS data. Flexibility of the method is desired as we do not expect the vast majority to behave entirely according to one of the three models. Our models are normative and many different concerns come into play when interpreting a question and selecting an answer. We are happy to be able to explain regularities and tendencies. Thus we wish to choose a method that allows different questions to be assigned different degrees of importance in explaining regularities. Our behavioral models should be able to compete with alternative systematic ways of responding to the questions. At the same time we need to be able to correct for the fact that a less restrictive definition of a type will always explain the data more accurately. We choose to perform a latent class analysis as this gives us the desired flexibility and the means to investigate the role of degrees of freedom in explaining the data.

Previous investigations on European attitudes towards uncertainty (Hofstede, 2001) have relied on counting the number of answers consistent with each behavioral model. Each question is given equal weight. However, the questions asked in surveys are just approximations of the question we are really interested in ("Do you mind ambiguity?" or "What is your type?"). To ignore that some questions are better approximations than others is to ignore that the explanatory power of the questions may vary. For example, consider the extreme case that one of the included questions has no explanatory power. By just counting the correct answers this irrelevant question has the same impact as the other questions. Our method instead does not impose equal explanatory power of the questions but determines their relevance endogenously.

Latent class analysis is a method to find classes (or clusters) in the data when the relative importance of the questions is not known[2]. In our model a class represents a DM type. In addition to our three types we add an unrestricted (or *free*) type in order to pick up regularities not predicted by our models. Individuals belonging to a given type are assumed to answer any given question according to a probability distribution that only depends on the type and on the question. The proportions of types as well as the probability distributions over the answers for each question and type are estimated by maximizing the likelihood of the data. So the estimated behavior within a class (or associated with a type) is a distribution of answers to each question. Allowing for probabilitistic distributions over the set of answers introduces the desired flexibility to be able to capture the importance of each question for each type. Flexibility is limited by the bounds imposed by the prediction we make for the given type. For instance, when the type imposes that one of two answers is more likely than the other and one estimates that the two probabilities coincide, then we find that this question plays no special role in explaining responses.

We seek to compare attitudes towards uncertainty across Europe. To this end we pool the data across all countries and estimate behavior of each type as well as the proportions across Europe. Using the country specific information we then derive the induced distribution of types for each country. A separate analysis for each country would not be useful for this objective as estimated types identified with a probability distribution for each question would then be difficult to compare across countries.

4.2.2 Details of the Method

In the following we briefly describe latent class analysis assuming, first of all, that data comes from a single country. Let K denote the set of DM types. Denote the sample proportion of type $k \in K$ by π_k. Of course, $\pi_k \in [0,1]$ and $\sum_{k\in K} \pi_k = 1$.

Let Q be the set of questions. For question $q \in Q$, let A_q denote the set of possible answers. We assume that, conditional on the type of individual, the answers across questions are uncorrelated. Let the probability that a k-type individual answers question q with answer $a \in A_q$ be denoted by $\pi_{kq}(a)$ where $\pi_{kq}(a) \in [0,1]$ and $\sum_{a\in A_q} \pi_{kq}(a) = 1$.

Let π be the parameter vector containing the sample proportions π_k for $k \in K$ and the probabilities of the answers $\pi_{kq}(a)$ for $k \in K$, $q \in Q$ and $a \in A_q$.

[2]Lazarsfeld (1950), Goodman (1974) and Haberman (1979) are the classic references. See Hagenaars and McCutcheon (2002) for an overview of recent innovations.

Up to this point a type is simply an index. However each type is typically associated to an exogenously imposed set of restrictions, which determine for each question a set of possible answers or distributions over answers. The restrictions we impose for our three behavioral types are specified in the next section. Let Π be the set containing the allowed parameter vectors π that are compatible with the exogenously imposed restrictions on the types.

Denote the set of individuals by I. Let the answer of individual i to question q be denoted by $a_{iq} \in A_q$. The loglikelihood of the sample $\{a_{iq}\}_{i\in I, q\in Q}$ of answers for each individual across all individuals is then given by

$$\mathcal{L}(\pi) = \sum_{i\in I} \log\left(\sum_{k\in K} \pi_k \prod_{q\in Q} \pi_{kq}(a_{iq})\right).$$

The maximum likelihood estimator is then simply $\arg\max_{\pi\in\Pi} \mathcal{L}(\pi)$.

When the sample consists of all countries, a correction is needed to preserve representation due to the different country sizes. Ideally one would like to have a sample that consists of the same proportion of individuals relative to the total population in each country. In order not to throw away observations, one considers averages when too many individuals have been sampled. The adjustment is as follows. Denote the set of countries by C. For country $c \in C$, let the set of individuals be denoted by I^c, the number of individuals in the sample by n_c and the number of inhabitants from which the sample was (potentially) drawn by p_c. Let $r = \min_{c\in C} n_c/p_c$ be the minimum ratio of individuals per inhabitant. It then follows that $rp_c \leq n_c$ for all $c \in C$. We now act as if the proportion r was selected from each country, rescaling the above likelihood by rp_c/n. With this correction term the loglikelihood becomes

$$\mathcal{L}(\pi) = \sum_{c\in C} \frac{rp_c}{n_c} \sum_{i\in I^c} \log\left(\sum_{k\in K} \pi_k \prod_{q\in Q} \pi_{kq}(a_{iq})\right).$$

All test statistics are corrected for the population size likewise.

4.2.3 Imposing the Types

Each class is associated with a type. We have developed predictions for the behavior of each of the three models in dealing with uncertainty (see Table 8). In the following we show how these predictions enter the empirical model to be estimated.

If an individual is of a particular type, we expect her to answer the questions in line with that type. However, the vast majority of individuals will not answer the

Question	Content	Bayesian	Maximin	M. regret
C013	job security	$\frac{1}{4} \leq \pi(M) \leq \frac{3}{4}$	$\pi(M) \geq \frac{1}{2}$	$\frac{1}{4} \leq \pi(M) \leq \frac{3}{4}$
D026	long-term relation	$\frac{1}{4} \leq \pi(A) \leq \frac{3}{4}$	$\pi(A) \geq \frac{1}{2}$	$\frac{1}{4} \leq \pi(A) \leq \frac{3}{4}$
C016	initiative	$\pi(M) \geq \frac{1}{2}$	$\pi(M) \geq \frac{1}{2}$	$\frac{1}{4} \leq \pi(M) \leq \frac{3}{4}$
A029	independence	$\pi(M) \geq \frac{1}{2}$	$\pi(M) \geq \frac{1}{2}$	$\frac{1}{4} \leq \pi(M) \leq \frac{3}{4}$
C012	no pressure	$\pi(M) \leq \frac{1}{2}$	$\pi(M) \geq \frac{1}{2}$	$\pi(M) \geq \frac{1}{2}$

Table 9: Conditions to impose on types

questions in a way that exactly coincides with one of the DM types. As we said above, our estimation method assumes that a DM type has a probability distribution over the answers of each question. If an individual is of a particular type, we thus expect that she answers the questions in line with that type with a *high probability*. In other words we predict the answer of the majority.

We now derive the restrictions imposed on the answer frequencies for each of the three behavioral types. Consider first the case where our model reveals a unique prediction, for example a maximin utility DM is predicted to answer 'mention' to the job security question. To allow flexibility, so that different questions can have different degrees of predictive power, we contrast our predictions to a hypothetical random DM who is equally likely to choose all answers. So the random DM chooses 'mention' (M) and 'not mention' (NM) with the same probability. The maximin utility DM should then outperform the random DM and answer 'mention' with a probability of at least 1/2. In other words, we predict that there are more maximin utility DMs answering 'mention' than 'not mention'. We also have to specify how we deal with 'ambiguous' predictions, for instance with Bayesians' attitude towards job security. We want to rule out that our estimation allows for all Bayesians in Europe to like job security or for all to dislike job security. Hence, we impose heterogeneity by requiring that the proportion of Bayesians liking job security lies between 25% and 75%. In this fashion we translate each of our theoretical predictions (see Table 8) into restrictions on the parameter set Π shown in Table 9 (where A stands for 'agree').

5 The Findings

We now show the results of the maximum likelihood estimation on the pooled data of the EU15 countries. We have estimated both the 'unconstrained' model, that is the model only made up of free types, and the 'constrained' model, that is the model obtained by imposing the conditions on the types specified in Table 9. Moreover, assuming that Bayesians (B), maximin utility (M) or minimax regret (R) are the only DM types in Europe might be too strong. We therefore also perform the estimation of the constrained model with one free type (F). The model with four free types is also estimated.

Models statistics are summarized in Table 10, where C($\cdot$) and U($\cdot$) respectively denote the constrained and the unconstrained models, while $3, 4$ stand for the number of types included in the estimation.

Statistic	Model			
	C(3)	C(4)	U(3)	U(4)
Akaike	58.0	−12.2	−10.1	−9.1
Bayesian	−53.2	−97.2	−101.6	−61.4
Pearson	93.6	13.8	17.7	6.9
Likelihood Ratio	92.0	13.8	17.9	6.9
$\chi(0.95)$	27.6	22.4	23.7	15.5
df	17	13	14	8

Table 10: The statistics for the constrained and unconstrained models with three and four types. In addition the table shows the 95% critical values and the degrees of freedom.

For the unconstrained model, with both three and four types, the test statistics are within the 95% confidence range, suggesting that the model performs fairly well in explaining the data. The constrained model with only three types has instead rather high Pearson and Likelihood Ratio values, which are known to be conservative measures for large samples, and is acceptable only under the Bayesian criterion, which favors small models in terms of parameters. However, by adding one free type, the constrained model fits the data accurately according to each of the four criteria. Moreover, the Akaike criterion actually prefers the constrained model with four types to the unconstrained models with either three or four types, while for the Bayesian criterion this constrained model performs better than the unconstrained model with four types and only slightly worse than the unconstrained model with three types. Overall we find that the data can be explained by imposing the behavioral conditions on the types dictated by our normative models.

We will now focus on the constrained model with four types. Let us proceed with the analysis of the parameters estimated. The first row of Table 11 reports the types of decision makers included in the estimation, the last row contains their sample proportions for all of Europe. We find 69% of the answers can be explained by one of our three models, the majority, equal to 39%, fall within the Bayesian model. Maximin utility and minimax regret types account for 23% and 6% respectively. The rest of the sample, 31%, is captured by the free type.

Answers M/A in %	Bayesian	Maximin U.	M. regret	Fourth	Avg
Job security	59	98	25*	58	62
Long-term relation	49	61	25*	67	53
Initiative	50*	79	48	26	51
Independence	76	50*	57	16	51
No pressure	13	87	50*	10	32
Proportion	39	23	6	31	

Table 11: The estimation results for the model with four decision maker types. For each decision maker type the table shows its population proportion and the percentage probabilities of answer ‘agree’ (A) to the question on ‘long-term relation’ and of answer ‘mention’ (M) to the other four questions.

The other rows of the same table show the answer percentage probabilities of the four types for each question. In particular, the row associated with question on long-term relationship reports the probability for each type of answering ‘agree’ (A), while the other rows report the probabilities for each type of answering ‘mention’ (M) to the other four questions. Needless to say, subtracting these numbers from 100 gives us the percentage probabilities of answering respectively ‘not agree’ and ‘not mention’ for each type. Asterisks are added to indicate that the constraints are binding. It is worth comparing the probabilities in Table 11 with the answer frequencies for the whole sample reported in the final column of the same table, which can be interpreted as describing the ‘average’ type. For instance, consider the question about preference for ‘not too much pressure’ at work. Around 32% of the sample do mention ‘not too much pressure’ as an important aspect in a job (that is, answer M), while the remaining 68% do not mention it. We obtain rather extreme and opposite predictions about the behavior of Bayesian and the free type on the one hand, and the behavior of the maximin utility type on the other. The estimated value for the minimax regret type is instead binding at the boundary imposed by predicting that this type chooses M with probability greater than or equal to 50%. Overall, allowing for different behavioral types gives us the possibility of capturing the behavioral heterogeneity behind the answers to this question. Roughly, the same reasoning holds for the other four questions, for which our types perform reasonably well in uncovering the heterogeneity in the answers. At the same time our types impose limits to this heterogeneity and the constraints are binding in 1/3 of the cases with most constraints binding for the minimax regret type.

As mentioned above, latent class analysis allows us to give different weights to different questions. For instance, while we include the questions on job security and on long-term relationship both within the safe versus risky model, we allow for different degrees in which the value of safe options influences the answers to these questions. For the maximin utility type we find a very close alignment between mentioning job security and the value of safe options as indicated by the value of 98%. On the other hand, long-term relationship and safety seem to be less aligned given our estimated 61% of choosing M.

Consider the degree to which our estimates correspond to our predictions. Looking at the number of binding constraints we see that the Bayesian and maximin type predict fairly well across questions, each with one binding constraint respectively on ‘initiative’ and ‘independence’ questions. Minimax regret behavior only predicts well for the questions within the freedom of choice category.

Finally Table 11 also allows us to delineate the behavioral traits of the free or fourth type. This type neither likes nor dislikes safe options, dislikes freedom of choice and does not mind pressure and thus, in our interpretation, places a relatively low importance on acquiring additional information. One may interpret this type as one who does not use initiative and, hence, does not care about either the availability of safe options or the pressure that may be present if she had to make a choice.

Table 12 reports the sample proportions of the decision maker types for each country in EU15. The table can also be interpreted as providing the average relative probabilities that an individual belongs to each of the decision maker types. In Figures 1 and 2 we plot the frequencies of, respectively, the maximin utility types against the

Bayesians, and the free against minimax regret types.

Country	Type			
	B	M	R	F
Austria	54.0	15.4	4.9	25.7
Belgium	36.9	17.7	8.5	37.0
Denmark	61.9	10.6	7.9	19.6
Finland	44.7	22.4	7.8	25.0
France	35.8	9.2	5.7	49.3
Germany	53.1	18.8	4.5	23.5
Greece	32.3	37.4	5.4	24.8
Ireland	34.9	34.0	7.4	23.6
Italy	25.2	45.7	6.0	23.1
Luxemburg	38.6	23.3	7.0	31.1
Netherlands	44.9	13.6	16.9	24.7
Portugal	27.7	20.4	4.8	47.1
Spain	27.4	29.4	5.4	37.8
Sweden	47.1	20.7	10.1	22.2
United Kingdom	41.7	20.3	7.6	30.4
EU15	39.3	23.2	6.4	31.1

Table 12: The sample proportions of the decision maker types for the constrained model with four decision maker types. The table shows the average relative probabilities that an individual belongs to each of the decision maker types.

While Bayesians average 39.3% of the entire sample, their proportions range from a minimum of 25.2% in Italy to a maximum of 61.9% in Denmark. The pattern followed by the proportions across European countries is evident. Southern European countries, namely Greece (32%), Italy (25.2%), Portugal (27.7%) and Spain (27.4%), are all clearly below the average. Most Continental countries, namely Austria (54%), Germany (53.1%) and the Netherlands (44.9%), as well as all Scandinavian countries, namely Denmark (61.9%), Finland (44.7%) and Sweden (47.1%) are instead all clearly above the average. The group of countries 'around' the average is composed of Belgium (36.9%), France (35.8%), Luxembourg (38.6%), as well as Ireland (34.9%) and the UK (41.7%).

Even if less clear-cut and with lower heterogeneity, a similar and reverse pattern emerges when looking at the proportions of maximin utility types. The average proportion is 23.2%, ranging from a minimum of 9.2% in France to a maximum of 45.7% in Italy. A group of southern countries, composed of Greece (32.3%), Italy (45.7%) and Spain (29.4%), plus Ireland (34%), stand clearly above the average. A fraction of Continental countries, specifically Austria (15.4%), Belgium (17.7%), France (9.2%), Germany (18.8%) and the Netherlands (13.6%), plus Denmark (10.6%), present instead relatively low proportions of maximin utility types. Around the average we find an admittedly rather heterogenous group of countries composed of Portugal (20.4%), Luxembourg (23.3%), Finland (22.4%), Sweden (20.7%) and the UK (20.3%).

From these estimates a clear and remarkable difference in the attitude towards uncertainty seems to emerge between southern European countries (plus Ireland to some extent) and continental/northern European countries. Relatively low proportions of Bayesians and high proportions of maximin utility types usually characterize the former group, while exactly the opposite pattern holds for the latter. Given our definitions of the two types, the former group is then populated with agents holding a relatively stronger preference for safe options and a greater aversion towards the feeling of pressure. In contrast, the latter group is characterized by agents more prone to face uncertain scenarios and less sensitive to situations of great pressure. This result, whose graphical intuition is provided in Figure 1, is consistent with the north-south interpretation common in the cross-country psychological literature.

The percentage of minimax regret types is around 6.4% across the whole European sample. The countries with respectively the lowest and highest estimated proportions of minimax regret types are Austria (4.9%) and the Netherlands (16.9%). We find some evidence supporting a weak 'northern pattern' for minimax regret types, as the highest proportions are located among northern countries. However, apart from the Netherlands and Sweden (10.1%), the estimated proportions are mostly concentrated near the average. We in fact find all countries except for the latter two in the range between 4.9% and 7.9%, suggesting that the behavioral pattern corresponding to the minimax regret type is not widespread and is rather homogenous across Europe.

Finally the free type averages 31.1% of the sample and ranges from a minimum of 19.6% in Denmark to a maximum of 49.3% in France. We find a remarkable east-west divide (see Figure 2). The highest values are taken by countries in the west, namely France, Portugal (47.1%), Spain (37.8%), Belgium (37%), Luxembourg (31.1%) and United Kingdom (30.4%). In fact, the fourth type captures majority behavior in the first four countries listed.

It may be worthwhile to interpret the results of our empirical analysis in light of the expanding literature on *ambiguity* (Ellsberg, 1961, Schmeidler, 1989, Gilboa-Schmeidler, 1989). Ambiguity is associated with the way the decision maker confronts an uncertain environment and, in particular, with the lack of prior over the states. Bayesians behave as if they formed priors to transform uncertainty into risk and, hence, deal with an unambiguous scenario. On the other hand, maximin and minimax regret types behave as if they did not have any prior and thus perceive the choice scenario as ambiguous. The distinctive feature of the free type seems to be its aversion towards freedom of choice which, as we claimed above, may also be related to the behavior exhibited in the other two choice contexts, 'safe vs. uncertain' and 'value of information'. The free type can then be classified among the non 'prior-friendly' decision makers to the extent that a natural implication of the prior-based approach is a positive taste for freedom of choice (see Subsection 3.2). The results of our empirical analysis suggest that, in individual decision making under uncertainty, ambiguity plays a prominent role across European countries, and the more so as we move to the south of Europe.

References

[1] Anscombe, F.J. and R.J. Aumann (1963). "A Definition of Subjective Probability". The Annals of Mathematics and Statistics 34, 199-205.

[2] Bell, D.E. (1982). "Regret in Decision Making under Uncertainty". *Operations Research*, 30 (5), 961-181.

[3] Ellsberg, D. (1961). "Risk, Ambiguity and the Savage Axioms", *Quarterly Journal of Economics* 75, 643-669.

[4] European Values Survey (n.d.). http://www.worldvaluessurvey.org/.

[5] Eurostat (n.d). http://www.ec.europa.eu/eurostat/.

[6] Gilboa, I. And D. Schmeidler (1989). "Maxmin Expected Utility with Non-Unique Prior". *Journal of Mathematical Economics* 18, 141-153.

[7] Goodman, L.A. (1974). "The Analysis of Systems of Qualitative Variables when Some of the Variables Are Unobservable. Part I: A modified latent structure approach", *American Journal of Sociology*, 79, 1179-1259.

[8] Haberman, S.J. (1979). “Analysis of Qualitative Data, Vol 2, New Developments”. New York: Academic Press.

[9] Hagenaars, J.A. and McCutcheon, A.L. (2002). “Applied Latent Class Analysis”, Cambridge University Press.

[10] Hofstede, G. (2001). “Culture’s Consequences: Comparing Values, Behaviors, Institutions and Organizations across Nations”. Thousand Oaks, CA: Sage.

[11] Knight, F. (1921). “Risk, Uncertainty and Profit”. Houghton Mifflin Company.

[12] Lazarsfeld, P.F. (1950). “The Logical and Mathematical Foundations of Latent Structure Analysis”, pp 362-412, and “The Interpretation and mathematical Foundations of Latent Structure Analysis”, pp 413-472 in *Measurement and Prediction*, S.A. Stouffer et al. (eds.). Princeton, NJ: Princeton University Press.

[13] Loomes, G. and R. Sugden (1982). “Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty”. *Economic Journal*, 92, 805-824.

[14] Luce, D.K. and H. Raiffa (1958). “Games and Decisions”. New York: John Wiley and Sons.

[15] Milnor, J. (1954). “Games against Nature”. In R. Thrall, C. Coombs and R. Davis (Eds), *Decision Processes.* London: John Wiley, 49-60.

[16] Rustichini, A., J. Dickhaut, P. Ghirardato, K. Smith and V.J. Pardo (2002). “A Brain Imaging Study of Procedural Choice”, University of Minnesota Working Paper.

[17] Savage, L.J. (1951). “The Theory of Statistical Decision”. *Journal of the American Statistical Association* 46, 55-67.

[18] Savage, L.J. (1954). “The Foundations of Statistics”. New York: John Wiley and Sons.

[19] Schmeidler, D. (1989). “Subjective Probability and Expected Utility without Additivity”. *Econometrica* 57, 571-587.

[20] Siniscalchi, M. (2006). “Dynamic Choice under Ambiguity”, mimeo, Northwestern University.

[21] Statistisches Bundesamt Deutschland (n.d.). http://www.destatis.de/.

[22] Stoye, J. (2006). "Statistical Decisions under Ambiguity", mimeo NYU.

[23] von Neumann, J. and O. Morgenstern (1947). "The Theory of Games and Economic Behaviour". 2nd ed. Princeton: Princeton University Press.

[24] Wald, A. (1950). "Statistical Decision Functions". New York: John Wiley.

[25] Zeelenberg, M. (1999). "Anticipated Regret, Expected Feedback and Behavioral Decision Making". *Journal of Behavioral Decision Making* 12, 93-106.

A The EVS Questionnaire

In Subsection 4.1 we have illustrated the questions selected from the EVS questionnaire and the associated answer options. In the following we report the exact way in which they appear in the questionnaire.

Questions C013, C016, C012 are structurally very similar. They all start with the following statement: "Here are some more aspects of a job that people say are important. Please look at them and tell me which ones you personally think are important in a job". Each question is then associated to a specific aspect which may or may not be mentioned. In particular C013 is associated to 'job security', C016 to 'opportunity to use initiative', and C012 to 'not too much pressure'. In the questionnaire there are fifteen more questions structured in this way, each specifying a different aspect. The answer options are $0, 1$ respectively standing for 'not mention', 'mention'. There is no limit in the number of aspects that can be mentioned out of the eighteen presented.

Question D026 asks the respondent how he/she feels about the following statement: "A marriage or a long-term stable relationship is necessary to be happy". The possible answers are $1, 2, 3, 4, 5$, standing respectively for 'agree strongly', 'agree', 'neither agree or disagree', 'disagree', 'strongly disagree'. We combine answers $1, 2$ and answers $3, 4, 5$, label them respectively with 'agree' and 'not agree' so as to have two answer options only. Answer 3 is allocated to 'not agree' to obtain groups of roughly equal size.

Finally question A029 belongs to a group of seventeen questions all starting as follows: "Here is a list of qualities that children can be encouraged to learn at home. Which, if any, do you consider to be especially important?". Each of the seventeen questions mentions a different quality. The one corresponding to question A029 is 'independence'. The answer options are $0, 1$ respectively standing for 'not mentioned', 'important'. The respondent cannot answer 'important' to more than five questions.

B Data Treatment

The data comes from the fourth wave of the European Values Survey. Our country choice is based on the member states of the European Union in this period. In all countries the survey took place in 1999, except for Finland where it was held in 2000.

The total sample consists of fifteen countries and 20,729 individuals (see Table 13). Within a country the individuals are stratified according to geographic population density. The documentation of the EVS provides more information about the sampling

country	original n	relevant n	population	adjusted n
Austria	1,522	1,467	8.0	112
Belgium	1,912	1,669	10.2	144
Denmark	1,023	941	5.3	75
Finland	1,038	1,010	5.2	73
France	1,615	1,582	58.5	823
Germany	2,036	1,784	82.0	1,154
Greece	1,142	1,109	10.9	153
Ireland	1,012	966	3.7	53
Italy	2,000	1,895	56.9	800
Luxembourg	1,211	745	0.4	6
Netherlands	1,003	974	15.8	222
Portugal	1,000	961	10.2	143
Spain	1,200	1,131	39.8	560
Sweden	1,015	962	8.9	125
United Kingdom	2,000	1,698	58.6	824
Europe	20,729	18,894	312.0	5,264

Table 13: The number of individuals in the original sample, the number of individuals after dropping non-citizens and incomplete observations, the population sizes (in millions) and the number of observations adjusted for the country size.

procedures and non-response. See European Values Survey (n.d.).

To get a clearer image of country specific behavior, we drop all individuals who are not citizens of the country they live in. This eliminates 6.2% of the observations from the whole sample, and apart from Luxembourg (37.3%), the United Kingdom (12.5%), Belgium (11.2%) and Denmark (5.3%), less than 5% of the observations at the country level.

For our 5 questions the answer options have been reported in Appendix A. There exist However, four additional possibilities exist for each of the questions, namely: -4 = not asked in the survey; -3 = not applicable; -2 = no answer[3]; -1 = don't know. We have dropped all these data out of our sample. This eliminates 2.7% of the

[3] 'No answer' does not mean that the respondent does not know but that she explicitly prefers not to answer the question.

observations in the whole sample, and apart from Germany (9.6%) less than 5% of the observations at the country level.

Due to historical reasons, individuals living in one of the states of former East Germany are overrepresented in the sample. We use data from the Statistisches Bundesamt Deutschland (n.d.) to construct the population proportion of the combined formerly East German states in 1999. Individuals of these states are then weighted so that their sample proportion equals their population proportion. This reduces the effective number of observations for Germany to 1,154.

In the pooled estimations we weigh individuals so that the effective sample size of a country is proportional to its population size in 1999 (data obtained from Eurostat, n.d.). Germany is the country with the lowest ratio observations per inhabitant. The adjusted number of observations for the other countries follows from multiplying the country populations by this ratio. The effective number of observations when correcting for population size equals 5,264 (rounded).

In order to have a better representation of the country population, we weigh the individuals of a country such that the fraction females in the sample equals the population fraction in 1999 (data obtained from Eurostat). For Germany this weighting is performed for former East and West Germany individually (data obtained from Statistisches Bundesamt Deutschland).

C Test Statistics

For a specific country, the standard Pearson and likelihood ratio statistic measure the difference between the theoretical predictions and the data. Our Pearson and likelihood ratio statistics are the sum of these country specific statistics. Our statistics thus measure the aggregate difference between the theoretical predictions and the data for each country. In our case, each of the five questions has 2 relevant answers. The total number of cells thus equals $15 \times 2^5 = 480$. Note that we have 5,264 effective observations, which is enough for the statistic to be informative.

The degrees of freedom are the number of independent cells for Europe, $2^5 - 1 = 31$, minus the $k-1$ independent type specific proportions, minus the $k \times 5 \times (2-1) = 5k$ independent answer probabilities plus the number of binding constraints.

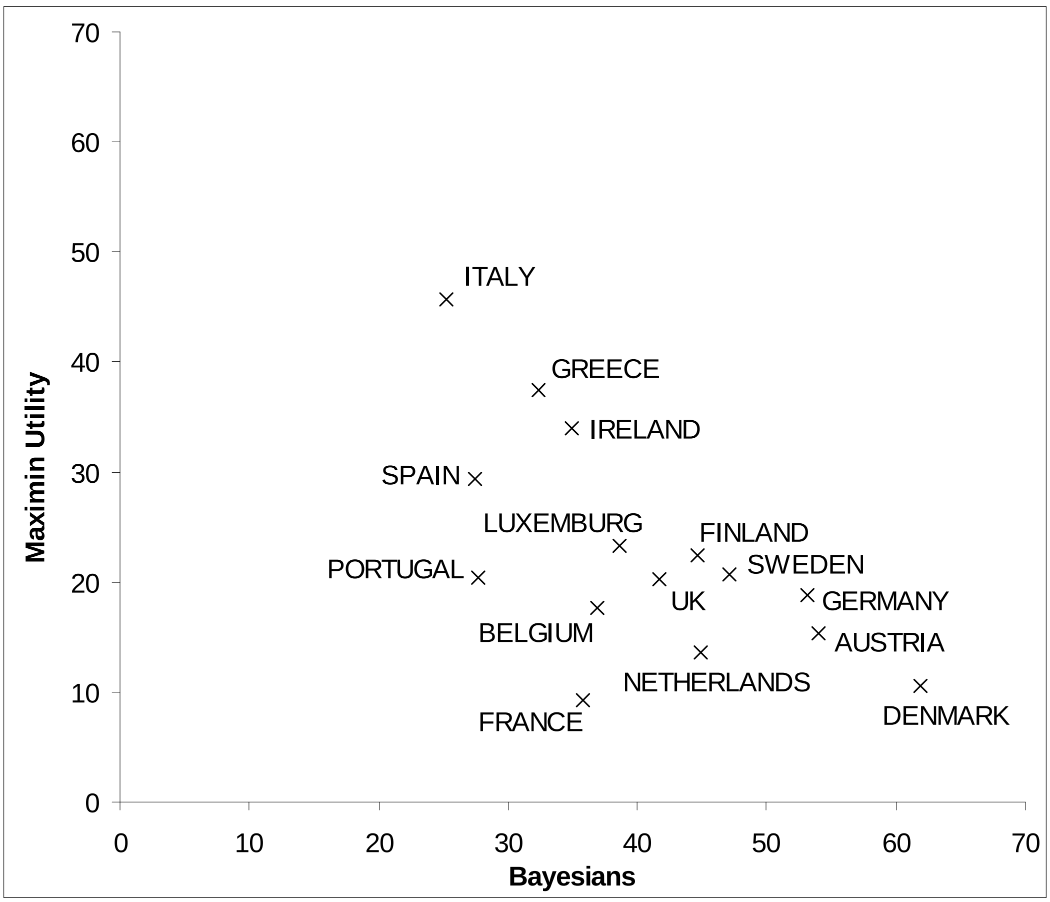

Figure 1: Scatter plot of frequencies of Bayesian and maximin utility type behaviors.

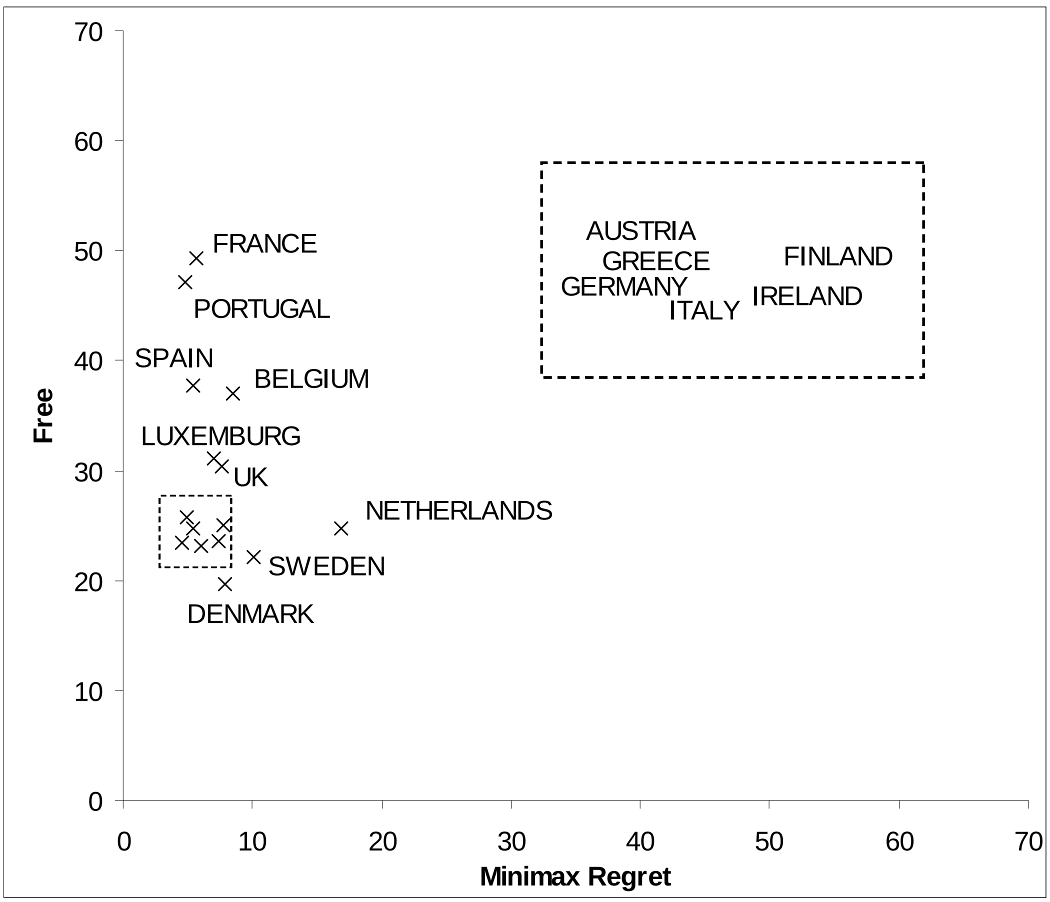

Figure 2: Scatter plot of frequencies of minimax regret and unconstrained type behaviors (labels belonging to the points within the box are contained in the box on the upper right).