

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Ambiguity-Reduction: a Satisficing Criterion for Decision Making

Permalink

<https://escholarship.org/uc/item/8sp8c7w2>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 28(28)

ISSN

1069-7977

Authors

Couyoumdjian, Alessandro
Pezzulo, Giovanni

Publication Date

2006

Peer reviewed

Ambiguity-Reduction: a Satisficing Criterion for Decision Making

Giovanni Pezzulo (giovanni.pezzulo@istc.cnr.it)

Institute of Cognitive Science and Technology - CNR
Via S. Martino della Battaglia, 44 - 00185 Roma, Italy

Alessandro Couyoumdjian (alessandro.couyoumdjian@uniroma1.it)

University of Rome "La Sapienza"
Piazzale Aldo Moro, 9 - 00185 Roma, Italy

Abstract

In the domain of decision making under uncertainty we propose the *Multiple Source Evaluation Model*, focusing on (1) how information coming from different evidential sources, either converging or diverging, is integrated and (2) how Ignorance, Uncertainty and Contradiction are evaluated and reduced before deciding in ambiguous domains. We argue that these operations involve a unique satisficing strategy, *ambiguity-reduction*. We introduce the *Two Cards Gambling Game* experimental paradigm and we present two experiments for exploring the main empirical implications of the model.

Keywords: decision making, uncertainty, ambiguity, heuristic

Introduction

Which strategies do individuals use for deciding in ambiguous and uncertain situations? In the fifties Herbert Simon introduced the term "satisficing" (Simon, 1957), a combination of the words "satisfying" and "sufficing". A satisficing solution is a sub-optimal one, i.e. allowing reasonable approximation of the complexity of reality within given constraints, one that is "good enough" as a basis for deciding. This means that people do not only consider the knowledge in the domain, but also have some prior knowledge and expertise about how to reason within the domain, including which sources to consider, what strategy to adopt and when to stop reasoning and decide. We assume that the procedural expertise of decision making in a domain consists of applying a set of rules, called here *epistemic actions*, which aim mainly at strengthening a belief structure before deciding, for example by reducing uncertainty and ignorance. Some good examples of epistemic actions are the fast and frugal heuristics such as "Take The Best" introduced by Gigerenzer and Todd (1999), but also strategies such as "ask for new information" or "revise reliability of a source".

We will argue in favor of a unique satisficing principle: *individuals tend to select the epistemic action resulting in a more stable basis for deciding in order to reduce ambiguity to an acceptable degree*. This criterium needs a limited amount of knowledge and processing time, without pre-calculating them, according to the desiderata of the "bounded rationality" research program (Gigerenzer & Todd, 1999). Consistently, we will also claim that many coarse-grained heuristics such as representativeness and availability (Kahneman, 2003) are phenomena which emerge from this more fine-grained strategy, depending on contextual factors.

In the first Section we propose a model of decision under uncertainty, the *Multiple Source Evaluation Model* (MSEM), providing a comprehensive account of how decision is based on beliefs and belief sources and which are the most relevant

metacognitive dynamics during decision making. In the second Section we present two experiments using the Two Cards Gambling Game (TCGG) paradigm, permitting to study how information is integrated and exploited under uncertainty: participants can bet on cards shuffled at different velocities, knowing the bet of other Gamblers having variable reliability.

The Multiple Source Evaluation Model

In the judgment and decision making domain several computational models have been recently developed (for a recent review see Busemeyer and Johnson (2004)). Among the most influential ones, there are decision field theory (DFT, Busemeyer and Townsend (1993)) and the Leaky, Competing Accumulator Model (LCAM, Usher and McClelland (2001)). Briefly, DFT is a connectionist model based on psychological rather than economic principles. It consists of a three layer neural network. The connections, linking the inputs (affective evaluations of the possible consequences of a decision) to the first layer represent an attentional weighting process, that selects the options processed at present time. Differently, the connections between first and second layers perform comparisons among weighted values of the options; the outputs of the second layer are the valences, that represent the advantages/disadvantages of each option over the other options. Finally, the connections, between the second and third layers, and the interconnections in the third layer (lateral inhibition), by integrating valences over time, produce a preference state for each option. When one of the preference states reaches a threshold bound, the choice and the deliberation time of a decision are determined.

Even if, in such a way, DFT can simulate various psychological effects (i.e. loss aversion, preference reversal), it, as most of the other computational models, does not explicitly take into account some aspects of cognition that usually influence behaviour: metacognition and background knowledge. Here, we introduce a model of how different sources of information are evaluated and integrated in decision making contexts, extending (Castelfranchi, 1996; Pezzulo, Lorini, & Calvi, 2004). The fundamental claim is that the result of a decision depends on the supporting epistemic structure, composed of: (1) beliefs *in* the domain of decision (base beliefs about information sources and their reliability); (2) beliefs *about* the domain of decision (the meta-beliefs: evaluation of ambiguity in the current decision); (3) background knowledge (declarative and procedural expertise in the domain of decision).

Beliefs in the domain depend mainly on evidences provided by information sources, weighed according to their reliability. These beliefs are also used at the *metalevel* for evaluating the

frame of choice: for example, for an estimation of the levels of ignorance or uncertainty. For example, if two conflicting choices receive the same amount of support, the situation can be evaluated as uncertain. Moreover, top-down influences of *background assumptions*, which we model as the framing of the situation, are a valuable part of the process. According to the *Competence Hypothesis* in (Heath & Tversky, 1991), even with the same data decision makers behave differently depending on how they evaluate their competence. As a result of the metalevel evaluative processes, ambiguity of the situation is quantified. Our assumption is that, before deciding, a decision maker exploits its procedural expertise in order to minimize ambiguity: this typically leads to *epistemic actions*, such as trying to acquire new information, reject an evidence or revise the reliability value of one or more sources, which in turn modify the belief structure at the base level. For example, if a choice is framed as simple, the participant can decide to rely only on the most reliable source (say perception); or, if he has unreliable information, he can decide either to ask for more information or not to bet at all. These epistemic operations over own belief structure are attempts to assess the belief structure, establishing solid and uncontradictory supports for the decision and reducing cognitive dissonance (Festinger, 1957). Thus, according to our model, final decision does not depend only on a base level source evaluation, but also on a metalevel evaluation of the decision context, as well as on the strategy selected for resolving belief conflicts and evidence contradictions. If the final epistemic structure structure is solid (i.e. there are strong reasons to believe) the decision will be rapid and expressed with great confidence; if the structure is less solid (in the presence of uncertainty or ignorance) the decision will be less rapid and prone to errors, mainly because the participant will try to strengthen it (minimizing ambiguity) before deciding.

The Base Beliefs

The knowledge structure of a participant can be described as a network of beliefs, whose edges are “sustain/activation” or “contrast/inhibition” relationships. Beliefs have a certain *strength*, i.e. the degree of confidence of reliability people assign them (Castelfranchi, 1996). We consider many kinds of beliefs and sources, including *inside* evidences, focused on the contingent situation (e.g. perceptual data) and *outside* evidences, focused on categorical knowledge (e.g. previous similar situations, statistical information, etc.); this distinction is also presented in (Kahneman, 2003). For the sake of simplicity, consider as a good analogy a *Balance with two Plates* where evidences are “put on the plate”, each one weighted with its “relevance”, which is proportional to the reliability of the evidential source. Fig. 1 shows a sample network realizing this model with nodes for evidences and (weighted) edges for their influence; it calculates the strength values for “left” and “right” by integrating information from perception as well as from other Gamblers; vertical and horizontal edges represent “sustain” and “contrast” relations. (Pezzulo et al., 2004) describes an implementation of the Balance which is similar to (Usher & McClelland, 2001).

The Meta Beliefs

People not only use information about the problem (e.g. “I have seen that the left card is red”) but also what we call




Figure 1: The Balance with Two Plates

meta-beliefs, i.e. beliefs about the domain, that are an evaluation of their metacognitive state (e.g. “I still do not have enough information”). In many studies “lack of information” as well as various of phenomena related to uncertainty and ambiguity have been shown to affect the decision process; for example ambiguity aversion in subjects has been identified (see Camerer and Weber (1992) for a review). Here we focus on three meta beliefs: *Ignorance, Uncertainty and Contradiction*, showing how crucial their role is in understanding human decision making.

Ignorance can be defined as a measure of the number of information sources; however, here we do not focus on a statistically optimal measure of ignorance, but on *Perceived Ignorance*: how much a participant feels ignorant with respect to a given task and domain. It is mainly related to the *informativeness* of his sources, i.e. how much he estimates new sources to be able to “change his mind”. Informativeness of a source is proportional to its reliability (more reliable sources have higher weights). So, a participant can feel subjectively not ignorant (e.g. in a weather forecast task) even when consulting a limited amount of influential sources (e.g. two reliable forecast TV channels). We will argue that subjective ignorance is a subjective evaluation of actual lack of information on the basis of cognitive evidential models; it is thus a complex cognitive measure and not a simple account of the number of available sources. As a first approximation, in a game with a limited amount of sources (e.g. in TCGG), perceived ignorance is lower if reliability of the sources is high (because this leads to high values in the strength of the involved beliefs in the Balance model). Pezzulo et al. (2004) describe a case study about Ignorance in open worlds.

Uncertainty is a measure of the relative evaluation of many conflicting hypotheses in a given situation: for example, in the TCGG, if two contradictory beliefs (the right card is red, the left card is red) are both strong, there is high uncertainty, i.e. many “good reasons” for betting on each one; another case of high uncertainty is assuming a new belief that contradicts existing knowledge. Instead, uncertainty is minimized when the difference in strength between the hypotheses is high. It is not always the case that having more information is better for choosing; on the contrary, as we will show in the next Section, using an intermediate amount of knowledge can be better for deciding; this is also described by (Gigerenzer & Todd, 1999) as the *less-is-more-effect*.

Contradiction evaluates the incompatibility of beliefs (a special case of uncertainty over a certain threshold). In a cog-

nitive system contradiction is not always strictly logical; two beliefs can be logically consistent but perceived as contradictory. Moreover, it is not binary but graded. In the TCGG, there are always two logically inconsistent hypotheses (and choices): if there are evidences for both, the belief system can become contradictory. In a certain measure contradiction is subjectively perceived, depending on the participants tendency to “choose risking to fail” (this attitude is similar to risk-avoidance in Kahneman (2003) but for epistemic risks); we model it with a subjective contradiction threshold. The contradiction-reduction activity provides feedback on source evaluation (e.g. “I have seen well, the source x is wrong”).

Ignorance, Uncertainty and Contradiction are calculated on the basis of beliefs about the task; for example, in the TCGG, if there are evidences for both right and left, ignorance can be low but uncertainty is high. If there are few evidences, ignorance is high. If the two values are very different, uncertainty is low. Before deciding, these values have to be reduced under domain-specific *Thresholds*.

Background Knowledge

In the literature many heuristics are described (Heath & Tversky, 1991; Camerer & Weber, 1992) which depend on how participants evaluate their competence in a given domain. Interestingly, in “ambiguous” situations (where a participant’s competence is questioned), a participant is uncertainty-avoidant: when choosing between two otherwise equivalent options, one in which the probability information is stated and another in which it is missing, most people avoid the option with missing probability information (Camerer & Weber, 1992). This substantial, frequently replicated tendency is known as the *ambiguity effect*. Instead, in “non-ambiguous” situations (where the participant is confident in himself or about how to choose his sources) there is no uncertainty-avoidance. These evidences account for participants’ ability to evaluate their “ability to decide” in a domain. This competence is highly domain-specific and involves many abilities: choosing sources, reasoning, resolving contradictions, etc.

Consistently, we assume that people store, use and evaluate their domain-specific information. Framing a situation thus involves “loading” context knowledge, both declarative and procedural. In our model a frame is associated to three sets of domain-specific information, called *classes*, which can be seen intuitively as the main source of evaluation of competence in a domain: “being able to obtain information for a task” (CAI), “knowing enough about a domain” (CAB) and “being able to reason in a domain” (CEA).

A *Class of Acceptable Ignorance (CAI)* is a domain-specific set of sources which the participant knows are useful for deciding; e.g. in the TCGG, perception or an expert Gambler’s opinion can be useful sources.

A *Class of Accepted Beliefs (CAB)* contains the “core knowledge” of a domain, i.e. the salient information that is useful to assume in a given domain. CABs can also be seen as a source of information, providing *outside* evidences; as different from *inside* evidences, which are focused on single cases and in the TCGG are provided by perception and Gamblers, outside evidences refer to categories of cases (see also Kahneman (2003)); they are provided by a memory of past similar situations; prior knowledge (e.g. statistical informa-

tion about Gambler’s reliability; base rates such as: the left and right card have the same base probability to be the red one); and reasoning (e.g. if the red card is on the left, it is not on the right). Also the Balance and the reliability values are data in the CAB.

A *Class of Epistemic Actions (CEA)* contains the set of “epistemic actions” useful in a domain. Epistemic actions are operations that modify the participants epistemic apparatus (e.g. adding or subtracting a belief). They are the “productive” parts of the process of decision and they are exploited for reducing Ignorance, Uncertainty and Contradiction before deciding. Epistemic actions are domain specific: for example, in a given situation it is better to ask for more information, to suspend the decision, or to decide quickly. For example, (Gigerenzer & Todd, 1999) describes strategies for dealing with lack of information such as the *recognition heuristic*. There are three main sets of epistemic actions: (1) for deciding if accepting, rejecting or integrating new knowledge, e.g. if the structure is solid and new information is not reliable, reject it; (2) for assessing the epistemic structure during decision, e.g. search for new evidences or background knowledge; reinforce lateral inhibitions; reject a previously accepted information; (3) for revising knowledge, during decision making or after knowing the outcomes, e.g. lower the reliability of a source who is wrong, erase unreliable sources from the CAI, “chunk” a successful case.

Evaluating Background Knowledge. The three classes we have described have an associated reliability value, representing how much the subject evaluates the overall information in the set, which is used in the evaluation of the ambiguity of the domain, too. When knowledge in the classes is actually used in the domain, part of their reliability value is added to the thresholds of Ignorance (for the CAI), Uncertainty and Contradiction (for the CAB and the CEA). The main effect of raising thresholds is that more ambiguity is tolerated, even if the information remains the same; for example, the value of Ignorance in a case of simple decision can be very low not only because many sources are queried, but particularly because all the relevant ones are. Or, Uncertainty can be very low because the situation is consistent with all relevant background information. Expertise in a domain can thus counterweight ambiguity, consistently with Heath and Tversky (1991).

Ambiguity Reduction

Before deciding individuals evaluate ambiguity and try to reduce it; *there is a unique mechanism for epistemic action selection that is ambiguity reduction*; this mechanism also bounds resources such as time and costs, presenting a unique stopping criterium. An individual is ready to decide when Ignorance, Uncertainty and Contradiction are under domain-dependent thresholds. As above described, thresholds can be reduced during decision, too, since when individuals use epistemic strategies knowing that they are appropriate, they rely upon their strategies and not only upon data. Here we recap the five principle phases of the process:

1. *The choice situation is framed* (difficult or easy; number and kind of sources; etc.) and the frame-specific background knowledge (CAI, CAB, CEA) is loaded. We assume that a limited set of frames is available and the

selection depends on the features of the situation (e.g. cards shuffled slowly or quickly) and similarity with other choices which were previously “chunked”.

2. *Reliability values are assigned to the evidences* (e.g. “I am quite sure to have seen the red card on the left”).
3. *Evidences are summed up using the Balance*; the strength of the two conflicting hypotheses (“the right card is red”, and “the left card is red”) is calculated on the basis of their sources and their reliability.
4. *Meta beliefs are calculated* on the basis of the values of the conflicting hypotheses and the background knowledge used. If Ignorance, Uncertainty and Contradiction are under their thresholds the participant is *ready to decide*; otherwise an (additional) epistemic action for reducing them is selected from the CEA and executed. The selected epistemic action is the one whose effects are expected to minimize the current values of ambiguity. An alternative stopping condition exists: if the cognitive load becomes too high, the decision is skipped.
5. *Decision is done*. Depending on the final values of the meta beliefs the decision will be to bet a low or high amount or not to bet at all (if the values are too high). Decision time depends on the length of the process and in particular on the number of epistemic actions actually used.

Some epistemic actions are also specific for learning after knowing the outcomes of the decision. Learning depends on an evaluation, too; for example, the reliability of the source is more likely to be revised after a wrong bet if the error is attributed to the source, and not, for example, to chance.

The Two Cards Gambling Game

In the *Two Cards Gambling Game* experimental paradigm the participant¹ is shown a movie representing two cards, one red and one black. The cards are then turned over (the backs are identical) and shuffled (at different velocities, depending on the experiment). The participant is instructed to look at the movie and bet on the placement of the red card. In some experiments participants were also presented with information about how one or two Gamblers (depending on the experiment) had bet; in these cases it was also shown the competence of the Gamblers, either “novice” or “expert” (the Gamblers are simulated; their bets are biased: experts bet better). In deciding the participant has to bet as quickly and accurately as possible; he has an initial pool of 50 Euros. The participant has 5 choices: “bet 5 Euros on right card”; “bet 5 euros on left card”; “bet 10 euros on right card”; “bet 10 euros on left card”; “do not bet”. After the bet, the outcome is shown: if the participant gave the correct response, he gains the bet, otherwise he loses the same amount of money. The TCGG experimental paradigm permits to study how the difficulty of a decision making task varies depending on the levels of ambiguity, the degree of accord between the sources and their reliability.

¹Fifteen undergraduate students at the University of Rome “La Sapienza” participated in each condition of the experiments. Participants were presented with a set of 40 short movies showing shuffling cards, balanced between “easy” and “difficult”, in random order. Items were presented the center of a computer screen.

Experiment 1: Sources Integration

Experiment 1 allows us to study the role of multiple sources of evidence in the decision process; it is split into three cases: the case 1A investigates the only perceptual source; the case 1B investigates also the contribution of an external source (one Gambler); the case 1C investigates also the contribution of two external sources (two Gamblers).

According to the model, by integrating different sources (as in cases 1B and 1C) participants have less Ignorance, so they should be facilitated in the task (this also involves betting more and taking less time). In experiment 1C, however, we introduced disaccord between the external source; the hypothesis is that when extra sources in contrast with the previous one are introduced, this also introduces an higher amount of Uncertainty, thus resulting in worsening the participants’ performance. Thus the hypothesis is that in case 1B there will be more amount of bet and dell reaction time, while in case 1C there will be less amount of bet and more reaction time, since the cognitive operation of reducing ambiguity (at the metalevel) takes longer.

Method and Results The experimental conditions resulted from a factorial combination of conditions (1A, 1B and 1C) and difficulty of the task (easy or difficult). Separate analysis of variance (ANOVA) with mean correct responses (in percentage), mean amount of bet (in Euros) and mean reaction time (in seconds) as dependent variables were carried out. Experiment (1A or 1B) and difficulty of the task (easy or difficult) were the factors (the former between-subjects, the latter within-subjects).

Correct Responses. The main effect of experiment is significant, $F(2,42)=7,11$; $p<,0018$. Participants in experiment 1B give more correct responses than participants in experiment 1A ($p<,0018$) and 1C ($p<,03328$). There is also significant interaction between difficulty of the task and experiment for percentage of correct responses, $F(2,42)=8,61$; $p<,0005$. A posteriori analysis (Duncan test) shows that in easy tasks participants in Experiment 1B give more correct responses than those in experiment 1A ($p<,0126$), and 1C ($p<,000124$). In difficult tasks participants in Experiment 1B and 1C give more correct responses than those in experiment 1A, ($p<,000527$) and ($p<,000252$) respectively. Results are shown in Table 1.

		DIFFICULTY	EXP.	RESP.
		easy	1A	,783
		easy	1B	,866
		easy	1C	,724
		difficult	1A	,405
		difficult	1B	,527
		difficult	1C	,538
(a) Responses			(b) Interactions	

Table 1: Mean Correct Responses: 1A, 1B and 1C

Amount of Bet. The main effect of experiment is significant, $F(2,42)=14,31$; $p<,00001$. Participants in experiment 1B bet more than participants in experiment 1A ($p<,001969$)

and 1C ($p < .0000078$). There is no significant interaction between difficulty of the task and experiment for amount of bet, $F(2,57) = .79$; $p < .4594$. Results are shown in Table 2.

		DIFF.	EXP.	BET
EXP.	BET	easy	1A	8,232
1A	7,309	easy	1B	9,437
1B	8,625	easy	1C	8,038
1C	6,712	difficult	1A	6,387
		difficult	1B	7,812
		difficult	1C	5,385

(a) Amount of the Bet (b) Interactions

Table 2: Mean Amount of Bet: 1A, 1B and 1C

Reaction Time. The main effect of experiment is significant, $F(2,42) = 116.41$; $p < .00001$. Participants in experiment 1B take less time in answering than participants in experiment 1A ($p < .000112$) and 1C ($p < .000111$); participants in experiment 1A take less time in answering than participants in experiment 1C ($p < .00006$). There is also significant interaction between difficulty of the task and experiment for reaction time, $F(2,42) = 3.90$; $p < .0258$. A posteriori analysis (Duncan test) shows that in easy tasks participants in Experiment 1B take less time than those in experiment 1A and 1C, ($p < .000061$) and ($p < .000032$) respectively. Moreover, participants in Experiment 1A take less time than those in experiment 1C ($p < .000116$). In difficult tasks participants in Experiment 1B take less time than those in experiment 1A and 1C, ($p < .008823$) and ($p < .000032$) respectively. Participants in Experiment 1A take less time than those in experiment 1C ($p < .000060$), too. Results are shown in Table 3.

		DIFFICULTY	EXP.	TIME
EXP.	TIME	easy	1A	5,667
1A	5,707	easy	1B	4,390
1B	4,743	easy	1C	7,742
1C	7,742	difficult	1A	5,747
		difficult	1B	5,097
		difficult	1C	7,742

(a) Reaction Time (b) Interactions

Table 3: Mean Reaction Time: 1A, 1B and 1C

Discussion Our results show participants are influenced in a significant way by the difficulty of the task (the influence of the expertise of the Gamblers is not shown here). An external source of evidence adds confidence, in particular in difficult tasks: participants bet more and better. Interestingly, participants in experiment 1B take less time in deciding than participants in case 1A, even if they manipulate more information; this result is only surprising if we consider decision as a “flat” process of calculating the evidences; on the contrary, our model describes one more phase where meta-beliefs are lowered before deciding. This means that when the context of choice is ambiguous, *processing time does not depend on*

the number of sources but on the level of ambiguity. The case without external Gamblers (Experiment 1A) gives the advantage (in terms of time) of considering less information at the base level; but his advantage is overcome in the case with an external Gambler (Experiment 1B) by the minor quantity of Ignorance to reduce at the metalevel. By introducing two external sources of evidence (Experiment 1C) the participants’ performance worsens (our findings, not shown here, indicate that this is mainly due to the case of disaccord between the external sources).

Experiment 2: Searching for Information

When does somebody feel ready to decide? when does he search for more information? In the preceding experiment this aspect indirectly emerges from reaction times; in this experiment we allowed participants to ask for more information, i.e. to ask to see more Gamblers. The first hypothesis is that, while in easy tasks participants will be satisfied even only by their perceptual source, in difficult tasks they will ask for a significant amount of external sources in order to reduce Ignorance. The second hypothesis is that the requested information is very salient, thus it will be taken into account more with respect to non requested information (as in Exp. 1B).

Method and Results Differently from the previous experiment, after having seen the cards shuffling, and before betting, participants had also the opportunity to see the bet of one or two Gamblers by paying 0,1 Euro each. The experimental condition resulted from a factorial combination of difficulty of the task (easy vs. difficult) and competence of Gamblers (expert vs. novice). One set of 40 movies (balanced between easy and difficult) was used. Analysis of variance (ANOVA) has been conducted with average number of requested Gambler as dependent variable. Competence of the Gambler (Novice or Expert) was the within-subjects factor. Analysis of variance (ANOVA) has been conducted with mean percentage of Accord with the Gambler (i.e. how many times the bets of the participant and the Gambler are the same) as dependent variable. Experiment (1B or 2) and Competence of the Gambler (Novice or Expert) were the factors (the former between-subjects, the latter within-subjects).

Requested Gamblers. The main effect of difficulty of the task is significant, $F(1,14) = 47.98$; $p < .00001$. Participants request more Gamblers in difficult than in easy tasks. Results are shown in Table 4.

MOVIE	REQUEST
Easy	,404
Difficult	,700

Table 4: Mean Number of Requests: 2

Percentage of Accord The main effect of experiment is significant, $F(1,28) = 7.18$; $p < .0122$. Participants in Experiment 2 are more in accord with Gamblers than those in Experiment 1B. There is no significant interaction between Competence and Experiment, $F(1,28) = .39$; $p < .5383$. Results are shown in Table 5.

		GAMBLER	EXP.	ACC.
EXP.	ACCORD	Expert	2	,712
2	,652	Expert	1B	,679
1B	,61	Novice	2	,592
		Novice	1B	,541
(a) Accord		(b) Interactions		

Table 5: Mean Percentage of Accord: 1B and 2

Discussion Difficult tasks require more (convergent) information. An hypothesis is that in easy tasks the player can be “ready to decide” even with the only perceptual source; on the contrary in a difficult task (and with high uncertainty) he will ask for more information before deciding. This position is also consistent with the use of non compensatory heuristics such as “take the best” in bounded rationality (Gigerenzer & Todd, 1999) and with satisfaction criteria (Simon, 1957). More precisely, a typical epistemic action for minimizing Ignorance, Uncertainty and Contradiction is “ask for more info”. In the case of easy tasks, the single (very reliable) perceptual source makes perceived ignorance very low (the informativeness of any new source will be low, because it will change the beliefs not so much). So, there is no need of “reducing ignorance” -on the contrary, any new source can potentially raise uncertainty. In the case of a difficult task, both perceived ignorance and uncertainty are high with the only perceptual source, and he will probably ask for more information. But, when ignorance lowers as an effect of more info, there are possibilities of raising uncertainty and contradiction -so the situation can be even worse (Gigerenzer and Todd (1999) call this the “less is more effect”). The MSEM model indicates as the main goal of epistemic actions to reduce ambiguity and not to maximize information; this unique mechanism permits to treat the Ignorance-Uncertainty trade-off. Participants in Experiment 2 are more in accord with Gamblers with respect to those in Experiment 1B. This result suggests that explicitly *requested data are more useful, since they resolve a need for information* (reducing Perceived Ignorance). According to the MSEM model, epistemic structures are organized in a network of relationships; epistemic actions are conducted in order to solidify the structures before deciding, so they search for information that is required to fill in the epistemic needs.

Conclusions

According to the MSEM model, before deciding, participants try to minimize ambiguity adopting domain-specific strategies, the *epistemic actions*. This operation leads to more or less fast and more or less confident decisions. With respect to the model of Usher and McClelland (2001), our model involves two levels, the base and meta level: while the Balance describes a standard, compensatory way to integrate new information, either in accord or disaccord with previous information, individuals can employ various cognitive strategies for this integration. With respect to the literature about heuristics (Kahneman, 2003; Gigerenzer & Todd, 1999) our model individuates an unique satisficing criterion for heuristics selection which is ambiguity-reduction; moreover, it focuses on

epistemic actions which are generally more fine grained than heuristics. Typically heuristics are claimed to model all the decision process; on the contrary, we claim that the decision making process requires many phases and in particular a cognitive evaluation precedes the choice of the epistemic action.

Our findings indicate that the epistemic process of representing an ambiguous context of choice involves not only a “flat” integration of information in the domain, but also reasoning about the domain, representing its ambiguity and activating epistemic strategies for reducing it. In a series of experiments (only two are presented here) using the TCGG paradigm we have also found many epistemic actions in play; for example, comparing Exp. 1B and 2 it emerges that different strategies are in play for accepting or rejecting information, depending on information needs; this result is also replicated in an experiment about belief revision. Our findings also indicate that ambiguity-reduction produce many biases, heuristics and cognitive illusions described in the literature. All the results have also been simulated in Pezzulo (2006) with a system including all the five phases above described.

References

- Busemeyer, J. R., & Johnson, J. G. (2004). Computational models of decision making. In D. J. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment & decision making*. Malden, MA, USA: Blackwell.
- Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, *100*, 432–459.
- Camerer, C., & Weber, M. (1992). Recent developments in modeling preferences: Uncertainty and ambiguity. *J. Risk Uncertainty*, *5*.
- Castelfranchi, C. (1996). Reasons: belief support and goal dynamics. *Mathware & Soft Computing*, *3*.
- Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford University Press.
- Gigerenzer, G., & Todd, P. M. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.
- Heath, C., & Tversky, A. (1991). Preference and belief: Ambiguity and competence in choice under uncertainty. *J. Risk Uncertainty*, *4*.
- Kahneman, D. (2003). A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist*, *58*, 697-720.
- Pezzulo, G. (2006). *Into the gambler’s frame of mind: Decision making under uncertainty in the two cards gambling game*. Unpublished doctoral dissertation, University of Rome “La Sapienza”.
- Pezzulo, G., Lorini, E., & Calvi, G. (2004). How do i know how much i dont know? In *Proceedings of cogsci 2004*. Chicago.
- Simon, H. A. (1957). *Models of man - social and rational*. New York: John Wiley and Sons.
- Usher, M., & McClelland, J. L. (2001). On the time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, *108*(3).