

Commitment to Explicitly Reported Choices: Evidence in Decision Making under Uncertainty

Giovanni Pezzulo (giovanni.pezzulo@istc.cnr.it)
Institute of Cognitive Sciences and Technologies - CNR
Via S. Martino della Battaglia, 44 - 00185 Roma, Italy

Abstract

We present some experimental evidence about the different roles and values of information in belief building and decision making under uncertainty. We adopt the *Two Cards Gambling Game* experimental paradigm and we present to the participants choice situations in which multiple sources of information, having different reliability, are available, in different experimental conditions: one or more sources, presented either together or at different times. Our results show that decisions depend on the availability and quality of their sources are are biased by commitment to explicitly reported choices as well as by the (subjectively felt) need for information.

Keywords: decision making, uncertainty, belief revision, commitment

Introduction

In several related fields a distinction between two kinds of knowledge exists, the former being more related to “inert” knowledge and the latter being more related to knowledge either used for action (such as decision making), or under attentive control, etc.

In philosophy and AI (Dennett, 1987; Dummett, 1991) there is a distinction between *implicit* and *explicit* knowledge; even if these terms can assume different meanings, one of the main claims is that we have direct access to explicit knowledge and we can express it in a linguistic format (Dummett, 1991). Humans have a great amount of potentially available knowledge (which can be inferred when required (Dennett, 1987)) as well as some *tacit* knowledge which is exploited but not represented at all (Chomsky, 1965). Only some knowledge is explicitly available, typically in a declarative format; heuristics such as *availability* (Kahneman, 2003) testify that knowledge can be more or less ready to be used.

In decision making (Kahneman, 2003) there is a related distinction between *preferences* and *choices*. Even if they can influence behavior, preferences can be revised or negotiated before a choice. On the contrary choices are explicitly reported, are considered definitive, imply a certain commitment and have a resistance to be modified; some phenomena such as *cognitive dissonance* (Festinger, 1957) testify that we are able to revise post-hoc our preferences in order to explain our choices.

A similar distinction between *information*, *beliefs* and *acceptances* appears in the field of belief revision (Castelfranchi, 1997; Paglieri, 2006). Information basically means data having a source and supported by reasons, stimuli, evidence, etc. A belief instead is an “internalized” information, i.e. information which has undergone a cognitive process of evaluation (e.g. against previously known information) and selection (e.g. on the basis of a lack of information). These operations can be either deliberated or automatic (Castelfranchi, 1997). Typically a cognitive agent firstly builds its belief structure depending on its information and their

sources, and then decides on the basis of its beliefs. However, it is only by using a certain belief for action (e.g. for a decision) that it assumes the role of an *acceptance*. It is not strictly necessary that something is believed to be used as a basis for deciding (i.e. as an acceptance), even if this typically the case of rational agents. It is not the case that all the information or the beliefs are used for deciding, too: this is why cognitive agents typically exploit sophisticated cognitive operations (such as selecting the best sources, negotiating contradictions, etc.) before deciding, having the goal to make the belief structure supporting decision very strong. However, after deliberation a cognitive agent either acts or intends to act according to its decision; in this case it is typically *committed* to its decision. As (Cohen & Levesque, 1990) puts it, *intention is choice with commitment*. Consequently, we claim that this is a side-effect of the operations for strengthening the belief structure: choices are *accepted* and involve a certain degree of commitment since they have been integrated in the belief system and it is costly to revise them.

The aim of this paper is twofold. (1) We present some experimental evidence about commitment to explicitly reported choices which depends on the reliability of the sources of information. In the discussion we propose that these data can be due to different strategies for accepting or rejecting information in the different cases. (2) In order to better frame these findings, we then introduce the *Multiple Source Evaluation Model (MSEM)*, which predicts that individuals adopt different cognitive strategies for dealing with uncertain and ambiguous situations depending on contextual parameters such as the kinds and levels of ambiguity, uncertainty, ignorance, etc. In particular, the MSEM indicates that before deciding individuals adopt *epistemic actions* (Kirsh & Maglio, 1994) (i.e. actions aiming to modify its belief structure, either by accepting new information or by modifying/rejecting beliefs). Epistemic actions have the effect to strengthen the initial beliefs supporting decision making, and are thus candidates for explaining the commitment effect we have found.

The Experiments

The inspiration for the experiments presented here comes from our past work. In Pezzulo (2006); Pezzulo and Couyoumdjian (2006) we investigated how humans deal with multiple sources of information, either in accord or in disaccord, in the field of decision making under uncertainty. The experiments were conducted by using the *Two Cards Gambling Game* experimental paradigm, which we use also for the experiments we present here. In the TCGG the participant 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 manipulate the number of sources and the levels of ambiguity, and 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.

Review of former experiments The first experiment we have conducted with the TCGG paradigm aims at investigating role of multiple sources of evidence in the decision process. The experiment is split into three cases: case 1A investigates the only perceptual source; case 1B investigates also the contribution of an external source (one Gambler); case 1C investigates also the contribution of two external sources (two Gamblers). Participants thus deal with ambiguous situations: for example, the opinion of the participant can contrast with the opinion of the Gamblers; or two Gamblers can bet differently.

Our findings indicate that individuals base their decision on available information; however, different strategies for integrating them can be adopted. In particular participants integrating two different sources of evidence (case 1B) have less Ignorance and are facilitated in the task, thus they bet significantly more and take significantly less time. We have interpreted these findings according to the MSEM (introduced later on), indicating that specific cognitive operations, the *epistemic actions*, are adopted for reducing ambiguity before deciding, especially in the case of disaccord between two sources (this explains certain patterns in response time).

In case 1C, by introducing two external sources of evidence (two Gamblers), the participants’ performance worsens. Our findings indicate that this is mainly due to the case of disaccord between the external sources; again, we have interpreted these findings according to the MSEM, which indicates that the cause is an higher level of Uncertainty. Thus it is not the case that ‘more information is better’, but on the contrary an optimal trade-off between amount of information and coherence is required in order to decide accurately and fast. Gigerenzer and Todd (1999) presents similar evidence of ‘less is more effect’: choosing only the most relevant cues is better for deciding in environments having a non compensatory information structure, and if extra information is added on it can worsen the performance. An alternative explanation is the use of the *anchoring* heuristic (Kahneman, 2003): participants remain committed to their perceptual source, because the external sources are in disaccord and thus are unreliable.

Explicit vs. Implicit knowledge It has to be noted that the kind of contradiction found in case 1C has a peculiarity: it is *explicit*, i.e. it involves already established/accepted beliefs of other agents (their *choices* and not simply *preferences*).

Consider in fact that during the decision process contradiction very often arise, due to contrasting evidence (e.g. perceptual vs. external source). During the process, however, this is likely a contrast of preferences, in the sense that their values are assumed to be tentative, defeasible and not definitive; the participant is not yet committed to any choice. The point is that this kind of contradiction can be negotiated and reduced by many kinds of epistemic actions (e.g. by strengthening or lowering some factors). Instead, once contradictory information is explicitly reported as in case 1C, it can not be reduced except by a selection of the sources, and by selectively accept or reject incoming information, i.e. choose one or another Gambler; and this is a costly epistemic operation.

How information is integrated before or after a choice In order to further investigate the ‘explicit vs. implicit’ issue, here we present two experiments using the TCGG experimental paradigm. The former is condition B of the already introduced experiment, with one external source of information (one Gambler). In the latter the same materials are used, but subjects report their choice before looking at the bet of the external source; they can decide whether or not to revise it after looking at the external source’s bet. Two sets of fifteen 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. We are interested in comparing the strategies for integrating information in the two different cases: information is the same in both studies, but while in the former it is available before any choice, in the latter part of it is only available after a choice has been explicitly reported.

Experiment 1: Perceptual Source plus One Gambler

In experiment 1 participants, before betting, see the bet of another Gambler. It serves to investigate the combination of two different source of information, the former perceptual and the latter external (one Gambler). Sixty movies varying in difficulty (thirty easy and thirty difficult) were presented, together with information about how a Gambler (either expert or novice) has bet. The main hypothesis is that in difficult tasks, in which uncertainty (in the perceptual source) is higher, the role of the external source is more relevant. However, our main aim here is investigating the mean ‘percentage of accord’ between the participants bets and the Gamblers bets, i.e. the cases of the same bet between the two¹; this is done in order to compare this experiment with the next one.

Method and Results The experimental conditions resulted from a factorial combination of difficulty of the task (easy vs. difficult) and competence of the Gambler (novice vs. expert). An analysis of variance (ANOVA) has been conducted with the mean ‘percentage of accord’ (between the participants bets and the Gamblers bets) as dependent variable. Difficulty of the task and competence of the Gambler were the within-subjects factors. The main hypothesis is that there is more accord with the most reliable external source, experts.

The main effect of competence is significant,

¹ It don’t indicate if the perceptual and the external source are concordant, but if the subject and the Gambler bet on the same card.

$F(1,14)=47,29$; $p<,00001$: participants are more in accord with expert Gamblers than with novices. Difficulty of the task is not significant, $F(1,14)=,46$; $p<,5080$; interactions are not significant, too, $F(1,14)=,54$; $p<,4730$.

Results are shown in Table 1.

MOVIE	ACC.	GAMBLER	ACC.
difficult	,606	expert	,679
easy	,615	novice	,541
(a) Difficulty		(b) Competence	

Table 1: Mean Percentage of Accord in Experiment 1

Percentage of Accord and Reaction Time In order to better understand the cases of accord and disaccord, we also analyzed the reaction times in accord and disaccord: an analysis of variance (ANOVA) has been conducted with reaction time as variable. Competence of the Gambler (novice or expert) and Accord between participant and Gambler (accord or disaccord) were the within-subjects factors. We want to investigate if the accord influences the cognitive load of the participant before decision. The main hypothesis is that in the case of accord reaction time is lower.

The main effect of accord is significant, $F(1,14)=17,43$; $p<,0009$; participants bet faster when they are in accord with the Gambler. There is also significant interaction between competence and accord, $F(1,14)=6,50$; $p<,0232$. From a posteriori analysis (Duncan Test) it emerges that participants bet faster when they are in accord (vs. disaccord) both with novices ($p<,0076$) and experts ($p<,00007$).

Results are shown in Table 2.

ACCORD	TIME	GAMB.	ACC.	TIME
accord	4,190	expert	accord	4,05
disaccord	5,118	expert	disaccord	5,317
(a) Accord		novice	accord	4,330
		novice	disaccord	4,918
		(b) Interactions		

Table 2: Mean Percentage of Accord and Reaction Time in Experiment 1

Discussion Participants are in accord with expert Gamblers more than with novice ones in both easy and difficult tasks. Moreover, reaction times are significantly lower in the case of accord with an expert, indicating that there is no ‘epistemic conflict’ to resolve and thus the level of uncertainty is low. The case we are more interested in is that of a conflict between the perceptual and the external source; in this case, in fact, the external evidence can be either accepted (revising a belief), or rejected (we can call those operations *epistemic actions*). Different strategies can be used: for example, a dubious perceptual evidence could be challenged by evidence provided by a reliable expert Gambler, but not by a novice one. We have also analyzed our results by introducing another factor, correct or incorrect responses. It is interesting

to note that the higher reaction time (7,23 seconds) is in the cases of ‘accord with an expert, incorrect response’ (significant with all the other times). It seems to be the case of a high ‘computational cost’ of revising a previously correct and strong belief, formed on the basis of perceptual evidence, relying upon an expert. This example indicates that belief revision do not rely on a generic rule for information negotiation, otherwise there would have been no difference in time. On the contrary, it is possible that depending on the strength of the initial belief, the reliability of the sources and the level of uncertainty specific strategies are adopted for deciding which information or belief to accept as a basis for deciding, and which to discard. The next experiment, about belief revision, further investigates this point.

Experiment 2: Effects of Revision

Experiment 2 is about belief revision; the materials are the same of experiment 1 but now there are two phases, called *Before* and *After*. In the first phase (Before) the participant bets only seeing the cards shuffling. In the second phase (After) after the bet he sees the bet of one Gambler and can choose either to confirm the bet or to change it. Since participants can decide to change their bet after assuming new information, this is a case of belief revision; our main hypothesis is that expert Gamblers will lead to more belief revision. Since the decision process has as a precondition a solid belief structure, we assume that it is more difficult for successive information to influence it. Our main hypothesis is that the degree of accord in this case is lower than what happens in experiment 1, in which both sources are presented together. In fact, by explicitly expressing a preference the player both explicitly reports a belief and commits to it; so, there is an additional cost in changing opinion.

Method The experimental conditions resulted from a factorial combination of difficulty of the task (easy vs. difficult) and expertise of the Gambler (novice vs. expert). Analysis of variance (ANOVA) has been conducted with mean percentage of accord with the Gambler and mean amount of bet as dependent variables. Phases of the experiment (Before or After) and competence of the Gambler (Novice or Expert) were the within-subjects factors.

Percentage of Accord The main effect of the phase is significant, $F(1,14)=13,78$; $p<,0023$. In the After phase participants are more accord with the Gambler. An important note: since in the Before phase participants do not see any Gambler, we only calculated the ‘hypothetical accord’ in order to match it with the After phase, in order to see if their choice is influenced by the Gambler (i.e. if their accord grows).

There is significant interaction between the competence of the Gambler and percentage of accord, $F(1,14)=5,14$; $p<,0397$. From post hoc analysis (Duncan Test) it emerges that with expert Gamblers participants are more in accord in the After than in the Before phase ($p<,000732$); this does not happen with novice Gamblers ($p<,29456$).

Results are shown in Table 3.

Amount of Bet The main effect of phase is significant, $F(1,14)=5,75$; $p<,0310$. In the After phase participants bet more. There is significant interaction between competence of the Gambler and amount of bet ($F(1,14)=18,45$; $p<,0007$).

		GAMBLER	PHASE	ACC.
PHASE	ACC.	Expert	Before	,517
Before	,528	Expert	After	,567
After	,561	Novice	Before	,540
		Novice	After	,555

(a) Accord (b) Interactions

Table 3: Mean Percentage of Accord in Experiment 2

From post hoc analysis (Duncan Test) it emerges that with an Expert participants bet more in the After than in the Before phase ($p < .000302$); this does not happen with a Novice. Moreover, while in the first phase (before) there is not significant difference, in the After phase participants bet more with an Expert than with a Novice ($p < .000108$).

Results are shown in Table 4.

		GAMBLER	PHASE	BET
PHASE	BET	Expert	Before	7,261
Before	7,260	Expert	After	7,906
After	7,526	Novice	Before	7,260
		Novice	After	7,147

(a) Bet (b) Interactions

Table 4: Mean Amount of Bet in Experiment 2

Discussion We can see that when there is an Expert the participants modify significantly their bet; also the percentage of accord before and after changes significantly; moreover, they modify in a significant way the amount of bet. On the contrary in presence of a Novice data are not significant. This means that only Experts are able to produce an appropriate change in the strength of beliefs, while the (late) influence of Novices is too weak. These findings indicate that participants commit to their explicitly reported choices, which are thus more difficult to revise. Later on we compare this experiment with the previous one and, as we will see, a more marked commitment effect emerges.

Comparison between Experiments 1B and 2.

Analysis of variance (ANOVA) has been conducted with mean percentage of accord with the Gambler as dependent variable. Experiment (1B and 2-After) and competence of the Gamblers (Novice or Expert) were the factors (the former between-subjects, the latter within-subjects).

The main effect of experiment is significant, $F(1,28)=4,34$; $p < .0465$. Participants in Experiment 2 (phase After) are less in accord with Gamblers than participants of Experiment 1B.

Significant interaction was found, too, $F(1,28)=5,44$; $p < .0270$. From post hoc analysis (Duncan Test) it emerges that participants in Experiment 1B are more in accord with expert Gamblers than those of Experiment 2 ($p < .007003$); this does not happen with novice Gamblers ($p < .706332$)².

Results are shown in Table 5.

²A control experiment with successive phases but without ex-

		GAMB.	EXP.	ACCORD
EXP.	ACCORD	Expert	2	,567
2	,561	Expert	1B	,679
1B	,61	Novice	2	,555
		Novice	1B	,540

(a) Accord (b) Interactions

Table 5: Mean Percentage of Accord, comparing Experiments 1B and 2

Discussion We can see that there is less accord with the Gambler in the After phase of Experiment 2 than in Experiment 1. This means that there is a tendency of participants in Experiment 2 to maintain their choices which are explicitly reported. Even if they assume new information, they are less prone to revise: we call it a *commitment effect*. Our data can be interpreted on the basis of the *anchoring heuristic* (Kahneman, 2003): we are better at relative than absolute thinking, so we try to adjust judgments starting from a familiar position or assumption. In this case, participants have a more stable opinion in experiment 2 since they already start from a reported (and stable) choice.

However, this is only part of the answer. Our findings indicate that beliefs can be maintained or revised depending on their strength (i.e. how much evidence they account for), and that under different conditions new information which contrasts with previous knowledge can be integrated or rejected. We therefore introduce the Multiple Source Evaluation Model (MSEM), which describes the belief formation and revision process during decision making, and use it for interpreting the results of our experiments. For a complete description of the MSEM we refer to (Pezzulo, 2006; Pezzulo & Couyoumdjian, 2006). Here we only describe its main components and focus on the role of epistemic actions in different contexts.

The Multiple Source Evaluation Model

The *Multiple Source Evaluation Model (MSEM)* focuses 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: 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 also claim that many coarse-grained heuristics such as representativeness and availability (Kahneman, 2003) emerge from this more fine-grained strategy, depending on contextual factors.

The fundamental claim of the MSEM is that the result of a decision depends on the supporting epistemic structure, composed of: (1) beliefs *in* the domain of decision (base beliefs

explicitly reporting the choice (no significant differences) has been conducted, too, showing no significant results.




Figure 1: The Balance with Two Plates

about information sources and their reliability); (2) beliefs *about* the domain of decision (the meta-beliefs: evaluation of ambiguity in the current decision); (3) declarative and procedural expertise in the domain of decision.

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, 1997). We consider many kinds of beliefs and sources, including *inside* evidence, focused on the contingent situation (e.g. perceptual data) and *outside* evidence, focused on categorical knowledge (e.g. previous similar situations, statistical information, etc.); this distinction is also presented in (Kahneman, 2003). A good analogy is a *Balance with two Plates* where evidence 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 evidence 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. The Balance shares some resemblances with the Leaky, Competing Accumulator Model (LCAM, Usher and 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 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 several studies ‘lack of information’ as well as 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). In particular, in the MSEM three meta beliefs, *Ignorance, Uncertainty and Contradiction*, are claimed to have a crucial role in decision making. According to the MSEM, before deciding the subject tries to minimize (under domain-specific thresholds) the values of these three meta-beliefs. For doing so, he performs one or more *epistemic actions*, such as ‘ask for more information’, or ‘reject information from an unreliable source’, until the levels of ambiguity are considered acceptable.

Declarative and Procedural Expertise The model also describes *declarative and procedural expertise* for dealing with

uncertain situations. Here we focus mainly on procedural expertise, consisting of applying a set of rules, the *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’.

The model introduces three kinds of *epistemic actions* for accepting, integrating and revising beliefs in presence of specific kinds of ambiguity: (1) epistemic actions for integrating new knowledge: they decide whether or not to assume knowledge, etc. (2) epistemic actions for assessing the epistemic structure during decision: searching new evidence, reinforcing lateral inhibitions, etc. (3) epistemic actions for revising knowledge: raising or lowering the reliability of sources or epistemic actions, etc.

In order to interpret the results of the experiments presented here, we focus only on the former kind of epistemic actions. In fact, while the Balance describes a standard, compensatory (Payne, Bettman, & Johnson, 1988) way to integrate new information, either in accord or disaccord with previous information, individuals can employ various cognitive strategies for this integration. New knowledge is not in fact merely added to the Balance, but must be *accepted*, i.e. integrated into the preexisting epistemic structure. In a similar way Boutilier, Friedman, and Halpern (1998) argues that information (e.g. incoming data, sensor stimuli, etc.) is only *accepted* under certain conditions such as being consistent with existing knowledge or, in the case of disagreement, being more likely than previous knowledge.

Integration of evidence is highly contextual and depends on many factors: possible conflicts, the amount of information already available, etc. There are many possible situations in which new information cannot simply be accepted by default, such as conflicts with strong previous knowledge; in these cases, specific epistemic actions are utilized for filtering new knowledge. For example, if new evidence conflicts with one priorly assumed, depending on the context (e.g. new sources are reliable or unreliable; agents ignorance is high or low) it can be accepted or rejected.

As an example, some of the epistemic actions used in the MSEM are presented here. The format is close to fuzzy logic³, which is also the language which has been used for implementing the model (see Pezzulo (2006) for details and for comparison with experimental data). All these epistemic action are of type (1): they serve to decide if accepting or rejecting new information.

```
(reliability: ,92)
IF new_evidence_conflicts_with_assumed_knowledge
AND source is reliable
THEN accept_new_evidence
EXPECT uncertainty is low
```

³Epistemic actions are fuzzy rules having 1) a context for applying the rule (the IF part); 2) an operation (the THEN part) which describes the procedure to be applied ; 3) an expectation (the EXPECT part) indicating the expected value of one or more meta-beliefs if the rule is applied in the given context ; 4) a degree of reliability indicating how much the EXPECT part is reliably produced when the rule is applied, and serving for rule selection.

```
(reliability: ,90)
IF new_evidence_conflicts_with_assumed_knowledge
AND source is unreliable
OR uncertainty is very_low
THEN reject_new_evidence
EXPECT uncertainty is low
```

As a consequence of the process of strengthening the belief structure before a choice, it is much more difficult to revise a belief structure after than a choice. This means that when a participant is asked to report explicitly a choice, in order to do so he firstly builds a ‘solid enough’ belief structure. This structure is now much more resistant to revisions. Our claim has some resemblances with the literature about *cognitive dissonance* (Festinger, 1957), that is a post-hoc revision in order to explain our choices. The difference in our interpretation is that most of the work or reinforcing the reasons to believe (i.e. the belief structure) is done *before* choosing and *in order to* choose. According to the MSEM, in facts, individuals try not to decide until uncertainty is considered acceptable.

Conclusions

We have found evidence that the levels of reliability of the sources influences the levels of uncertainty in decision making (Exp. 1); and that a *commitment effect* occurs after explicitly reporting beliefs (Exp. 2). In order to understand these phenomena we have proposed to look at which epistemic actions are adopted for accepting or rejecting information in different contexts, and we have introduced the MSEM.

The results of Exp. 1 can be explained in the MSEM as the ‘computational cost’ in reducing uncertainty, by means of successive epistemic actions, in the different conditions. The commitment effect of Exp. 2 can be explained in the MSEM by considering that before deciding the belief structure serving as support has to be strengthened in order to obtain tolerable levels of uncertainty, and this happens by applying specific epistemic actions. After an explicitly reported choice it is much more difficult to revise the belief structure. In particular, epistemic actions for accepting or rejecting new information have to be used; information has to be selected and it is no more possible to simply negotiate or adjust information or to simply use compensatory heuristics. It is worth noting that, differently from Festinger (1957), we do not interpret this phenomenon as a post-hoc revision, but as a side-effect of uncertainty-reduction occurring *before* choice.

There is a related aspect which influences the selection of information: the ‘need for information’. In (Pezzulo & Couyoumdjian, 2006; Pezzulo, 2006) we presented an experiment using the TCGG paradigm in which participants are not shown the bet of other gamblers, but can explicitly ask to see it (by paying 0,1 Euro each). Our results show that in that case participants are significantly more in accord with gamblers than in the experiments here presented: *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 people only search for information that is required to fill in the gaps

in the structure.

We think that our results can support the distinctions between different kinds of knowledge which exist in many literatures, and in particular between many *roles* that the same knowledge assumes before, during and after a choice, e.g. as evidence, as belief (integrated in a belief structure) or as acceptance (used for deciding). We have also shown that the order of presentation of information, and the cognitive strategies adopted for integrating or rejecting them, can influence decision making -and this aspect is often underestimated.

References

- Boutilier, C., Friedman, N., & Halpern, J. Y. (1998). Belief revision with unreliable observations. In *Proceedings of AAAI/IAAI 1998* (p. 127-134).
- Camerer, C., & Weber, M. (1992). Recent developments in modeling preferences: Uncertainty and ambiguity. *J. Risk Uncertainty*, 5, 325–370.
- Castelfranchi, C. (1997). Representation and integration of multiple knowledge sources: issues and questions. In S. e. T. Cantoni Di Gesu’ (Ed.), *Human and machine perception: Information fusion*. Plenum Press.
- Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.
- Cohen, P. R., & Levesque, H. J. (1990). Intention is choice with commitment. *Artificial Intelligence*, 42, 213–261.
- Dennett, D. (1987). *The intentional stance*. MIT Press.
- Dummett, M. (1991). *The logical basis of metaphysics*. Cambridge, MA: Harvard University Press.
- 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.
- Kahneman, D. (2003). A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist*, 58, 697-720.
- Kirsh, D., & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*, 18, 513-549.
- Paglieri, F. (2006). *Belief dynamics: From formal models to cognitive architectures, and back again*. Unpublished doctoral dissertation, University of Siena.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 14, 534–552.
- 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., & Couyoumdjian, A. (2006). Ambiguity-reduction: a satisficing criterion for decision making. In *Proceedings of Cogsci 2006* (pp. 669–674).
- Usher, M., & McClelland, J. L. (2001). On the time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, 108(3), 550–592.