

How do I Know how much I don't Know?

A cognitive approach about Uncertainty and Ignorance

Giovanni Pezzulo (pezzulo@ip.rm.cnr.it)

Istituto di Scienze e Tecnologie della Cognizione - CNR, viale Marx 15 – 00137 Roma, Italy
and Università degli Studi di Roma “La Sapienza” – piazzale Aldo Moro, 9 – 00185 Roma, Italy

Emiliano Lorini (e.lorini@istc.cnr.it)

Università degli Studi di Siena – via Banchi di Sotto, 5 - 3100 Siena, Italy

Gianguglielmo Calvi (calvi@noze.it)

Istituto di Scienze e Tecnologie della Cognizione - CNR, viale Marx 15 – 00137 Roma, Italy

Abstract

We propose a general framework for reasoning and deciding in uncertain scenarios, with possibly infinite source of information (open world). This involves representing ignorance, uncertainty and contradiction; we present and analyze those concepts, integrating them in the notion of *lack of confidence* or *perplexity*. We introduce and quantify the *strength of the beliefs* of an agent and investigate how he can do explicit *epistemic actions* in order to supply information lacks. Next we introduce a simple distributed game (RBG) and we use it as a testbed for comparing the performance of agents using the (classical) “expected utility maximization” and the “perplexity minimization” strategies.

Introduction

In an “open world” uncertainty and ignorance are difficult categories to deal with; how much can I be certain of a belief of mine? how much information there is that I have not considered and I should?

The first aim of the present work is to provide an analysis of epistemic dimensions such as *strength of belief*, *uncertainty*, *contradiction* and *ignorance* (or ambiguity). A special focus will be given to the third dimension. In Economical literature the notion of ignorance has been extensively investigated (Shackle, 1972) and ways to quantify it have been proposed (Shafer, 1976). In those approaches “lack of information” has been shown to affect the decision process and ambiguity aversion in subject has been identified; see Camerer & Weber (1992) for a review of the literature on decisions under ambiguity. We will argue in the following analysis that Ignorance is a subjective evaluation of actual lack of information on the basis of cognitive *evidential models*. The agent has a model (script) of his sources that allows him to evaluate that a certain type and a certain number of sources can provide *sufficient information* for reducing ignorance close to zero. In this way the strength of the belief and the (perceived) ignorance are two different measures, the second belonging to the meta-level. The second aim of the present work is to investigate the decision dynamics in an open world, with conflicting beliefs and multiple sources of information. We will formalize the process that leads the agent to acquire new

information from the world (from witnesses) and that leads the agents to be “ready” to decide. We will claim that this process involves strength of beliefs that are relevant for deciding, as well as uncertainty and ignorance. The results of the current work are suitable e.g. for MAS environments, where an agent has to take decisions in open worlds.

The Red-or-Blue Card Game (RBG)

We introduce a simple distributed game that is suitable for Multi-agent system simulations as well as for human experiments: the agents (players) have to bid on the color of a card (red or blue) and they have many sources of information (their perception and potentially infinite witnesses); the game can last an indefinite number of turns. The bidding game is the following: a card is shown (very quickly) to the player; it can be either red or blue and the player has to bid on the right color (he starts with 1000 Credits). We assume that he cannot be totally sure of his own perception (e.g. it is shown very quickly, or the lights in the room are low), but he is able to provide a degree of certainty about the color. Before bidding he can ask for help to a (potentially) infinite number of witnesses that have observed the scene and provide the answer “red” or “blue” (without degrees of certainty); those new information can lead the agent to confirm or revise his beliefs. Asking a witness has not a cost in Credits but it costs 1 *Time*. Credits and Times can be aggregated in different ways.

When he decides that he is “ready”, he can bid from 0 to 10 Credits on the color he wants. The true card color is shown: if he was right, he gains two times the bid; otherwise he loses the bid. The game lasts an indefinite number of turns; between the turns, depending on the result, the agent can revise the *reliability* he attributes to his sources: his perception and the witnesses (depending for example on the number of correct answers they provided) as well as his SCAI. Besides, his perception and the witnesses have *true reliability* values that determine the average correctness of their answers. True values are not known by the player; at the first game round they are initialized and they do not change during the game. At the end of the game the agent will collect a certain amount of Credits; a set of reliability values for his sources; a SCAI and he will have spent a

certain amount of Time. Using this game as a testbed, we compare different kinds of agents having many possible heuristics in order to collect information and Credits¹.

Theoretical Foundations

We present first an *evidential* notion of ignorance: ignorance is determined by the lack of belief sources. In our approach, following the cognitive approach of Castelfranchi (1995), the *strength of a belief* (i.e. how much I rely on one of my beliefs) depends on the *reliability* of its sources (i.e. the beliefs it is grounded on). Sources include: direct experience (such as perceptive evidences); information provided by other agents; reasoning (about other beliefs) and categorization (reasoning about classes and similarities). Since in an open world in principle there are infinite sources to take into account, the agent can never conclude that his own ignorance is zero. We propose in the present model a solution to the problem of ignorance quantification by identifying *Classes of Ignorance Acceptance* that reduce ignorance to finite values.

Uncertainty and contradiction have the same status of the notion of ignorance: they are meta-cognitive notions, i.e. agent's evaluations about his own "epistemic state". We are interested in this paper in investigating how those different epistemic states affect the decision process, and especially how they affect the way the agent decides to execute a pragmatic action (e.g. bet) or an epistemic one (e.g. query). For example, if the agent feels to be too much ignorant or uncertain he can decide to query, to bid a little amount, or not to bid at all. Here we describe the epistemic dimensions.

Ignorance

Intuitively ignorance depends on how much information I have with respect to how much it exists; in an open world there is a potentially infinite number of witnesses that have not been questioned; so if we calculate ignorance in this way the agent has always the maximum degree of ignorance. *The agent does not know how many witnesses he can consider at most or better he does not know how he can reduce his ignorance close to zero.* A qualitative and cognitive analysis is here required. Here we shift the issue to an evidential and subjective level². We introduce the notion of *Structure of Classes of Acceptable Ignorance* (SCAI).

¹Since it is an "open world" (there are an infinite number of witnesses and an indefinite number of turns) it is not possible to perform full search. More, it is not possible to perform a long-term maximization because the agents don't know when the game will end (this condition is called "shadow of the future").

²Our notion of ignorance is very close to the notion of ambiguity identified in some recent economical and psychological literature where is stressed that decision making is affected by the decision maker's evaluation of his or her actual available information and competence to make judgments in specific domains (Heat & Tversky, 1991). Instead, our notion of ignorance is quite far from *Sample Space Ignorance* in Support Theory (Tversky & Koehler, 1994) where it is claimed that people do not follow the extensional logic of conventional probability theory. In Support Theory an agent can actually "ignore" actual information in the sense that he

Classes of Acceptable Ignorance

Each agent has a SCAI that includes several *Classes of Acceptable Ignorance* (CAI) that include one or more sources (e.g. Witnesses), each having its reliability value. For instance, $CAI_1 = (\text{witness 1, witness 2, witness 3})$ could be one of those classes. Classes of acceptable ignorance can be intersected and unified (see Fig. 1): they have the normal properties of sets in set theory.

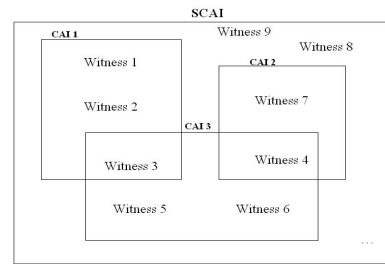


Fig. 1 Structure of Classes of Acceptable Ignorance (SCAI).

The agent knows that testing all witnesses in a given class is enough for making the ignorance acceptably close to value zero. Imagine for example that the agent wants to know if tomorrow will rain or will be sunny. He has several classes of acceptable ignorance. For instance he can believe that by acquiring information about tomorrow's weather from source 1 = "New York Times" and source 2 = "CNN" is enough for making ignorance acceptable. Moreover, the following points are crucial for understanding how the relation between SCAI and agents works.

A. The agent has explicit models (meta-level) of Classes of Acceptable Ignorance as shown in Fig.1. There are witnesses who are included in classes of acceptance but also witnesses who are not included in any class.

B. The agent can also make "queries" to witnesses who are not in the SCAI. Indeed the number of classes is finite even if, according to the agent, the set of witnesses he can make a query is indefinite. The agent can make a query to whatever witness even if this witness is not included in the structure of ignorance. In principle the agent can ask witness_10000 and he will always get an answer. Witnesses that are not in the structure do not have a value of reliability. A default value is assigned to them (through feature value assignment) whenever a query is made to them. After the query the witness belongs to the SCAI as a witness who is not included in any CAI (for example witness 8 in Fig.1).

C. The value of Ignorance is calculated at a meta-level whereas the value of reliability of a witness is calculated at a base-level.

is not explicitly evaluating that evidences concerning a certain event e1 are also evidences concerning another event e2. Indeed it has been shown that unpacking (making information available for explicit evaluation) a compound event into disjoint components tends to increase the perceived likelihood of that event. An immediate implication is that unpacking an hypothesis and/or repacking its complement will increase the judged likelihood of that hypothesis.

Quantifying Ignorance through SCAIs

Class-Ignorance is given for each class at a certain point of a query sequence (q_1, \dots, q_n) and is defined as the total number of witnesses in the class *minus* the number of tested witnesses in that class, weighted for the inverse of the total number of witnesses in the class.

Absolute-Ignorance is defined as the minimal value of Class-Ignorance among all CAIs.

$$\begin{aligned} \text{Class-Ignorance (Class } n, q_i) &= \\ & (n.\text{wit. (Class } n, q_i) - n.\text{queried.wit. (Class } n, q_i)_{\text{Agent } x, q_i}) / \\ & n.\text{wit. (Class } n, q_i)_{\text{Agent } x, q_i} \\ \text{Absolute-Ignorance (} q_i) &= \\ & \text{Min}_{\text{Class } x} (\text{Class-Ignorance (Class } x, q_i)_{\text{Agent } x, q_i}) \end{aligned}$$

We have already pointed out that after a query is made to a witness who does not belong to SCAI, the witness will be included in SCAI as a witness who is not included in any class (such as witness 8 and 9 in Fig.1). The measure of Absolute-Ignorance is not fixed: it depends on the single agent categorization and classes organization. That measure varies through learning, as new witnesses are added.

Uncertainty

Uncertainty is a measure of the difference between the value of strength of the belief “the card is red” and the value of strength of the belief “the card is blue”. When the difference is 0 the value of uncertainty is maximal, when the difference is 1 the value uncertainty is minimum. This dimension takes into account the difficulty of deciding when the two strengths of beliefs are too close.

Contradiction

Contradiction is a (logical) inconsistency in a belief set; for example I can not believe consistently that (in the previous example) the ball in an urn is both red and black. In a normal (statistical) analysis there is contradiction if the sum of the two strengths of beliefs is more than 1. In the evidential approach the threshold of perceived contradiction (α) can be fixed at different values depending on cognitive biases (e.g. more or less contradiction tolerant).

$$\begin{aligned} \text{If}(\text{Strength.belief(CardRed, } q_i) + \\ \text{Strength.Belief(CardBlue, } q_i)) \leq \alpha \text{ then} \\ \text{Perceived.contradiction}(q_i) = 0 \end{aligned}$$

$$\begin{aligned} \text{If}(\text{Strength.belief(CardRed, } q_i) + \\ \text{Strength.Belief(CardBlue, } q_i)) > \alpha \text{ then} \\ \text{Perceived.contradiction}(q_i) = \\ (\text{Strength.belief(CardRed, } q_i) + \\ \text{Strength.Belief(CardBlue, } q_i)) - \alpha \end{aligned}$$

Perplexity

Ignorance, Uncertainty and Contradiction are three meta-level epistemic information that an agent can take into account in order to “decide if he is ready to decide”. In order to model this kind of decisions we propose to integrate ignorance, uncertainty and contradiction in a single measure called **Perplexity** (i.e. lack of confidence). In calculating Perplexity, the three dimensions can be aggregated in

different ways, depending on some more cognitive biases (e.g. Agents that are biased to consider ignorance, or contradiction, or uncertainty). The basic heuristic is summing them (and normalizing).

Value of Information: Epistemic Actions

An Epistemic Action (EpA) is *any action aimed at acquiring knowledge from the world*; any act of active perception, monitoring, checking, testing, ascertaining, verifying, experimenting, exploring, enquiring, give a look to, etc. (Castelfranchi & Lorini, 1998). The notion of epistemic action has been extensively considered both in psychology and in economics. The centrality of this notion comes from the fact that epistemic actions have a role in different cognitive functions. In the present model an Epistemic Action is always towards a witness (i.e. making a query). Epistemic Actions are directed either to reduce perplexity (or one of its dimensions) given a certain “perplexity aversion” threshold of the agent (first function); or to acquire new information in order to make a better decision (second function).

In both cases a value is assigned to epistemic actions. The first value is a measure of the capacity of a given witness of reducing perplexity: we call it *informativeness*. The second value is called *value of information* and has been extensively investigated in economical literature in the sense of “how much the agent is disposed to pay for obtaining that information?” In that approach a possible way to calculate the value of information is given with respect to utility functions. These two notions can lead to different decision strategies; in order to compare them, we have designed the simulative testbed “Red-or-Blue Card Game” (see above).

Considering the Sources

Strength of beliefs depends on its sources (perception; more or less reliable witnesses). Those sources are not all equal: in order to represent their relative contribute, we aggregate them using Fuzzy Cognitive Maps (Kosko, 1986). In order to represent the fact that there are diverging sources (and they aggregate in a different way with respect to converging ones) our FCMs have two “competing” branches for representing the competing beliefs “the card is red” and “the card is blue”. FCMs are additive fuzzy systems with feedback, having nodes and edges. The weight of the nodes represents the strength of a belief (e.g. “I am pretty sure that the card is red”); the edges are weighted and they represent the impact of a belief over another. The FCM that we use can be seen as divided into two branches, each aggregating the values either for “red” or “blue”. These nodes receive input from intermediate nodes (“perception for red” and “witnesses for red” the first; “perception for blue” and “witnesses for blue” the second); these edges are weighted by two fixed factors κ, λ representing the relative impact of perception and witnesses. The nodes “perception for red” and “perception for blue” assume either the value 0 or 1 depending on the perceptual input; their edges have the value of *perception reliability* (according to the agent). The “witnesses for red” and “witnesses for blue” nodes receive as input the information of the queried witnesses (either 0 or

1); the edges between each witness and “witnesses for <color>” have the value of the *witnesses’ reliability*. There are also negative-weighted edges between the “red” and “blue” nodes, as well as for each source. In this way the contribute of diverging sources is modeled, because each positive evidence in a branch counts also as a negative one in the other branch. So, starting from the input values (the contributes of perception and the queried witnesses) the FCM calculates the final values for the strength of belief in “the card is red” and “the card is blue”. There is not fixed “sum 1” between the two final values, so it is possible to model contradictory beliefs (that the agent can reduce performing epistemic actions). The FCM structure is the same for all the agents, but at each step it can be updated (e.g. modifying the impact of the edges, i.e. reliability values).

Player Agents and Decision Strategies

Here we describe three classes of decision strategies, implemented into three Agents. *Normative Agent* and *Satisficing Agent* do not use the notions of ignorance, uncertainty and contradiction. *Perplexity Reducing Agent* uses them in order to select the witness to be tested.

Normative Agent

A normative agent decides either to bet a certain amount of credits on a given option (either “the card is red” or “the card is blue”) or to make a query to a specific witness as follows (this agent is not affected by perplexity).

The agent calculates the value of information obtainable from a given witness for all witnesses in the Structure of Classes of Ignorance Acceptance. *The value of information obtainable from witness z* is determined as: the average of the max value of expected utility given the information “the card is red” given by wit. z (which impacts on the agent’s beliefs) and the max value of expected utility given the information “the card is blue” given by wit. z *minus* the max value of expected utility given the actual information. Afterwards the agent is able to decide. If the max value of information obtainable from witnesses is more than 0 then the agent decides to make a query to the witness who maximizes that value; otherwise he decides to bet a quantity y of credits on “the card is x” that maximizes his actual expected utility. We have not included the costs of making a query in the utility function (we assume only the cost in Time).

$$\begin{aligned} \text{Potential-Chosen-Bet } (q_i) &= \text{BET.yONx} \\ \text{such that} \\ \text{Max}_{x,y} (\text{Strength.belief } (x, q_i) \text{ credits } (y, q_i)) \end{aligned}$$

where x is either “the card is blue” or “the card is red” and y is whatever sub-amount of the total amount of credits at a given point in the query sequence (q_1, \dots, q_n) . This agent has a very time consuming policy (minimizing the lack of information) and is not well suited for real time situations. Another agent can be introduced that limits Time spent.

$$\begin{aligned} \text{Value-Information } (\text{wit.z, } q_i) &= \\ (\text{Max}_{x,y} ((\text{Strength.belief } (x, q_{i+1}) \leftarrow \\ \text{speech } (\text{wit.z, CardRed, } q_{i+1})) \text{ credits}(y, q_i)) + \\ \text{Max}_{x,y} ((\text{Strength.belief } (x, q_{i+1}) \leftarrow \end{aligned}$$

$$\begin{aligned} \text{speech } (\text{wit.z, CardBlue, } q_{i+1})) \text{ credits}(y, q_i)) / 2 - \\ \text{Max}_{x,y} (\text{Strength.belief } (x, q_i) \text{ credits } (y, q_i)) \end{aligned}$$

Effective-Choice (q_i) =

1. If $\text{Max}_{\text{wit.z}} (\text{Value-Information } (\text{wit.z, } q_i)) > 0$ then
Effective-choice (q_i) = QUERY.wit. z such that
 $\text{Max}_{\text{wit.z}} (\text{Value-Information } (\text{wit.z, } q_i))$
2. If $\text{Max}_{\text{wit.z}} (\text{Value-Information } (\text{wit.z, } q_i)) \leq 0$ then
Effective-choice (q_i) = BET.yONx such that
 $\text{Max}_{x,y} (\text{Strength.belief } (x, q_i) \text{ credits } (y, q_i))$

Satisficing Agent

The Satisficing Agent makes sequential search through the witnesses in his SCAI. He starts with a given threshold γ for expected utility. At each step, he randomly calculates either the expected utility value associated with BET.yONx or the expected utility value associated with BET.yONx after that a given witness will be questioned. This value is the average of the expected utility value associated with BET.yONx in case the witness will say “Red Card” and the expected utility value associated with BET.yONx in case the witness will say “Blue Card”. The first option during the sequential search that overcomes threshold γ is chosen. If no suitable option is found after n (fixed value) steps, the agent lowers the threshold of a certain value $\Delta\delta$. With respect to the Normative Agent, the Satisficing Agent makes less queries and it is better suited for open worlds (Simon, 1990).

Perplexity Reducing Agent

The Perplexity Reducing agent has the goal to reduce the level of perplexity below a given threshold δ before betting. Since the only way to reduce perplexity is through queries, the agent starts choosing the witness to test: he makes a sequential search on witnesses and takes the first witness whose information is able to reduce perplexity under the threshold. If not suitable witness is found, the agent reduces the value of the threshold of a certain value $\Delta\delta$ and restarts with the same strategy. The expected capacity of a witness of reducing (or augmenting) perplexity represents the expected informative contribute of the epistemic act of querying him. This value is called *expected informativeness* and it is calculated as the actual value of perplexity *minus* the average of the value of perplexity after that witness z says “the card is red” and the value of perplexity after that witness z says “the card is blue”³.

$$\begin{aligned} \text{Expected-Informativeness } (\text{wit.z, } q_i) &= \\ (\text{Subj.unconfidence}(q_i)) - \\ ((\text{Subj.unconfidence}(q_{i+1}) \leftarrow \text{speech } (\text{wit.z, CardRed, } \\ q_{i+1})) + (\text{Subj.unconfidence}(q_{i+1}) \leftarrow \text{speech } (\text{wit.z, } \\ \text{CardBlue, } q_{i+1}))) / 2 \end{aligned}$$

Expected informativeness is quite different from the *value of information* as defined for the Normative Agent. The difference between those two definitions indicates two different theoretical perspectives: while the Normative Agent maximizes utility values (for bidding) the Perplexity Reducing Agent uses a cognitive theory of sources in order

³ It follows from the definition that there could be negative values of expected informativeness.

to consider the contribute of the witnesses in the cognitive dimensions of uncertainty, ignorance and contradiction. The Perplexity Reducing Agent is implicitly biased to make queries to witnesses that are in the CAIs, since by definition they lower the value of absolute ignorance more than witnesses that are not in any CAIs. The Perplexity Reducing Agent should be combined with the two others agents (Normative or Satisficing). Once the level of perplexity is under the threshold, he could decide either which color and how much to bid or decide to make a query to another witness using his optimization methods. However, in order to simplify our experiments we did not allow perplexity reducing agents to carry on making queries to witnesses once the degree of perplexity was reduced under the threshold δ . This simplification is plausible for maintaining completely distinct the 2 different functions of epistemic actions: the function of *perplexity reduction* and the function of “*increase*” of *expected utility*.

Learning During the Game

The RBG game has many turns, so it is possible to learn between them. In the epistemic perspective, it is interesting to model how agents revise information about sources of beliefs.

Updating Reliability Values

All the agents have a representation of the witnesses reliability and are able to update these values depending on past interactions. Since reliability updating strategies are outside the scope of this paper, we used a linear statistical heuristic for all players: witnesses' reliability is lowered if they furnished a wrong advice, augmented otherwise, of a fixed amount $\Delta\phi$.

Updating Classes of Acceptable Ignorance

The Perplexity Reducing Agent is also able to change its SCAI adding or removing the witnesses in the Classes of Acceptable Ignorance. At the beginning of the game the SCAI is set randomly (e.g. the one shown in Fig.1) and it can be updated after each turn extending or contracting its CAIs. Imagine that the agent has queried in sequence w_1, w_2, w_3, w_4, w_8 before deciding. Imagine he has verified that after the second test the value of perplexity has not changed so much (i.e. less than a threshold α). Since w_1 and w_2 belong to the same Class1, Class1 can be *contracted* eliminating w_2 (that resulted not very informative). Imagine also he has verified that after the fifth test the value of perplexity has changed quite a lot (over a threshold β). Since w_4 and w_8 do not belong to the same Class, the class of w_4 can be *extended* adding w_8 , that proved to be so informative. We do not describe here the full algorithm for CAIs contraction and extension⁴. We want only to present verbally its structure.

⁴ The variable ϕ for reliability updating, as well as thresholds α e β in classes of Acceptable Ignorance updating depend from cognitive biases towards belief revision. It is relevant to notice that for

1. Given a previous sequence of queries (q_1, \dots, q_n), if during that sequence there were two queries q_i and q_{i+1} for *witness A* $\leftarrow q_i$ and *witness B* $\leftarrow q_{i+1}$ and witness A and witness B belong to the same CAI x and the degree of perplexity did not vary so much (in absolute value given threshold α) from q_i to q_{i+1} then the *witness B* is taken out from CAI x .
2. Given a previous sequence of queries (q_1, \dots, q_n), if during that sequence there were two queries q_i and q_{i+1} for *witness A* $\leftarrow q_i$ and *witness B* $\leftarrow q_{i+1}$ and witness A belongs to CAI x whereas witness B belongs to the SCAI but not to CAI x , and the degree of perplexity varied a lot (in absolute value given threshold β) from q_i to q_{i+1} then *witness B* is inserted into CAI x .

Experimental Setting and Variables

Here we show the comparison between three players: Normative (N), Satisficing (S), Perplexity Reducing (E). There are also two baselines: Random Bidder (B1) that chooses at random to test or to bid (and how much); and Perceptive Bidder (B2) that bids only according to his perceptive input.

The three independent variables we use are: *perception reliability (PR)*; *average witnesses reliability (AWR)*; *witnesses' convergence (WC)*. The first one describes how reliable in absolute is the perception of the agent; the second one indicates how reliable are in average the witnesses answers. They reflect also the “difficulty” of the task. The third one describes how convergent are the answers of the witnesses; this influences the final uncertainty value. We have built three scenarios: *good perception* (where PR is higher than AWR); *good witnesses* (the inverse); *high uncertainty* (where WC is set to a low value, and PR and AWR have the same value)⁵.

Results and Discussion

In the following tables we present the preliminary results of our experiments (for Credits and Time) of the three Scenarios (250 simulations, 100 bid turns)⁶. As an indirect

relatively high values of α and relatively low values of β and ϕ the agent is relatively *closed-minded* and conservative (he is less biased to revise the structure of classes of acceptance and the reliability values). But for relatively low values of α and relatively high values of β and ϕ the agent is relatively *open-minded*. This distinction is very close to the typology of cognitive epistemic styles in (Sorrentino et al., 1986).

⁵We use many thresholds and variables in our model: *Close Mind agents* vs. *Open Mind agents* in SCAI revision strategies (thresholds α and β); *strong* vs. *weak* need for low degree of perplexity (threshold δ); *degree of satisfaction* in expected utility (threshold γ); different way to weight different kinds of sources (bias towards perception or witnesses). In order to eliminate their effects we have randomly varied them through the experiments (three dimensions for each variable on average).

⁶The simulations were performed using the cognitive architecture AKIRA, developed at ISTC-CNR (<http://www.akira-project.org/>).

measure of “algorithm performance”, we introduced also Hypothesis Time: it measures how many witnesses an Agent has considered (but not questioned) before deciding.

Table 1: *Good Perception Scenario.*

Agent	Credits	Time	H. Time
B1	981	102	0
B2	1202	0	0
N	1641	6112	112453
S	1388	987	13936
E	1622	409	10681

Table 2: *Good Witnesses Scenario.*

Agent	Credits	Time	H. Time
B1	1009	101	0
B2	799	0	0
N	1306	9207	144582
S	1102	997	19103
E	1298	603	13190

Table 3: *High Uncertainty Scenario.*

Agent	Credits	Time	H. Time
B1	1007	99	0
B2	999	0	0
N	1803	8834	137866
S	1551	1156	21033
E	1563	673	15943

In the first and second Scenarios the Perplexity Reducing Agent performs very well with respect both to gained Credits and temporal measures (Time and Hypothesis Time): it performs at the same level of Normative agent with respect the final amount of credits but his temporal measures are much better. The comparison with the Satisficing agent is even better. Not surprisingly, in the third Scenario he needs to query more witnesses and it is not able to perform as the Normative. Results in **bold** are significant with respect to the Perplexity Reducing Agent. These results show that Perplexity Reducing Agents are very suited in open world conditions where search of new information is in general very costly.

Moreover, a qualitative analysis allows to get a nice result about SCAIs updating: the final SCAIs are in average populated with small CAIs of very reliable witnesses: the average reliability changes from 0.5 to 0.7 in the three scenarios and the number of witnesses remains less to 20 in all simulations. The fact that final CAIs are small and include reliable witnesses is in accordance with the way we learn about belief sources. The more you know an environment, the less you need to question. Moreover, you prefer to question very reliable sources.

Conclusions and Future Work

We have proposed a theoretical foundation of some cognitive categories such as ignorance, uncertainty and contradiction that are generally difficult to quantify in an

open world. We have introduced a MAS game (RBG) as a simulation setting in order to compare many agents that take or do not take into account epistemic dimensions. Our preliminary results show that perplexity reduction is a good heuristic for dealing with open world scenarios, and the Structure of Classes of Acceptable Ignorance can be used in order to quantify ignorance and reasoning about it. It would be interesting to test mixed decision strategies (e.g. considering the perplexity in the utility function; or using the Perplexity Reducing Agent as a filter). Another interesting direction is comparing simulation data with data from human experiments; actually the RBG game is being used as an experimental setting in order to collect such data.

Acknowledgments

We are indebted to Cristiano Castelfranchi for many insightful discussions about Ignorance and Uncertainty.

References

- Camerer, C., Weber, M. (1992). Recent Developments in Modelling Preferences: Uncertainty and Ambiguity. *Journal of Risk and Uncertainty*, vol. 5, pp. 325-370.
- Castelfranchi, C. (1995). Representation and integration of multiple knowledge sources: issues and questions. In Cantoni, Di Gesu', Setti e Tegolo (Eds.), *Human & Machine Perception: Information Fusion*, Plenum Press.
- Castelfranchi, C, Falcone, R. & Pezzulo, G. (2003). Trust in information sources as a source for trust: a fuzzy approach. *AAMAS 2003*: 89-96.
- Castelfranchi, C., Lorini, E. (2003). Cognitive Anatomy and Functions of Expectations. *IJCAI '03 Workshop on Cognitive modeling of agents and multi-agent interaction*, Acapulco, Mexico.
- Ellsberg, D. (1961). Risk, Ambiguity and the Savage axioms. *Quarterly Journal of Economics*, 75, pp. 643-669.
- Heath, C., Tversky, A. (1991). Preference and belief: Ambiguity and competence in choice under uncertainty. *Journal of Risk and Uncertainty*, 4, pp. 5-28.
- Kirsh, D., Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*, 18, pp. 513-549.
- Kosko, B. (1986). Fuzzy Cognitive Maps. *International Journal Man-Machine Studies*, vol. 24, pp. 65-75.
- Shackle, G. L. S. (1972). *Epistemics and Economics*. Cambridge: Cambridge University Press.
- Shafer, G. (1976). *A Mathematical Theory of Evidence*. Princeton: Princeton University Press.
- Simon, H. (1990). Invariants of human behavior. *Annual Review of Psychology*, 41, pp.1-19.
- Sorrentino, R. M., Short, J. C. (1986). Uncertainty orientation, motivation, and cognition. In R. M. Sorrentino, E. T. Higgins (Eds.), *Handbook of motivation and cognition: Foundations of social behaviour*, vol. 1, Guilford Press, New York, pp. 379-403.
- Tversky, A., Koehler, D. J. (1994). Support theory: A nonextensional representation of subjective probability. *Psychological Review*, 101, pp. 547-567.