Intention Reconsideration like uncertain dichotomous choice model

Cecilia Sosa Toranzo, Marcelo Errecalde, and Edgardo Ferretti

Laboratorio de Investigación y Desarrollo en Inteligencia Computacional
Departamento de Informática, Universidad Nacional de San Luis
Ejército de los Andes 950 - (D5700HHW) San Luis - Argentina
e-mails:{csosatoranzo, merreca, ferretti}@unsl.edu.ar

Abstract. A key issue in the design of Belief-Desires-Intentions (BDI) agents is that of finding an appropriate strategy for Intention Reconsideration (IR). Traditional approaches to IR defines the policy in the agent's design stage, which makes impossible to modify it in execution time. This is clearly not a practical solution for agents operating in dynamic and changing environments. Besides, IR typically involves considering multiple criteria. That is why, in this work, we propose a novel approach to IR based on a dichotomous choice model. This approach allows changing commitments to intentions depending on how the environment evolves and involves multi-criteria aggregation for IR.

Keywords: Intention Reconsideration, BDI Architecture, Dichotomous Choice, Voting

1 Introduction

The BDI model, as a practical reasoning architecture aims at making decisions about what to do based on cognitive notions as Beliefs, Desires and Intentions. A very important design issue in BDI agents concerns defining the intention reconsideration policy [1, 2]. This policy will define under which circumstances the BDI agent will use computational resources to deliberate about its intentions. At present there is no consensus on when or how an agent should reconsider its intentions. Current proposals consider the agents' *commitment levels*, which range from *cautious* agents (which reconsider their intentions after each action execution) to *bold* agents (where no reconsideration is performed until the current plan has been completely executed).

In [3], the efficiency of these policies in different kind of environments are investigated, but the intention reconsideration policy is defined in the agent's design stage, which makes impossible to modify this policy in execution time. It is clear that this is not a practical solution for agents operating in dynamic and changing environments. In this respect, in [2], a framework is proposed that allows the agent choosing by itself what policy to follow based on the current state of the world. The main idea underlying this work is that an intention reconsideration (IR) policy can be conceived as a meta-level control process

which selects whether to deliberate or act. This proposal is based on the *discrete deliberation scheduling* framework [4], where deliberations are treated as if they were actions.

In [2], the proposed model incorporates decision making within the IR process of a BDI agent. To determine the best possible action, the maximum expected utility is considered; that is, only one criterion is taken into account to solve the decision problem. However, IR typically involves considering multiple criteria. In this respect, Sosa et al. [5] extends [2] by applying voting-based multi-criteria decision making to choose whether to act or to deliberate. In [5], multi-criteria decision making is applied to make background decisions at different stages within BDI architecture, namely: choosing among conflicting desires, choosing between plans and intentions reconsideration. In opposition, our work solely focuses on posing IR as a multi-criteria voting-based approach, a very relevant step that was not fully approached by Sosa et al.

It is worth mentioning, that voting sucessfully works on problems where heterogeneity in individual preferences exist. However, various problems associated with the aggregation of individual decisions also arise when individuals share identical preferences, but have to make decisions given their different decision capabilities, i.e., the problem of aggregating individual decisions under an uncertain dichotomous choice setting. When individuals sharing identical preferences operate in an uncertain environment, the problem of preference aggregation is no longer relevant, but the issue of choice aggregation is still pertinent. In an uncertain environment individuals may have identical or different decision skills, but in any case they might fail and make incorrect decisions. Therefore, in this context, the aggregation problem is conceived as "What is the most appropriate collective decision rule for implementing the common objective of the individual decision makers?". The answer to this question is developed in Sect. 3.1, as proposed by Nitzan and Paroush [6, 7].

We argue, that if IR is conceived as a meta-level control process which selects whether to deliberate or act, and multiple criteria are involved, the IR strategy can be modeled as an uncertain dichotomous choice setting [6]. This is the main proposal of our work and it is formalized in Sect. 4.

Our model considers voting schemes whereby each voter is required to opt for one of two available alternatives (to deliberate or to act, in the case of IR), and assumes a set of individuals with identical objectives but whit possibly different abilities to identify the alternative for the attainment of their common will (IR being optimal). Since individual preferences are identical, one of the two alternatives is preferred by all of them, but in the current setting, the identification of that alternative, "the correct alternative" or "the most preferred alternative" is unknown.

The rest of the paper is organized as follows. Section 2, briefly introduces the BDI model to provide the background concepts underlying intention reconsideration. Section 3, introduces the uncertain dichotomous choice model and the optimal voting rule that will be applied within the IR mechanism (Sect. 4). Finally, Sect. 5 draws the conclusions and briefly describes possible future work.

2   BDI Model

Belief-Desires-Intentions (BDI) models have been inspired from the philosophical tradition on understanding *practical reasoning* and were originally proposed by Bratman et al. [8]. This kind of reasoning can be conceived as the process of deciding what action perform next to accomplish a certain goal. Practical reasoning involves two important processes, namely: deciding *what* states of the world to achieve and *how* to do it. The first process is known as *deliberation* and its result is a set of intentions. The second process, so-called *means-ends reasoning* involves generating actions sequences to achieve intentions.

The mental attitudes of a BDI agent on its beliefs, desires and intentions, represent its informational state, motivational state and decision state, respectively. The BDI architecture defines its cognitive notions as follows:

- **Beliefs:** Partial knowledge the agent has about the world.
- **Desires:** The states of the world that the agent would ideally like to achieve.
- **Intentions:** Desires (states of the world) that the agent has been *committed* (dedicated resources) to achieve.

These cognitive notions are implemented as data structures in the BDI architecture, which also has an interpreter that manipulates them to select the most appropriate actions to be performed by the agent. This interpreter performs the deliberation and means-ends reasoning processes aforementioned, and it is shown in Algorithm 1 (as proposed in [1]). It assumes that an explicit representation of desires (D), belief (B) and intentions (I) exist, within the agent. The agent's perceptions will be represented by ρ and regarding plans, they will be referred as "recipes" to achieve intentions. Therefore, π will be used to denote plans. Lines $1-3$ initialize beliefs, intentions and the plan. The main control loop comprises lines $4-19$. In lines $5-6$, the agent perceives and updates its beliefs while in line 7 it decides whether to reconsider or not. In lines $8-12$ it deliberates and in line 11, it generates a plan to achieve its intentions. Lines $14-17$ show that an action of the current plan is executed. Because the purpose of the functions in this loop can be easily derived from their names, due to space constraints we omit their formalizations but the interested reader should refer to [1].

2.1   Intention Reconsideration

Intentions can be adopted, maintained and modified. Intentions are maintained by means of a commitment strategy. Intention maintenance concerns the commitment to single intentions that have already been adopted; whether an intention should be abandoned or maintained. Intention modification happens "in the light of changing circumstances" through re-deliberation. Intentions resist change, but there are conditions under which a rational agent must consider whether its intentions are still worth committing to. That is, an agent should *reconsider its intentions* when it is reasonable to do so. However, IR is a computationally costly process, and is a kind of meta-level reasoning. It is therefore necessary to

---

Algorithm 1 BDI Agent control loop

1: B ← B_0
2: I ← I_0
3: π ← $null$
4: **while** true **do**
5:     get next percept ρ
6:     update B on the basis of ρ
7:     **if** reconsider(B,I) **then**
8:         D ← options(B,I)
9:         I ← filter(B,D,I)
10:        **if not** sound(π,I,B) **then**
11:            π ← plan(B,I)
12:        **end if**
13:    **end if**
14:    **if** $\pi \neq \emptyset$ **then**
15:        α ← hd(π)
16:        execute(α)
17:        π ← tail(π)
18:    **end if**
19: **end while**

---

fix upon an IR strategy that makes optimal use of the available computational resources. Whereas a commitment strategy says when an individual intention should be kept or dropped, a reconsideration strategy says when to deliberate. Reconsideration is equivalent to re-deliberation. That is, when an agent decides to reconsider, it activates its deliberation process.[1]

Situation number	Chose to deliberate?	Changed intentions?	Would've changed intentions?	$reconsider(\cdot,\cdot)$ optimal?
1	No	-	No	Yes
2	No	-	Yes	No
3	Yes	No	-	No
4	Yes	Yes	-	Yes

Table 1: Interactions between $reconsider(\cdot,\cdot)$ and deliberation.

To try capture trade-off between to deliberate or to act (continue with current intentions), Algorithm 1 incorporates an explicit meta-level control component: $reconsider(\cdot,\cdot)$ function. Possible interactions between deliberation and meta-level control (whether to deliberate again)[2] are summarized in Table 1. The analysis elucidates when reconsideration is optimal([8, 9]). In situation 1, the agent does not deliberate, but if it did, it would not have changed its intentions anyway. This situation is desirable. In situation 2, the agent does not deliberate, but if it did, it would have changed its intentions. Here the agent gets bad

---

[1] In Algorithm 1, deliberation process is considered as composed by $options(\cdot,\cdot)$ and $filter(\cdot,\cdot,\cdot)$ functions.

[2] It is assumed that the cost of executing the reconsider function is much less than the cost of the deliberation process itself.

advice from $reconsider(\cdot, \cdot)$. The agent chooses to deliberate in situations 3 and 4. When it does not change its intentions in situation 3, the agent is wasting time deliberating. The $reconsider(\cdot, \cdot)$ function is not behaving optimally. The agent does change intentions in situation 4, which means it was a good idea to deliberate, and $reconsider(\cdot, \cdot)$ has done well. In conclusion, the purpose of $reconsider(\cdot, \cdot)$ function is to deliberate when it pays off to deliberate (i.e., when deliberation will lead to change intentions), and otherwise not to deliberate, but to act.

As mentioned above, our aim in this paper is to consider the question of when to deliberate (i.e., to reconsider intentions) versus when to act (viz. continue to execute its current intentions) like a dichotomous choice situation.

3 Uncertain dichotomous choice model

Nitzan and Paroush [6, 7] focus on dichotomous choice situations in which a set of n individuals is required to select one of two alternatives such as "support" or "reject" a certain proposal or answering "yes" or "no" to a certain question. Let us consider two generic alternatives a and b. The choice of individual i is represented by the decision variable x_i, where $x_i = 1$ means that individual i chooses alternative a, and $x_i = -1$ means that alternative b is chosen instead. The decisions of the group members are represented by a decision profile $x = (x_1, \cdots, x_n)$. In the dichotomous setting, the collective decision rule f assigns to any possible profile x one of the values $\{1, -1\}$, where $f(x) = 1$ or -1 means that, given the profile x, the group chooses alternative a or b, respectively.

In the uncertain dichotomous choice model, there are two possible states of nature: "1" the state of nature in which a is the correct alternative that should be chosen, and "-1" the state of nature where b is the correct choice.[3] We denote by $B(t; s)$ the utility from choosing t in state of nature s. The common preferences of the individuals are thus represented by the payoff matrix B:

$$\begin{bmatrix} B(a; 1) & B(b; 1) \\ B(a, -1) & B(b, -1) \end{bmatrix}$$

Let α and $1-\alpha$, with $0 \le \alpha \le 1$, denote the a-priori probabilities of realization of states of nature 1 and -1, respectively.[4] Let p_i denote the probability of the i-th expert making the right choice.[5] We assume that individuals are independent in their choices and vector $p = (p_1, p_2, \cdots, p_n)$ is called the vector of abilities or skills. Without loss of generality, we assume that if $i < j$, then $p_i \ge p_j$. This means that individual 1 can be referred to as the expert and individual n as the least competent.

---

[3] Given state of nature 1, a is the correct alternative if $B(a; 1) > B(b; 1)$. Given state of nature -1, b is the correct alternative if $B(b; -1) > B(a; -1)$.

[4] Notice that, α and $1 - \alpha$ are the a-priori probabilities that a and b are the correct alternatives.

[5] We assume that this probability is equal in the two states of nature and $p_i > 1/2$.

3.1   The optimal decision rule

The optimal collective decision rule f^* which maximizes the expected utility of every individual i, given the decision profile x, the a-priori probabilities of the two states of nature α and $1 - \alpha$, the payoff matrix B, and the decision skills of individuals, $p = (p_1, ..., p_n)$, is:

$$f^* = sign(\sum_{i=1}^{n} \beta_i x_i + \gamma + \delta) \qquad (1)$$

where $\quad \delta = \ln \dfrac{B(1)}{B(-1)}, \quad \gamma = \ln \dfrac{\alpha}{1 - \alpha}, \quad \beta_i = \ln \dfrac{p_i}{1 - p_i}, \quad sing(t) = \begin{cases} 1 & t > 0 \\ -1 & t < 0 \end{cases}$

$B(1) = B(a; 1) - B(b; 1)$ is the net utility of making a correct decision in state of nature 1 and $B(-1) = B(b; -1) - B(a; -1)$ is the net utility of making a correct decision in state of nature -1. The optimal decision rule f^* is a weighted qualified majority rule f_{q^*} (function 2), such that $q^* = (\gamma + \delta)/\sum_{i=1}^{n} \beta_i$.

$$f_{q^*}(x) = \begin{cases} -1 & (- \sum_{i=1}^{n} \beta_i x_i / \sum_{i=1}^{n} \beta_i) \geq q^* \\ 1 & \text{otherwise} \end{cases} \qquad (2)$$

The two parameters γ and δ determine the desirable bias in favor of one of the alternatives (alternative a in this case). The parameter γ specifies the asymmetry between the a-priori probabilities of the two states of nature. The parameter δ specifies the asymmetry between the net utilities obtained when making a correct decision in the two states of nature.

   If individual decision skills are identical, then the optimal decision rule f^* is a qualified majority rule f_{k^*} (function 3) where $N(b)$ is the number of individuals choosing alternative b, and $k^* = \frac{1}{2}(1 + \frac{\gamma + \delta}{n\beta})$.

$$f_{k^*}(x) = \begin{cases} -1 & N(b) \geq nk^* \\ 1 & \text{otherwise} \end{cases} \qquad (3)$$

   When there is symmetry between the states of nature, $\gamma + \delta = 0$, the optimal collective decision rule is the weighted majority rule (function 4). In addition, if the individual decision skills are identical, the optimal collective decision rule is the simple majority rule (function 5). When there is symmetry between the states of nature, the expert rule (function 6) is optimal if and only if $\beta_1 > \beta_2 + \cdots + \beta_n$.

$$f^* = sign(\sum_{i=1}^{n} \beta_i x_i) \quad (4) \qquad f^* = f^m = sign(\sum_{i=1}^{n} x_i) \quad (5) \qquad f^e(x) = x_1 \quad (6)$$

3.2   Partial information on the individual decision skills

The optimal decision rule (function 1) depends on several parameters, in particular, on the decision skills of individuals. These skills have been assumed to be known, although quite often information on these skills is very hard to get and skills estimation can be difficult. A more plausible assumption is that individual

decision skills are only partly known. Thus we assume that the correctness probabilities p_i or, equivalently, the logarithmic expertise levels β_i are independent random variables, distributed according to some known distribution function. If the ranking of the members in the group is known, then one can follow rules based on this ranking. The extreme cases are the majority and expert rules (functions 5 and 6).

A general comprehensive study of weighted majority rules is a very complicated task, since the class of such rules becomes very large as the number of group members increases.[6] Table 2 provides the probabilities of weighted majority rules to be optimal under the assumption of uniform distribution on $[1/2; 1]$ of the p_i's, obtained by Nitzan and Paroush [10]. Table 3 represents the optimal probabilities of the same rules under the assumption of exponentially distributed logarithmic expertise levels β_i, as reported in [11]. In both cases the expert rule is far more likely to be optimal than the majority rule. Note that for $n = 5$ the "leader" was the expert rule, $(1, 0, 0, 0, 0)$, for the exponential distribution while for the uniform distribution, the "leader" was a rule "close" to the expert rule, called balanced expert rule, $(2, 1, 1, 1, 0)$. In both situations, the "loser" was the majority rule, $(1, 1, 1, 1, 1)$.

n			Rule and corresponding probability				
3	(1,0,0)	(1,1,1)	-	-	-	-	-
	0.675	0.325	-	-	-	-	-
4	(1,0,0,0)	-	(1,1,1,0)	-	(2,1,1,1)	-	-
	0.373	-	0.277	-	0.350	-	-
5	(1,0,0,0,0)	(1,1,1,1,1)	(1,1,1,0,0)	(3,1,1,1,1)	(2,1,1,1,0)	(3,2,2,1,1)	(2,2,1,1,1)
	0.199	0.022	0.175	0.107	0.229	0.194	0.074

Table 2: Optimality probabilities of all weighted majority rules for $n \leq 5$ for $p_i \sim U[\frac{1}{2}; 1]$

n			Rule and corresponding probability				
3	(1,0,0)	(1,1,1)	-	-	-	-	-
	0.75	0.25	-	-	-	-	-
4	(1,0,0,0)	-	(1,1,1,0)	-	(2,1,1,1)	-	-
	0.5	-	0.25	-	0.25	-	-
5	(1,0,0,0,0)	(1,1,1,1,1)	(1,1,1,0,0)	(3,1,1,1,1)	(2,1,1,1,0)	(3,2,2,1,1)	(2,2,1,1,1)
	0.312	0.010	0.157	0.104	0.208	0.157	0.052

Table 3: Optimality probabilities of all weighted majority rules for $n \leq 5$ for $\beta_i \sim E(\lambda)$

The probabilities $P_e(n)$ and $P_m(n)$ of the expert and majority rules being optimal, respectively, were calculated or estimated in a series of papers for a variety of distributions (in [12] a summary is shown). The comparison of $P_e(n)$ and $P_m(n)$ shows that typically (but not for any distribution) the expert rule has a much better chance of being optimal than the majority rule, especially for large n.

---

[6] For example, for committee size $n = 3$ it includes 2 weighted majority rules; for $n = 4$, 3 rules; for $n = 5$, 7 rules; for $n = 6$, 21 rules; for $n = 7$, 135 rules; for $n = 8$, 2470 rules and for $n = 9$, 172958 rules.

Berend and Sapir [13] have
found the optimal probabilities
for the family of restricted ma-
jority rules, expert and major-
ity rules (extreme rules) as spe-
cial cases, and the family of bal-
anced expert rules. Then, the
two families are compared, the
rules within each family, and all
rules of the two families with the
extreme rules (Figure 1). In each
family, the rule is determined by
the number of group members
(x axis) having an influence on
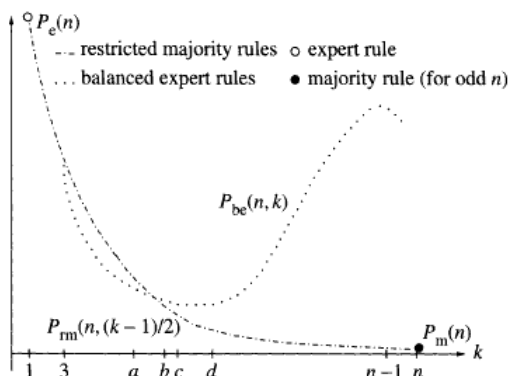the group decision.

Fig. 1: Comparing the two families

However, under the restricted majority rules, each of these members is equally
influential, the balanced expert rule gives the top member almost all the power,
and the top member is outvoted only if is opposed by all other influential mem-
bers. In [13], it is assumed that logarithmic expertise levels β_i are independent
exponentially distributed random variables. Figure 1 provides a schematic draw-
ing of the graphs of the functions $P_{be}(n,k)$ and $P_{rm}(n, \frac{k-1}{2})$ as functions of k
(number of group members) for large n. Here, P_{rm} and P_{be} are the probabilities
of the restricted simple majority rule and the balanced expert rule being opti-
mal, respectively. In particular, the probability of the expert rule being optimal
is higher than that of all other members of the family of restricted majority rules,
while that of the majority rule is smaller than that of all others. Moreover, for
$n \geq 3$, $P_{be}(n,k) < P_e(n)$; and for sufficiently large n, $P_{be}(n,k) > P_m(n)$.

In [10] it was also shown that for three-member groups, under complete
ignorance (an extreme case of partial information) regarding individual compe-
tences, simple majority rule is always preferable to the expert rule (equivalent
to an even-chance lottery on individual skills).

4   Intention reconsideration like dichotomous voting

There are various mechanisms that an agent might use to decide when to recon-
sider its intentions. We shall call such a mechanism a *reconsideration strategy*.
Besides, the reconsideration of intentions can be seen as a decision-making be-
tween to deliberate and to act, consisting of several criteria when choosing one
of the alternatives (i.e., a multi-criteria decision problem).

Multi-criteria decision problems and collective decision making are similar. In
both cases multiple preference orders on the alternatives exist and it is necessary
to combine them into a unique global preference ordering. Therefore, techniques
from the field of collective decision making, as voting, can be used to make
multi-criteria decisions.

In this section, we propose to implement the agent's reconsideration strategy
as dichotomous voting (uncertain dichotomous choice model). Given that we are

interested in answering the question of whether to reconsider current intentions, the model must choose between alternatives $a = yes$ and $b = no$. The correct alternative is the one that makes optimal $reconsider(\cdot, \cdot)$ function. Therefore, the two states of nature correspond to the possible outputs of the deliberation function. That is, we will denote by 1 the state of nature in which deliberate process does modify its current intentions, and by -1 the state of nature in which deliberate process does not change intentions. The common preferences of individuals are presented in Table 4.

	yes (a)	no (b)
changed intentions (1)	1	0
not changed intentions (-1)	0	1

Table 4: Payoff matrix for Intention Reconsideration

In this case, the net benefit from a correct decision is independent of the state of nature (i.e., $B(1) = B(-1)$). That is, there is symmetry in the individual utility in case of making a correct decision in the two possible states of nature, hence $\delta = 0$.

Let α denote the a-priori probability that $reconsider(\cdot, \cdot)$ chooses to reconsider intentions and deliberation process does change intentions (situation 4 of Table 1). Conversely, $1 - \alpha$ denotes the a-priori probability that $reconsider(\cdot, \cdot)$ function does not choose to deliberate, but if it did, it would not have changed its intentions anyway (situation 1 of Table 1). Note that the label of an alternative, a or b, does not convey any indication as to which of the two is the correct one. Consequently, alternatives might be considered a priori as equi-probable, i.e., $\alpha = 1/2$ and therefore $\gamma = 0$. In this way, under this assumption we get that $\gamma + \delta = 0$, so alternatives are perfectly symmetric, and function 4 is the optimal rule to be applied for IR.

Our work applies multi-criteria decision making, so we assume that n individuals correspond to the criteria $C = \{c_1, \cdots, c_n\}$, considered in decision making. These criteria should include dynamism, determinism and accessibility of the environment in which the agent is located, together with features of the current intention. Each criterion has a threshold θ_i from which it votes "yes" or "no" about reconsider intentions. Namely, if value of criteria v_i is higher or equal than θ_i then it votes "yes", otherwise it votes "no". Each p_i represents the probability of making the correct decision considering only the criterion c_i. This probability can be estimated or considered partially known, so that it could be applied to the rule most likely to be the optimal for any distribution (see Sect. 3.2). In the particular case that $|C| = 3$, we can assume complete ignorance about the vector p and use rule 5.

5 Conclusion

In this work, we have presented an intention reconsideration strategy implemented as an uncertain dichotomous choice model. This model allows us to identify "the correct alternative", i.e., the one that makes optimal reconsidera-

tion of intentions. In contrast to fixed strategies [3], our approach changes commitments to intentions depending on how the environment changes, and besides environmental dynamism, other features such as determinism and accessibility, are contemplated.

In [14], the performance of different reconsideration strategies are empirically evaluated in complex environment with varying degrees of dynamism, accessibility and non-determinism. Therefore, as future work, we intend to perform the same analysis with the reconsideration strategy proposed in this work, to assess its effectiveness.

Future work will also include research on the comparison of this model and those utilized in [15] as intention reconsideration strategy within a BDI agent. Moreover, we will study the optimality of the proposed reconsideration strategy following the guidelines from [9].

References

[1] Wooldridge, M.: Reasoning about Rational Agents. The MIT Press (2000)
[2] Schut, M., Wooldridge, M.: Intention reconsideration in complex environments. In: 4th International Conference on Autonomous Agents. (2000)
[3] Kinny, D.N.: Commitment and effectiveness of situated agents. In: 12th International Joint Conference on Artificial Intelligence (IJCAI). (1991) 82–88
[4] Russell, S., Wefald, E., Karnaugh, M., Karp, R., Mcallester, D., Subramanian, D., Wellman, M.: Principles of metareasoning. Artificial Intelligence (1991)
[5] Sosa-Toranzo, C., Errecalde, M., Ferretti, E.: On the use of agreement technologies for multi-criteria decision making within a BDI agent. In: 14th Ibero-American Conference on Artificial Intelligence (IBERAMIA). (November 2014) Accepted as full paper.
[6] Nitzan, S., Paroush, J.: Optimal decision rules in uncertain dichotomous choice situations. International Economic Review **23**(2) (1982) 289–97
[7] Nitzan, S.: Collective Preference and Choice. Cambridge University Press (2009)
[8] Bratman, M., Israel, D., Pollack, M.: Plans and resource bounded reasoning. Computational Intelligence **4**(4) (1988) 349–355
[9] Wooldridge, M., Parsons, S.: Intention reconsideration reconsidered. In: Intelligent Agents V: Agents Theories, Architectures, and Languages. Volume 1555 of LNCS. Springer (1999) 63–79
[10] Nitzan, S., Paroush, J.: Partial information on decisional competences and the desirability of the expert rule in uncertain dichotomous choice situations. Theory and Decision **17**(3) (1994) 275–286
[11] Sapir, L.: The optimality of the expert and majority rules under exponentially distributed competence. Theory and Decision **45**(1) (1998) 19–36
[12] Berend, D., Sapir, L.: Range of asymptotic behaviour of the optimality probability of the expert and majority rules. Journal of Applied Probability **43**(1) (2006)
[13] Berend, D., Sapir, L.: Between the expert and majority rules. Advances in Applied Probability **35**(4) (12 2003) 941–960
[14] Schut, M., Wooldridge, M.: Intention reconsideration in complex environments. In: 4th International Conference on Autonomous Agents, ACM Press (2000)
[15] Schut, M., Wooldridge, M., Parsons, S.: The theory and practice of intention reconsideration. Journal of Experimental & Theoretical Artificial Intelligence **16**(4) (2004) 261–293