

Improving Choice Mechanisms within the BVG Architecture

Luis Antunes, João Faria and Helder Coelho

Faculdade de Ciências, Universidade de Lisboa

Campo Grande, 1749-016 Lisboa, Portugal

Ph: +351-21-7500087

Fax: +351-21-7500084

{xarax@di., jfaria@, hcoelho@di.}fc.ul.pt

Abstract. The BVG agent architecture relies on the use of values (multiple dimensions against which to evaluate a situation) to perform choice among a set of candidate goals. Choice is accomplished by using a calculus to collapse the several dimensions into a function that serialises candidates. In our previous experiments, we have faced decision problems only with perfect and complete information. In this paper we propose new experiments, where the agents will have to decide in the absence of all the needed and relevant information. In the BVG model, agents adjust their scale of values by feeding back evaluative information about the consequences of their decisions. We use the exact same measures to analyse the results of the experiments, thus providing a fair trial to the agents: they are judged with the same rules they can use for decision. Our method, based on values, is a novel approach for choice and an alternative to classical utilitarian theories.

1. Introduction

We consider a setting in which multiple agents interact in a shared environment. Usually, this environment is computer-simulated. Sometimes it is self-contained and agents are used in experiments to draw conclusions about socially relevant phenomena; in other cases, there is a user to whom the agent responds to, and a certain amount of subservience is expected from the agent.

Whichever the complexity of agents, they must possess a flexible decision component to clarify all the deliberation processes. Even a compile-time pre-specified agent will be of little use if it is not ready for a certain extent of non-forecast possibilities. As the environment gets more demanding in terms of unpredictability (at least a priori unpredictability), more complex should our agent be in what respects to decision adaptability. The designer must have the means to specify what is expected from the agent even in a new environment s/he has never considered. With the advent of mobile computation and huge, varied artificial environments (such as the Internet), we have to enhance our agents with autonomous and reliable decision skills.

When confronted with a situation, an agent is defined as rational if he decides in such a way that pursues his self-interest. A classical way of defining self-interest is by adopting utility theory [11], that requires the agent to know in advance all possible situations and be prepared to express his preference between any two states of the world. Not only do these conditions seem difficult to be fulfilled, but also this theory leads to interesting decision paradoxes that show its limitations [9, 8].

An attempt to escape from this kind of rationality was the BDI (Belief, Desire, Intention) agent model [10]. Here, commitment to past decisions is used as a way to decrease complexity, since committed intentions constrain the possibilities for the future, and are only abandoned when fulfilled or believed impossible to fulfil. The preferences of the agents are represented by their desires, and these will be transformed in intentions through a deliberation process. In [13], nothing is said about how this deliberation is done. Most BDI approaches consider the problem of commitment, but choice is seldom addressed.

Simon [12] proposed the idea of aspiration levels along multiple, non comparable dimensions that characterise a decision problem. Aspirations are the minimum standards that some solution must meet in order to be adopted. The agent adopts and selects for execution the first solution that meets all of the aspiration levels. In a similar line of reasoning, we have addressed the issue of choice, as one of the central components in the agent's decision machinery [1, 2, 3]. We have proposed the use of multiple values to assess a decision situation. A value is a dimension against which a situation can be evaluated. By dimension we mean a non empty set endowed with an order relation. Most interesting situations from the decision standpoint will have several such dimensions, and so most decisions are based on multiple evaluations of the situation and alternative courses of action. The agent's choice machinery becomes more clear, as agents express their preferences through the use of this multiple value framework. Choice is performed by collapsing the various assessments into a choice function, that cannot be considered equivalent to a utility function, since it is computed in execution time. The multiple values framework we defend can encompass Simon's aspiration levels, but it is more general, allowing for further flexibility, as is shown in [2].

The coexistence of these values in a mind further allows the enhancement of the adaptability decision capabilities by feeding back assessments of the quality of the previous decision into the agent's decision process. Our agents' decisions no longer depend solely on the past events as known at design time. Instead, events are incorporated into the decision machinery as time passes, and the components of those processes evolve continuously to be aggregated just when a decision is needed. This is done by feeding back evaluative information about the results of the decision taken by the agent. In [2] this assessment of (the results of) the decision was done by using some measure of goodness (an abstract higher value). We also suggested other alternatives, such as the agent's own values, or the designer's values (which amounts to looking for emergent features in the agent's behaviour, that is, agents decide by using some system of values, but the designer is interested in what happens to *another* set of values, to which the agents do not have access). Even in the simplest version, the power of adaptability shown by this schema surpasses by far that of choice based

on the maximisation of expected utility. It is enough to remember the demands made on the agents by utility theory: they must know *in advance* all available alternatives and preferences between any two of them [7, 11].

In this paper we expand on the functionality of the BVG (Beliefs, Values, Goals) architecture to enhance the decision mechanism to cope with more complex situations, namely the agent has to decide in absence of all the relevant evaluations. We also propose functions for the several kinds of update of the decision process, and experimentally assess them. In section 2 we briefly present the BVG agent architecture, while in section 3 agent autonomy is address with a special focus on choice. Section 4 presents a case study that is explored through sections 5 and 6, where the choice calculus is presented. Section 7 concludes by pointing out the most important contributions.

2. The BVG architecture

The BVG architecture roughly follows Castelfranchi’s principles for autonomy contained in his “Double Filter Architecture” [6]. We include herein the main ideas behind this architecture, for a more exhaustive description refer to [1, 2, and 3].

The reference schema of the BVG architecture for decision-making includes goals, candidate actions to be chosen from, beliefs about states of the world, and values about several things, including desirability of those states. Values are dimensions along which situations are evaluated, and appropriate actions selected.

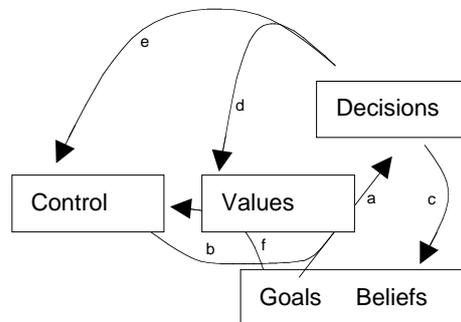


Fig. 1: The BVG architecture. (a) value-informed decision making; (b) control over decision; (c) watching consequences of decisions; (d) feeding back information to enhance decision machinery; (e) getting emotional information about the results of decisions; (f) cognitive evaluations.

Decision is a complex, multi-staged process in an agent’s mind. One stage deals with the origin of goals. Agents can either adopt goals from other agents or generate goals as a result of internal processes. In another stage, goals are considered against other objects in the agent’s mind, such as beliefs (which include plans about their feasibility) and classified accordingly. For instance, we can have suspended goals, active goals, etc. Finally, among the active goals, the agent has to serialise them into execution. This is the choice phase, our primary focus in this paper. We don’t want to

over-determine the behaviour of the agent. However, it is fundamental to reduce its deliberation time, which is possible by having several partial ordered sets of actions, and computing the final order only when the moment arrives.

In the model proposed in [3], emotions have a determinant role in the control of the choice process (see fig. 1). One main role of emotions is to set the amount of time available for a decision to be made. We proposed a cumulative method, that improves the quality of the decision when time allows it. The idea is based on the different importance of the relevant values. Options are evaluated against the most important values first.

The other key issue in the BVG architecture is the update of the choice machinery based on assessments of the consequences of the previous decisions. These assessments can be made in terms of (1) some measure of goodness (the quality of the decision, measured through its consequences on the world); (2) the same dimensions that the agent used for decision; and (3) a different set of dimensions, usually the dimensions that the designer is interested in observing, what amounts to look for emergent behaviour (cf. [2]).

3. Autonomy and character: focus on choice

When we want to define autonomous agents, we must address three processes: creation, selection and choice, which raises the issue of the agent's character. The character can be defined as the set of collective qualities, especially mental and moral, that distinguish a person or entity. Autonomy implies difference, the agents should be free to decide differently from one another even in the same situation: agents can have different characters. Either in the creation or in the choice, the reasons can vary from agent to agent. It is in this two phases that we can locate the personality traits that define the character. In the creation phase, some agents will adopt goals whereas others don't, some agents will generate goals in one way, while others may not generate goals at all, and only obey goals from their creator. In the choice phase, we look at different rationalities in different agents. Some choice is not necessarily irrational just because we cannot identify the reasons that led to it [7]. Even considering agents endowed with the same set of goals and beliefs, their different sets of values should be enough to produce differentiated behaviours.

We sustain that the multiple values approach represents a step towards autonomy. Here we will be more concerned with executive autonomy [6], in [3] we further addressed "cognitive mediation" issues, by proposing a mechanism for value-based goal acquisition, thus assuring goals autonomy. Autonomy should pay no tribute to a super-imposed rationality: autonomy means exactly the liberty to adopt one's own rationality.

In [2], our agent had a goal to fulfil that was characterised by some values. The candidate actions were all comparable according to the same values, so choice was straightforward. The agent just computed a real function of those values for all the alternatives, and the highest scorer was chosen. Afterwards, the agent looked at the results of his action by assessing the resulting state of the world against a dimension of goodness, and updated the values appropriately.

We now propose to tackle choice situations where the agent doesn't have all the relevant information. Imagine some goal G is characterised by targets for three (out of five) values: $V_1=\omega_1$, $V_2=\omega_2$, $V_3=\omega_3$. Let the candidate actions be represented by sub-goals G_1 and G_2 , with associated values, respectively: $V_1=\omega_{11}$, $V_4=\omega_{14}$, and $V_1=\omega_{21}$, $V_2=\omega_{22}$, $V_5=\omega_{25}$.

As long as choice is performed by using a linear combination of some function of the paired (e.g. ω_1, ω_{11}) values, like in [2], one can just omit the values outside of the intersection of the goal and the candidate characterisation, thus rendering them not redundant. Or we can assign mean values to the unknown fields. If other types of choice functions are used one must proceed with more care. Anyway, even in the simple case above, there are open problems to be dealt with. First of all, it is necessary to characterise the new adopted goal, say G_2 . Should G_2 include values for V_3 ? Should it keep values for V_5 ? We think that the answer to both questions is positive: V_3 should be kept (with target ω_3) because we are adopting G_2 *just because* of G . So it is only fair that we keep whatever guides the attempts at achieving G , and those are the values V_1 , V_2 , and V_3 . For an analogous reason we should include V_5 . It could be the case that V_5 represents important evaluative notions to be considered during the attempts at G_2 , and so we mustn't give up V_5 for our future execution. In both cases, these values will help control the execution of the agent towards his goals, possibly allowing for revisions and recalculations if the chosen goals no longer serve the relevant values at stake.

Another key and general purpose issue is the necessity of multidimensional assessment of the quality of decisions. There is only limited adaptation one can perform over his decision processes when considering unidimensional measures. This does not mean optimisation of the decision model to better react to some reward function. The truth is that we assume no reward function to be learnt, unlike in reinforcement learning. The world is just out there and our agents need to constantly adapt. Multidimensional information is shown to enrich this adaptation process.

4. A case-study: choosing wines

We put our model to trial by using an experimental setting where choice has to be performed in purchasing a product, with later assessment of this decision. We simulate the case of wine selection (in the future we will have our agents going to virtual shops and selecting wines for us, yet hopefully not drinking them). Note that this example, because of the generality of its multidimensional form, is quite meaningful for tackling other problem domains.

4.1 An instance of BVG to select wines

Some consumer wants to purchase a bottle of wine, and so goes to a wine merchant. He has some knowledge about wine, but naturally does not know every wine that is available to him. He will evaluate candidate wines using a fixed set of five dimensions: V_1 is the knowledge and trust he has of the producer; V_2 is the quality of the region in the given vintage; V_3 is the price of the wine; V_4 is the ideal consumption time window (with $v_4 = (v_{4,1}, v_{4,2}) \in V_4$); and V_5 is the quality of the

specific wine. Domains for these dimensions will be $D_1 = D_2 = \{0, *, **, ***, ****, *****\}$, $D_3 = N$, $D_4 = N \times N$, and $D_5 = [0..20]$.

Now imagine our buyer has the goal of giving a formal dinner this evening, and wants to impress his guests with a nice bottle of wine. So, he specifies as minimum standards (4, 4, 1000, (0,5), 16), meaning that the wine should come from a known and trustworthy producer, price is no matter, should be ready to drink, and a quality of at least 16 points. Notice that different things could be tried: for instance, his guests might be impressed by a wine fulfilling this specification: (1, 5, 75, (0,5), 18), meaning that a very expensive wine from an unknown producer turns out to be very good. In any case, there could be lacking information at the time the choice must be performed.

4.2 Goals and alternatives

We propose to iterate the process just illustrated, and let the consumer adjust his choice machinery. In a first row of experiments, we will test the power of quantitative assessments by making this adjustment by feeding back the grade the wine achieves when it eventually gets tasted. Afterwards, we enhance the choice machinery update by feeding back multiple assessments of the quality of the decision (for instance, after several bad bottles from a given region in a given vintage, we can decide to downgrade the impression we had on that region/vintage, or even to pay less importance to this dimension in the future).

Goals to buy are those characterised as follows:

G₁: Goal ($V_1 = ****$, $V_2 = ****$, $V_3 = 20$, $V_4 = (5,10)$, $V_5 = 16$) Buy(wine)

(good wine, from top producer in a good vintage, to hold)

G₂: Goal ($V_1 = **$, $V_2 = *****$, $V_3 = 10$, $V_4 = (1,4)$, $V_5 = 14$) Buy(wine)

(medium wine, from new producer in a very good vintage, to drink)

G₃: Goal ($V_1 = *****$, $V_2 = *****$, $V_3 = 100$, $V_4 = (20,30)$, $V_5 = 19$) Buy(wine)

(excellent wine, from star producer in an excellent vintage, to hold a long time).

Goals to drink include these characterisations:

G₄: Goal ($V_1 = *$, $V_2 = ****$, $V_3 = 0$, $V_4 = (0,1)$, $V_5 = 14$) Drink(wine)

(everyday drinking)

G₅: Goal ($V_1 = ****$, $V_2 = *****$, $V_3 = 200$, $V_4 = (0,10)$, $V_5 = 17$) Drink(wine)

(special occasion)

G₆: Goal ($V_1 = *$, $V_2 = *$, $V_3 = 200$, $V_4 = (10,50)$, $V_5 = 10$) Drink(wine)

(wine tasting)

We have included about thirty wines in our wine shop, and will look first at what our agent buys, to see how he builds his initial wine cellar. This allows us to tune up the choice function by hand. When he starts drinking, the process enters a new phase, with updates to the choice mechanism.

5. A calculus for choice

Let us start by presenting the functions we have selected to tackle the decision problem we have just described, before we go into their makings with some detail.

The function that performs choice is (for v_k features of the goal, and v_{ki} features of alternative i ; and where $\|(x,y)-(a,b)\| = \sqrt{(x-a)^2+(y-b)^2}$, (the Euclidean norm), and

$$\tau(u) = \frac{1+u}{2}:$$

$$F_{\text{buy}}(v_k, v_{ki}) =$$

$$c_1 \tau\left(\frac{v_1 - v_{1i}}{2 + |v_1 - 3|}\right) + c_2 \tau\left(\frac{v_2 - v_{2i}}{2 + |v_2 - 3|}\right) + c_3 \frac{1}{1 + e^{-\left(\frac{v_{3i} - v_3}{v_3}\right)^3}} + c_4 \left(1 - e^{-\frac{\|v_{4i} - v_4\|}{10}}\right) + c_5 \frac{1}{1 + e^{-\frac{v_{5i} - v_5}{2}}}$$

$F_{\text{drink}}(v_k, v_{ki}) = F_{\text{buy}}(v_k, v_{ki})$ with c_4 altered to have more importance in the final decision (for instance, c_4 twice what it was above).

5.1 Designing the choice function

Let us look into the design of the choice function F with some detail. For simplicity reasons, we opted to keep the structure of this function a linear combination of several components. F is meant to be minimised, so is each of its components (call them F_k). All F_k s will be scaled to the interval $[0, 1]$ by appropriate linear transformations (e.g.

from $[-1, 1]$ to $[0, 1]$ we use $\tau(u) = \frac{1+u}{2}$). It seemed natural to keep F also inside $[0, 1]$.

So we kept the linear combination convex ($\sum c_k = 1$). Several choices were possible, and we are conscious that the choice we've made ($c_k = 1/n$) corresponds only to a possible attitude of the buyer. We can easily model other attitudes by changing these coefficients (see discussion in section 6).

F_1 and F_2 share the same structure. In both, we compare our goal with the alternative, and then bring this difference to a scale of $[-1, 1]$ by dividing by a quantity that relates the scale of the difference with the scale of the goal. For instance, in F_1 , it is preferable to obtain 5 when aspiring to 3 (with the best of scores, -1), than to obtain 3 when aspiring to 1 (with score -0.5).

For field F_3 (rating price discrepancies), and after a lot of discussion and experimentation, we opted for a sigmoid curve, $f(x) = \frac{1}{1 + e^{-x^3}}$, which has the advantage

of being approximately stationary behaviour in the neighbourhood of the origin. Its possibly not ideal symmetrical behaviour to the left and to the right of the origin could be corrected by using another branch for the function. However, we didn't think it was necessary, because the numerical significance is negligible. It suffices to notice

that f will be applied to $\frac{v_{3i} - v_3}{v_3}$, and price v_{3i} is always greater or equal to 0, so when the argument is negative, it always is in the interval $[-1, 0]$. With this design of F_3 , we slightly overlook the behaviour in face of discounts. However, the behaviour in face

of high prices is fully modelled (the case when $x \in [0, +\infty[$). When the price is exactly what we look for ($x=0$), f is $1/2$, that is, the medium point of the target interval. For small aggravations of the price, f grows slowly; f grows faster for prices far superior to the desired one, but refrains this growth rate when prices are very high, meaning that when the price is already much higher than the one we want, further increases only aggravate the choice function infinitesimally.

F_5 rates quality discrepancies and is similar to F_3 , although it doesn't suffer the cubic aggravation. The argument is only divided by a constant (2), in order to get a less steep curve. For example, for a goal of 15 points, F_5 will compute 0.182 for getting 18, 0.269 for getting 17, 0.378 for getting 16, 0.5 for getting 15 and 0.818 for getting 12.

Finally, field F_4 is only a reduction to the interval $[0, 1]$ of the Euclidean distance between the two time-intervals (goal and alternative). The distance is divided by a factor 10 (supposed to be a common value for the maximum norm) as a way of avoiding too much concentration of the function around 1.

Some of the values v_{ki} can be directly observed on the alternatives (e.g. price). Others depend from prior knowledge, and are assigned to each alternative by the consumer on the basis of previous contacts with the same product. We thought it was adequate to value the number of previous contacts, and so consider the pair (x, n) to represent respectively the rating for the value at stake, and the number of observations that gave rise to that rating. In the next section we will check on the scope of this representation.

5.2 The update mechanism

We have several alternatives to update the choice mechanism. First, we consider the result $r(s_i)$ of the selected alternative (s_i), to be an unidimensional classification of the tasted wine, in the scale of 0 to 20.

The update of the choice mechanism is thus done by function G , as follows:

$$G(r(s_i), v_{ki}) = [G_1(r(s_i), v_{ki}), G_2(r(s_i), v_{ki}), \text{id}, \text{id}, G_5(r(s_i), v_{ki})]^T,$$

where id is the identity function projected onto the appropriate component, and

$$G_1(r(s_i), v_{1i}) = \frac{\text{scale}(r(s_i)) + \text{no. observations} \times \text{rating}}{1 + \text{no. observations}}, \text{ with } v_{1i} = \frac{\text{rating}}{\text{no. observations}},$$

$$G_2(r(s_i), v_{2i}) = \frac{\text{scale}(r(s_i)) + \text{no. observations} \times \text{rating}}{1 + \text{no. observations}}, \text{ with } v_{2i} = \frac{\text{rating}}{\text{no. observations}},$$

$$G_5(r(s_i), v_{5i}) = \frac{r(s_i) + \text{no. observations} \times \text{grade}}{1 + \text{no. observations}}, \text{ with } v_{5i} = \frac{\text{grade}}{\text{no. observations}},$$

where scale performs the adequate transformation from $[0..20]$ to $[1..5]$.

We should note that some of these updates (namely G_1 and G_2) must be applied to all alternatives containing the same information (e.g. one must update the producer rating not only for the wine just tasted, but also in all other wines of the same producer). Note also that we chose to perform a complete average of all the

observations, although other options could be used (e.g., moving averages, or filtering through a degradation function, like in [4]).

Another alternative to perform this update operation would be to distribute credit equally among all the components of choice. This amounts to linearly transform the result $r(s_i)$ into the interval $[0, 1]$ (the interval where all F_k s have co-domains), and then incorporate this classification into the alternative, for instance by computing an average like we have done above, for all of the multidimensional components (or using some moving average to include memory dissipation effects). However, it is easily shown that this would amount to collapse all dimensions into one unique dimension, because the iteration of such an update process would quickly lead to equal values in every dimension (the weight of the initial values would be weaker and weaker with the addition of new values to the average).

One of the main theses of our work rests upon the use of multiple values to perform also the update of the choice machinery. This was only pointed out in [3] so here we propose another update function H to compare results of the two simulations. In the following definition of H , $r(s_i)$ will then be a vector of multiple assessments. We consider the same dimensions that were used for choice itself (cf. also [2]), although some things are not assessable, price for instance. So, we let $r(s_i)$ be $[r_k(s_i)]$, but consider only $k=4$ and $k=5$:

$$H([r_k(s_i)], v_{ki})=[H_1(r_k(s_i), v_{ki}), H_2(r_k(s_i), v_{ki}), \text{id}, H_4(r_k(s_i), v_{ki}), H_5(r_k(s_i), v_{ki}))]^T;$$

$$H_1(r_4(s_i), r_5(s_i), \frac{\text{rating}}{\text{no. obs.}}, v_{2i}, v_{3i}, v_{4i}, v_{5i})=\frac{(1+||r_4(s_i)-v_{4i}||)\times\text{scale}(r_5(s_i))+\text{no. obs.}\times\text{rating}}{1+||r_4(s_i)-v_{4i}||+\text{no. obs.}};$$

$$H_4(r_4(s_i), r_5(s_i), v_{1i}, v_{2i}, v_{3i}, v_{4i}, v_{5i})=\frac{r_4(s_i)+v_{4i}}{2}.$$

H_2 and H_5 are respectively obtained from G_2 and G_5 , in such a way as H_1 was from G_1 . Note that for V_4 we perform an outright average with previous value, because the most recent assessment reports on the current condition of the wine, and so should have a quicker impact on our decisions. The factor $(1+||r_4(s_i)-v_{4i}||)$ put on the quality is meant to represent the urgency of change. The more the wine is not to be consumed when we imagined it should be, the more weight the new assessment should have in the update.

5.3 Further enhancing choice procedures

As we describe the genesis of the choice function F , we should observe that values V_1 to V_5 are considered a priori, in design-time. Much as we want our choice model to tackle any situation in execution-time, it is to be expected that situations come around where our model doesn't cope. The BVG architecture encompasses methods for adaptation to enhance choice performance. Adaptation is done by updating the values along the relevant dimensions, but another possibility would have been to modify the choice function.

This option would amount to incorporate some realism into the model, in the sense that our agent would always consider observations of the alternatives at stake to be perfect, and would change the choice function in order to optimise its behaviour. We prefer the alternative idealist view: the choice function is not perfect, rather as good

as we can get; it is always possible to improve on the perception we have of alternatives, illuminated by the assessment of our previous choices. A combination of both alternatives would be preferable, but the one we picked seems more intuitive, and allows for a more parsimonious calculus.

Once chosen a model of choice and respective values, we can raise the point if those are the appropriate values, all of the appropriate values, what are the correct functions to account for the role of each of those values in the problem at hand, etc. Can our model ever be closed up and considered ready for every eventuality? We think not. The world is highly dynamic and unpredictable, and choice must conform to that. Decision is often time-dependent, and an additive form for the choice function allows for the process to be interrupted (or even pre-emptive) and some anytime character to be achieved (cf. [3]). On the other hand, mechanisms must be included to avoid ties between alternatives (for instance, in the case of major absence of evaluative information). Some choice must be done, whatsoever.

Usually, when people are confronted with massive lack of information and have to choose all the same, they recur to analogies with other lived situations. For instance, to pick some wine just because the bottle has a nice label is not irrational, it amounts to invoke the analogy upon which the whole Marketing discipline is based: if the producer cares about the quality of the packaging then *for sure* he cares about the quality of his product. Could it be then that we should always include the beauty of the label as a pointer of the wine's quality? Perhaps, but when information about the real quality is available why should we trust the label?

The conclusion is: when we don't have information, we can use the *Principle of Insufficient Reason* [5] to assign a mean value to this field. When we have information, whatever its origin, we just use it. Although we cannot at this time detail all of the mechanisms here involved, in the next section we present one such mechanism. We include a new field in the choice function that uses non-linearity to account for positive or negative effects of lacking information.

5.4 The effect of surprise

After we analysed values V_1 to V_5 and corresponding fields in the function F (let's call them F_1 to F_5), where common-sense and intuition set the optimisation rules, it was apparent to us that we needed to introduce a further dimension, V_6 , to account for features not grabbed by the initial fields. Broadly speaking, this new value accounts for the "surprise-effect." Let us explain through an example.

An unknown brand of automobiles issues a new model with a price far above the average for that market segment. A possible buyer's attention is drawn by the publicity effort and is surprised by the price. Probably, he will not buy this car, since the amount at stake is high. If, on the contrary, the new model costs much less than its competitors, the effect on the image of the product can also lead the careful customer not to buy it. If we consider a box of chocolates instead of a car, this "surprise-effect" can more easily lead the customer to a purchase: the potential loss is much less. In this case, the expectations caused by the high price (and other factors, for instance, a good presentation) surpass the negative effects.

Going back to our case study, the same wine can rise different expectations (and decisions) in different agents. H and A both see on the shelves a wine from a new producer, with good presentation, and slightly higher price than the average for its region. Naturally, for some of the F_k s there is no information. Then H buys the wine and A doesn't. Clearly, our original F was disregarded by H. The surprise-effect of the high price caused different evaluations (of attraction and of repulse) in the two agents.

F was built in such a way (with all F_k s increasing) that this behaviour of H could never occur. But the F_k s must be increasing, at the risk of arbitrarily losing transitivity in choice (cf. section 5.1). We propose to add a new component to F , called F_6 , to account for this strange effect. F_6 is sensitive to lacking information in the other F_k s. When this happens, F_6 sets up the agent's preference for novelty, by conceding extra points for big discrepancies in whatever evaluative information is available. For instance, in H's case, F_6 will give a positive contribution to F , for the simple reason that v_{3i} is high, and v_{5i} is low (in the absence of information about quality, H naturally assigns an average grade to v_{5i}).

The default value of F_6 is 0.5. When the surprise effect is triggered by lacking information, we'll have:

$$F_6(v_k, v_{ki}) = 0 \text{ if } v_3/3 < v_{3i} - v_3 \leq v_3 \text{ and } v_1 \leq 2 \text{ and } v_2 \leq 2 \text{ and } v_5 \leq 10;$$

$$F_6(v_k, v_{ki}) = 0.25 \text{ if } v_3/3 < v_{3i} - v_3 \leq v_3 \text{ and } (v_1 \leq 2 \text{ or } v_2 \leq 2 \text{ or } v_5 \leq 10).$$

6. Results

The first results we collected concern the effect of surprise (field F_6). To this end, we decided to use the very example that led to the inclusion of this field, that is, H and A behave differently in face of a possible purchase of a bottle of red wine Quinta da Ponte Pedrinha, Dão, 1997. H buys the wine, and A doesn't. We evaluated this particular wine, assuming the real conditions of insufficient information that H and A had to stand against. So we considered default mean values for the lacking information, and computed the choice function for goal G_2 with and without the surprise-effect, in a universe of 25 other wines. Obviously, we didn't expect surprise to put it into pole-position, but it jumped 5 positions up, from 23rd to 18th in the ordered list. Further positions could have been gained if the particular parameters for either the wine or the goal would have made it possible to get the maximum possible score for surprise.

In table 1, we list the best five wines to be chosen for each of the goals G_1 to G_3 . We had previously made the choice ourselves, based on the text description of those goals (see section 4.2). The wines from our list that are present in the BVG choice are underlined in the table. Some interesting results occur, especially in what concerns goal G_3 . Since this goal requires wines of great longevity, and the contenders weren't particularly adapted to that requirement, four out of the five best wines benefited from insufficient information leading to the assignment of mean default values (0.5) to the respective field of the choice function (F_4). This unavailability of information also sustains the results of three of the five wines for G_1 . Castello Banfi manages to get in the top for all three goals as a result of having already a non-default quality

assessment of 15, whereas most candidates have a default quality of 10. Nevertheless, other candidates with good quality grades didn't make that well, since had poorer grades on other values. A good fit between goals and characteristics of the candidates also favours a good score (e.g. Quinta de Pancas has a perfect fit with goal G_2 with respect to consumption time).

Table 1. Top-scorers for buying goals.

G_1	G_2	G_3
Castello Banfi, Chianti '98	<u>Castello Banfi, Chianti '98</u>	Castello Banfi, Chianti '98
<u>Quinta do Côtto 1997</u>	Quinta das Caldas 1996	Quinta Baceladas 1997
Quinta das Caldas 1996	Herd. Esporão, Arag. 1997	<u>Tapada Coelheiros 1996</u>
Herd. Esporão, Arag. 1997	Quinta da Terrugem 1996	<u>Capela-Mor 1997</u>
Ad. Coop. Portalegre 1997	Qta. Pancas, Cab. S. 1997	<u>Quinta do Côtto 1997</u>

Table 2 shows the five top-scorers (for goal G_2) when we progressively changed weights. As we increase the weight of V_4 in the final decision, we watch the candidates with the most approximate fit to G_2 in what concerns V_4 do better and better. We did some other experiments with V_1 and V_2 and the results were similar. Setting up weights seems to be a good way of fine-tuning the choice model.

Table 2. Fine-tuning by changing weights.

(1/6, 1/6, 1/6, 1/6, 1/6, 1/6)	(.1,.1,.1,.5,.1,.1)	(.05,.05,.05,.75,.05,.05)
Castello Banfi, Chianti '98	Castello Banfi, Chianti '98	Qta. Pancas, Cab. S. 1997
Quinta das Caldas 1996	Qta. Pancas, Cab. S. 1997	Castello Banfi, Chianti '98
Herd. Esporão, Arag. '97	Quinta da Terrugem 1996	Quinta da Terrugem 1996
Quinta da Terrugem 1996	Herd. Esporão, Arag. '97	Falcoaria 1997
Qta. Pancas, Cab. S. 1997	Falcoaria 1997	Borges 1996

Now we invert terms, and try to induce better results for Quinta da Ponte Pedrinha, by exaggerating the weight of the surprise-effect in the calculus of F . We kept $c_1=c_2=c_3=c_4=c_5=(1-c_6)/5$, and ranged c_6 through $\{0.15, 1/6, 0.2, 0.25, 0.4, 0.5\}$, and the position of Ponte Pedrinha changed respectively from 20th to 18th, 17th, 17th, 5th and 2nd. So, one way of getting different attitudes with respect to surprise (and so emulate our example buyers H and A) is to consider a heterogeneous distribution of weights c_k . This conclusion is extendable to any other field of F (e.g. different importance of money for different agents, including some function F_3 (price) that depends on exogenous variables, for instance agent's wealth). Note that we didn't set any desirability for the state of surprise, although this is easily done. The idea would be to simulate the consumer's curiosity, by giving a premium to options not yet tried and assessed.

In a similar line of reasoning, there is another way of adapting the agent's choice machinery without changing the function F at all: by judiciously setting up the goal specifications. We isolated value V_3 (price), and run several experiments with goal G_2 . The first thing we realised is that setting up a very low aim for V_3 is pretty much the same as setting up a very high one, and amounts to making F_3 irrelevant. This is caused by the shape of function F_3 : both very big or very small discrepancies between

v_3 and each of the v_{3i} tend to cause F_3 to equalise all alternatives. So we watched the ordering of alternatives when small variations were made around the original aim price of G_2 . For a significant number of wines, the non-linear shape of F_3 determines trajectories from the V_3 -irrelevant initial ordering ($v_3=1$) to that same order (resulting from $v_3=1000$), passing through several progressive changes. Table 3 summarises some interesting cases, where price is quoted after the name of the wine.

Table 3. Some order trajectories determined by aim-price change.

v_3	1	8	9	10	11	12	15	20	∞
Poças, Res. 1992 (9.85)	13	13	14	14	14	16	16	16	13
Qta. Maias, Res. '91 (4.8)	16	16	16	16	15	14	13	13	16
Qta Terrugem '96 (11.22)	5	4	4	4	4	5	6	5	5
Qta. Pancas, C.S. '97 (6.7)	6	5	5	5	6	6	4	4	6
Ad. Portalegre '97 (14.9)	4	9	6	6	5	4	5	6	4

It is sometimes difficult to predict the results of a change in the goal values. So preferably, the set up of the decision functions will be made before-hand by correctly judging the weight of the various dimensions in the final decision. Moreover, some preliminary experimentation is recommended with goal specifications, before decisions are taken. The power of this decision model remains in its capacity for performing updates.

Table 4. Top-scorers for buying goals after G-update.

G_1	G_2	G_3
<u>Quinta do Côtto 1997</u>	Poças, Reserva 1992	Finca Flichman, C. S. '97
Vallado 1998	Castello Banfi, Chianti '98	<u>Quinta do Côtto 1997</u>
Herd. Esporão, Arag. '97	Herd. Esporão, Arag. '97	<u>#Capela-Mor 1997</u>
Quinta das Caldas 1996	<u>Casa de Santar, Res. 1996</u>	<u>#Quinta Banceladas '97</u>
Castello Banfi, Chianti '98	Quinta do Côtto 1997	<u>#Tapada Coelheiros '96</u>

Our last results are obtained by letting the choice machinery decide purchases and tastings. Their results are taken from real world tastings. After 28 tastings of 17 (of the 26) wines, and without further calibration, updates were performed as described in section 5.2. Table 4 shows the top five scorers after updates are made, and is meant to be compared with table 1.

The fall of Quinta das Banceladas (with respect to G_1) from position 14 to position 20, (while, say, Castello Banfi has fallen only slightly), is partially explained by the fact that its producer, Caves Aliança, was assessed through other wines. So despite Banceladas wasn't tasted, it was penalised all the same. It is interesting that the top scores for goal G_3 are kept by wines that didn't get tasted (marked with #). So the fact that some wines got tasted and even got good scores might not be enough to score highly with such an out of the way goal. Also the fact that updates aren't performed on V_4 allows for good scoring for F_4 on basis of unknown information. Since updates are progressive and cumulative, changes in ordering are gradual and depend on the quantity of samplings for each wine. This is the reason why the medium-scorer Castello Banfi decays from his top position rather slowly. A solution for quicker

adjustment could be to assign other type of default values (like the average of a set of somehow related wines, e.g. from the same producer, or the same region/year), or the use of moving averages to disregard older observations.

These comments lead us to the conclusion that there is a clear need for another type of update, based on a richer set of information, like we proposed above with update function H. Experiments with this function are still being conducted, but we present in table 5 some preliminary results. We have redone experiment reported in table 4, but now using update function H.

Table 5. Top-scorers for buying goals after H-update.

G_1	G_2	G_3
<u>Quinta do Côtto 1997</u>	Poças, Reserva 1992	<u>Quinta do Côtto 1997</u>
Herd. Esporão, Arag. '97	<u>Casa Santar, Res. 1996</u>	Herd. Esporão, Arag. '97
Vallado 1998	Herd. Esporão, Arag. '97	<u>#Capela-Mor 1997</u>
Quinta das Caldas 1996	<u>Castello Banfi, Chianti '98</u>	<u>#Qta. das Baceladas 1997</u>
Poças, Reserva 1992	Quinta do Côtto 1997	<u>#Tapada Coelhoiros 1996</u>

We compared this final ordering with our initial intuitive choices, to discover that the top five we had chosen for each goal had an average position of 15, 8, and 5, respectively for G_1 , G_2 , and G_3 . We explain the rather poor results for G_1 with the fact that most of the wines in our group weren't tasted at all, and so did not have the chance of improving position. Anyway, we don't find the results unsatisfactory, since it was obvious that this new way of performing update will lead to much faster adaptation to the user's tastes as accounted. If these accounts are accurate, the choice function will eventually reflect the user's preferences.

Perhaps the main problem we found whilst conducting these experiments was the evaluation of their results. Unless some drastic surprise occurred, this could be (and in fact was) easily foreseen; the question is 'how should one evaluate the results of his evaluations?' These experiments succeed in exhibiting the potential of the BVG model, but they don't go far beyond that. The idea of looking for emergence as a result of the agent using exogenous values to update his decisions (cf. section 2) may provide help in escaping from this riddle.

7. Concluding remarks

We have restated the fundamentals of the BVG architecture, namely the use of multiple values to inform choice, and the feedback of assessment information to recalibrate the choice machinery. We have proposed an instance of BVG that performs decisions in a case domain, and a corresponding calculus for choice. The update prescribed in BVG was done by feeding back both unidimensional and multidimensional information. The latter proved experimentally to accomplish further adaptability and reactivity to important new information coming from the environment.

The choice of the interval $[0, 1]$ for co-domain of F was not naïve. In the future, we will study a stochastic approach concerning the agent's decisions. We will also test

the thesis that patterns of consumption grow also out of habit (by letting the choice mechanism update itself), as opposed to 'rational' assessment of previous decisions. Further experimentation is also needed to assess emotion-driven control. We will be trying this approach in a multiple agent scenario, with incidence on the expansion of an agent's system of values as a result of interaction with other agents.

Acknowledgements. This research has been carried out within the research unit LabMAC, and partially supported by projects Praxis 2/2.1/TIT/1662/95 (SARA) and Praxis XXI/Feder/FCT. We wish to express our thanks to the anonymous referees.

References

1. Antunes, L., Towards a model for value-based motivated agents, in Proceedings of MASTA'97 (EPIA'97 Workshop on Multi-Agent Systems: Theory and Applications), Coimbra, October, 1997.
2. Antunes, L. and Coelho, H., Decisions based upon multiple values: the BVG agent architecture, in Barahona, P. and Alferes, J. (eds.), Progress in Artificial Intelligence, Proceedings of EPIA'99, Springer-Verlag, Lecture Notes on AI no. 1695, September, 1999.
3. Antunes, L. and Coelho, H., Redesigning the agents' decision machinery, in Proceedings of IWAI'99, Siena, October 1999; extended and revised version will appear in Paiva, A. (ed.) Affective Interactions, Springer-Verlag, Lecture Notes on AI no. 1814, 2000.
4. Antunes, L., Moniz, L. and Azevedo, C., RB+: The dynamic estimation of the opponent's strength, in Proceedings of the AISB Conference, IOS Press, Birmingham, 1993.
5. Bernoulli, J., *Ars Conjectandi*, (posthumous publication, Basel, 1713), reprint in *Culture et Civilisation*, Bruxelles 1986.
6. Castelfranchi, C., Guarantees for autonomy in cognitive agent architecture, in Wooldridge, J. and Jennings, R., *Intelligent Agents, Agent Theories, Architectures, and Languages*, ECAI'94 Workshop (ATAL), Springer-Verlag, Lecture Notes in AI no. 890, 1995.
7. Doyle, J., Rationality and its roles in reasoning, *Computational Intelligence*, 8 (2), 1992.
8. Faria, J., A Discussion with Kolmogorov; a Survey in Defence of the Axiomatic Approach towards Subjective Probability, Tech. Rep. 1/89, CEAUL, Lisboa, 1989.
9. Hollis, M., *The Philosophy of Social Science - An Introduction*. Cambridge: Cambridge University Press, 1994.
10. Rao, A. S. and Georgeff, M. P., Modeling Rational Agents within a BDI-Architecture, in Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning, Morgan Kaufmann, 1991.
11. Russell, S. and Norvig, P., *Artificial intelligence: a modern approach*, Prentice Hall, 1995.
12. Simon, H., *The Sciences of the Artificial* (3rd edition), the MIT Press, Cambridge, 1996.
13. Wooldridge, M. and Parsons, S., Intention Reconsideration Reconsidered, in Proceedings of ATAL'98, 4-7 July, Paris, 1998.