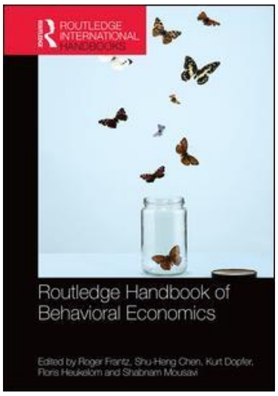

Routledge Handbook of Behavioral Economics

Roger Frantz, Shu-Heng Chen, Kurt Dopfer, Floris Heukelom, Shabnam Mousavi

Computational Behavioral Economics

Publication details
https://test.routledgehandbooks.com/doi/10.4324/9781315743479.ch21
Chen Shu-Heng, Ying-Fang Kao, Ragupathy Venkatachalam
Published online on: 27 Jul 2016

21

COMPUTATIONAL BEHAVIORAL ECONOMICS

Shu-Heng Chen, Ying-Fang Kao, and Ragupathy Venkatachalam

Introduction

Computational intelligence has been frequently applied to modeling artificial agents in agent-based computational economics. Commonly used applications include reinforcement learning (Chen, 2013), classifier systems (Vriend, 2002), genetic algorithms, genetic programming (Chen, 2002a,b), swarm intelligence (Boyer, Brorsen & Zhang, 2014), and instance-based learning (Pape & Kurtz, 2013). They are considered as alternative toolkits for the classical or Bayesian statistical models in modeling bounded-rationality and adaptive behavior (Sargent, 1993). However, these toolkits, except for reinforcement learning, are not explicitly grounded in psychology. It, therefore, remains to be seen whether these "machines" (artificial agents) are related to the bounded-rational agents as conceived by behavioral economists. Or, alternatively, to what extent can we relate the general principles or practices that are frequently applied in behavioral economics to the designs of these machines?

This issue has generally been ignored in the literature on behavioral economics, since machine learning and artificial intelligence remain a focus only for few branches of behavioral economics, specifically those following the legacy of Herbert Simon. On the other hand, this issue has not been well noticed in the literature on the machine learning community either. Although the machine learning community is well aware of the prevalence of ill-defined or poorly structured problems, this understanding is rarely extended to the context of economic decision making. Specifically, these two communities do not systematically share a background of the methodological controversy related to the divide between Homo Economicus and Homo Sapiens (Thaler, 2000). Therefore, given this dual ignorance, the fundamental connection between computational intelligence and behavioral economics is either missing or it only exists in an implicit manner.

The purpose of this chapter is to uncover this fundamental connection and to give it a systematic treatment. We attempt to do so by reviewing the behavioral economic principles behind computational intelligence tools. On the basis of this fundamental connection that we establish, we can see how agents, equipped with some "intelligence designs", substantiate the behavioral constraints and heuristics through implementable (computational) procedures. We refer to this substantiation or implementation and to the implied general approach as *computational behavioral economics*.

The rest of the chapter is organized as follows. The second section reviews some general features of decision making. This review motivates the framework used in this chapter.

The framework begins with routines, defaults or automated decisions. The third section addresses the role of computational intelligence in shaping this kind of decision process. This connection between computational intelligence and behavioral economics is illustrated by the instance-based decisions, such as K nearest neighbors, and other related algorithms, such as K-means, self-organizing maps and reinforcement learning. To cope with information or choice overload, heuristics based on instances need to be structured in a hierarchical form. The fourth section addresses how computational intelligence can be applied to examine this more advanced decision making behavior. The fifth section discusses the formation of novel heuristics, including the discovery of new attributes, new instances, and new hierarchies. The formation processes involve the idea of autonomous agents, whose behaviors are driven by the modularity heuristic. Computational modeling of these behaviors can be assisted by evolutionary computation, which provides an effective representation of behavioral heterogeneities among decision makers. Decision making can be affected by peers, colleagues, neighbors, and social norms. These behaviors have also been found in entomological experiments and some of them have been well formulated in computational intelligence. The sixth section provides a brief account of this development. The seventh section discusses some problems of treating randomization as a heuristic in decision making. Concluding remarks are presented in the final section.

Decision making and choices

Before we proceed, it may be useful to notice a common feature shared by both behavioral economists and machine learning scholars. For both, the "real world" is a world filled with ill-structured and vaguely defined problems. Many intelligent toolkits were proposed mainly to deal with these challenges. These challenges involve a kind of uncertainty, ambiguity or vagueness, which cannot be well formulated in a probabilistic environment and hence cannot be solved using standard rational (optimization) procedures that are built upon statistical decision theory or the von Neumann–Morgenstern expected utility maximization paradigm (von Neumann & Morgenstern, 1944). One of the most telling examples was given by Gerd Gigerenzer (Gigerenzer, 2007):

> A professor from Columbia University was struggling over whether to accept an offer from a rival university or to stay. His colleague took him aside and said, "Just maximize your expected utility—you always write about doing this." Exasperated, the professor responded, "Come on, this is serious."
>
> *(Ibid.: 3)*

A little reflection on this somewhat embarrassing situation highlights some important facets of decision making. First, many decisions are inconsequential, but some are not. Second, some choice or decision problems are encountered frequently; some less often. Accepting a new job offer or keeping the current job is not an inconsequential decision and is not the kind of decision which we make frequently; nevertheless, this kind of decision problem is prevalent in a normal economic life. Third, while it may be difficult to figure out the exact number of decisions that we make in a typical day, this number can be large and definitely larger than we might think (Wansink & Sobal, 2007). Fourth, we spend very little time making many choices or decisions and due to time constraints, many of us do not allow ourselves to spend too much time making those decisions (Mormann, Koch & Rangel, 2011). Fifth, many decisions are often made by processes that may be unclear for us, say, by emotion or gut feeling, or even automated (Damasio, 1994; Kahneman, 2011; Newell & Shanks, 2014). It is fortunate that many decisions do not take up much of our time or even need our conscious effort; therefore, we are still able to handle a sizable number of decisions in a typical day, including those with sizable consequences and for which we have very little past experience.

These facets of decision making problems suggest that there are two types of decision modes. The first are the *automated decision modes* that can handle frequently encountered decisions, specifically, those inconsequential ones. The second are the *manual decision modes* that can address less frequent, less experienced, but consequential decisions. The first type of decision mode typically refers to those *defaults* and *routines*, whereas the second type of decision mode is a meta-level decision model, which can identify novel elements, and constantly review and revise all routines and defaults, thereby facilitating the discovery of new routines or defaults.

Routine decision modes can be viewed as being organized in a hierarchical form (i.e., *the routine over routines*) as shown in Figure 21.1. The familiar decision problem will trigger our memory of the past similar situations, and the associated routines being followed in the past, but only the most relevant routine will be followed. The chosen routine will then be reviewed and revised based on its performance each time after its application, and will be added to the memory of routines. In this way, the set of routines can be updated, even occasionally. This hierarchy has often been mentioned in behavioral economics, but probably the most prominent quotation is the following one from Friedrich Hayek (Hayek, 1945).[1]

> We make constant use of formulas, symbols and rules whose meaning we do not understand and through the use of which we avail ourselves of the assistance of knowledge which individually we do not possess. We have developed these practices and institutions by building upon habits and institutions which have proved successful in their own sphere and which have in turn become the foundation of the civilization we have built up.
>
> *(Ibid.: 528)*

In the following sections, we elaborate more on this notion of hierarchical decision making processes, involving routines or rules, that are based on the experiences of the agents.

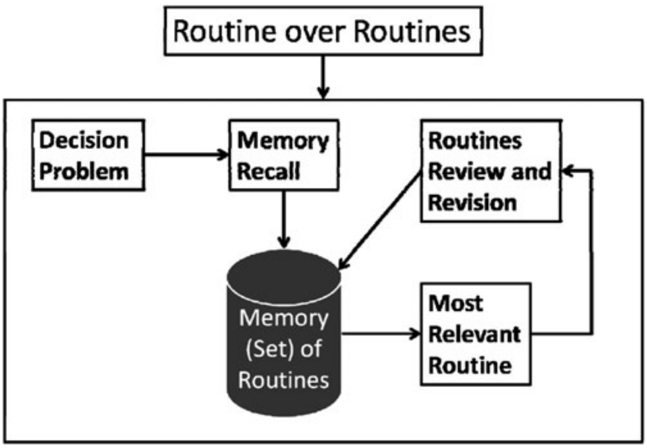

Figure 21.1 Routine formulation

Routines and instance-based decisions

The two-level hierarchical decision framework begins with the idea of defaults or routines, a subject well studied in behavioral economics (Thaler & Sunstein, 2008; Betsch & Haberstroh, 2014; Madrian, 2014). Routines help specify the rules concerning the default behavior for various problem instances. They allow us to economize on the time required for decision making and enhance the automated procedures for decision making. In this section, we shall address the behavioral features of using routines, and hence defaults, from the perspective of computational intelligence.

Routine formulation plays an important role in computational intelligence. The essence of the idea is that, until otherwise stated, similar simulations tend to evoke similar responses (decisions, actions, and choices). The key then is to consider an appropriate notion of similarity. David Hume, in his book *An Enquiry concerning Human Understanding*, has the following remark on experience and similarity.

> In reality all arguments from *experience* are founded on the *similarity* which we discover among natural objects, and by which we are induced to expect effects similar to those which we have found to follow from such objects. . . . *From causes which appear similar we expect similar effects.* This is the sum of all our experimental conclusions.
>
> *(Ibid.: section IV; italics added)*

Among many computational intelligence toolkits, an illustration concerning the first of the two modes (i.e., default or routine mode) that is most familiar to economists is the *case-based decision* (Gilboa & Schmeidler, 1995, 2001). In computational intelligence, the case-based decision is also popularly known as *instance-based learning* (Aha, Kibler, & Marc, 1991) or *lazy learning* (Aha, 1997).[2]

In instance-based decisions, the decision environment (instance) is characterized by its related features (attributes); for example, a vector $\boldsymbol{a}$ in an M-dimensional Euclidean space $\boldsymbol{R}^M$, $\boldsymbol{a} \in \boldsymbol{R}^M$. When the decision maker at time t faces a situation (instance) characterized by $\boldsymbol{a}_t$, we assume that she will recall her actions from her experience in similar situations in the past. Let $\boldsymbol{A}_t$ be the memory space of the past instances,

$$A_t = \{\boldsymbol{a}_s : S < t\}, \tag{1}$$

and $\boldsymbol{R}_t$ be the subset of similar instances, that is,

$$\boldsymbol{R}_t = \left\{\boldsymbol{a}_k : k < t, \left\|\boldsymbol{a}_k, \boldsymbol{a}_t\right\| = \epsilon_{k,t} < \epsilon\right\}, \tag{2}$$

where $\|\cdot\|$ is a metric which may be subjectively determined by the decision maker, and the distance ϵ, also subjectively determined, dictates what are perceived as similar instances by the decision maker. Furthermore, let d_t be the decision corresponding to an instance $\boldsymbol{a}_t$. The instance-based decision rule $d_t(\boldsymbol{a}_t)$ is then the function:

$$d_t = f(\boldsymbol{D}_t), \tag{3}$$

where the set $\boldsymbol{D}_t = \{d_k : \boldsymbol{a}_k \in \boldsymbol{R}_t\}$, that is., the set of all past decisions that were taken in "similar" instances.

Depending on the application domain, there are a number of possible functional forms that have been suggested in the literature. For example, if d_t is a numerical decision, that is, just a number, then a simple average of the past decisions under similar instances can form a new decision.

$$d_t = \frac{\sum_{d_k \in \boldsymbol{D}_t} d_k}{Card(\boldsymbol{D}_t)}, \tag{4}$$

where *Card* indicates cardinality. In addition to the simple average, weights or weighting functions can be further used to differentiate the similarity among different $\mathbf{a}_k$ to $\mathbf{a}_t$.

$$d_t = \sum_{d_k \in \boldsymbol{D}_t} w_k d_k, \tag{5}$$

where

$$w_k = \frac{g(\epsilon_{k,t})}{\sum_{\{s:\boldsymbol{a}_s \in \boldsymbol{R}_t\}} g(\epsilon_{s,t})}, \tag{6}$$

The function g is a transformation of the similarity index $\epsilon_{s,t}$. If we let $\pi_k(t)$ be the most updated *strength* of the rule d_k, that is, the past experience (evaluation) of the performance of the rule d_k, then in addition to similarity $\epsilon_{k,t}$, the weight can also be adjusted based on $\pi_k(t)$. Hence,

$$w_k = \frac{g(\epsilon_{k,t}, \pi_k(t))}{\sum_{\{s:\boldsymbol{a}_s \in \boldsymbol{R}_t\}} g(\epsilon_{s,t}, \pi_k(t))}. \tag{7}$$

If the decision is in the form of discrete choices, then the function can be given with a stochastic choice formulation.

$$\text{Prob}(d_t = d_k) = \frac{g(\epsilon_{k,t}, \pi_k(t))}{\sum_{\{s:\boldsymbol{a}_s \in \boldsymbol{R}_t\}} g(\epsilon_{s,t}, \pi_k(t))}. \tag{8}$$

The above general discussion of the instance-based decision, with slight modifications, applies to a number of computational intelligence algorithms. Equations (4) to (6) constitute the basic form of K nearest neighbors (Chan et al., 1999). Equation (8) is a more general version of reinforcement learning.

K nearest neighbors

The method of K nearest neighbors (KNNs) is a typical experience-based computational behavioral model. In KNNs the idea of neighborhood, that is, Equation (2), is altered and instead of imposing an upper limit ϵ to define the set $\boldsymbol{R}_t$, KNNs select the K most similar instances or the K most nearest neighbors. We can rank $\epsilon_{s,t}$ in an ascending order and let the rank of $\epsilon_{s,t}$ be denoted as $R(\epsilon_{s,t})$. The set of similar instances, Equation (2), is then modified as follows.

$$\boldsymbol{R}_t = \{\boldsymbol{a}_k : k < t, \boldsymbol{R}(\epsilon_{k,t}) \leq K\}. \tag{9}$$

KNNs has been initiated thrice by different academic communities, first, by engineers (Cover & Hart, 1967), then by statisticians (Stone, 1977; Cleveland, 1979), and finally by physicists (Farmer & Sidorowich, 1987). From these three origins, we can see how the *similarity heuristic* is introduced as a heuristic in information processing and statistics, and then later on to serve a computational model of behavioral economics (Chan et al., 1999).

When our knowledge of the environment is incomplete or vague, our decisions naturally rely on or are biased towards familiar or similar experiences. The nearest neighbor was first used by Cover and Hart (1967) to give a notion of *similarity*:

> In the classification problem there are two extremes of knowledge which the statistician may possess. Either he may have complete statistical knowledge of the underlying joint distribution of the observation x and the true category θ, or he may have no knowledge of the underlying distribution except that which can be inferred from sample . . .
>
> *(Ibid.: 21)*

In the second extreme case, "a decision to classify x into category θ is allowed to depend only on a collection of n correctly classified samples $(x_1, \theta_{1,})$, $(x_2, \theta_{2,})$, . . . , $(x_n, \theta_{n,})$, and *the decision procedure is by no means clear*" (Ibid.: 21; italics added). With the absence of a clear decision procedure, Cover and Hart (1967) proposed the following heuristic:

> Thus to classify the unknown sample x we may wish to weight the evidence of the nearby x_i's most heavily. Perhaps the simplest nonparametric decision of this form is the *nearest-neighbor* (**NN**) rule, which classifies x in the category of its nearest-neighbor.
>
> *(Ibid.: 21; bold and italics original)*

KNNs was later introduced in the literature on robust local regression by Cleveland (1979). However, instead of having closeness or similarity as the main pursuit, the key focus here is on *smoothness*, specifically, the smoothness of the conditional density function. As commonly seen in functional approximation; its main goal is to regulate the polynomial degree of curve fitting. However, in addition to functional approximation, it is also fundamentally connected to the pursuit of simplicity in the science of discovery (Li & Vitanyi, 2008).

The *smoothness heuristic* is related to the *closeness heuristic* under the instance-based reasoning principle, where similar inputs are expected to have similar outputs. This principle implies a response surface which is simple in terms of its *descriptive complexity* or *algorithmic complexity*.[3] In other words, the instance-based decision model helps the decision maker to give a more concise description of her decision making process, specifically explaining why such a decision is made. Without the closeness and smoothness constraints, the simplicity of the decision-response surface may be lost, and, given the increased complexity, an automated decision becomes hardly available, and the decision will have to be left to "the man on the spot" (Hayek, 1945: 524–5). Such kinds of non-smooth decisions may be time-consuming, but their frequency must be limited, given the time constraint to which each decision maker is subjected.

In agent-based computational economics, *NN agents* were first used in an agent-based artificial stock market (Chan et al., 1999). The NN agent forecasts the price based on a moving window with a length l, which is also known as the *embedding dimension*. Let $p_t = ln^{(P_t/P_{t-1})}$, where P_t is the asset price at time t and ln denotes the natural log. Furthermore, let

$$\boldsymbol{p}_t^l = (p_t, p_{t-1}, \ldots, p_{t-(l-1)}). \tag{10}$$

To forecast p_{t+1}, the NN agent will find the past K historical windows (instances) which are most similar to $\boldsymbol{p}_t^l$, that is,

$$\boldsymbol{R}_t = \left\{\boldsymbol{p}_k^l : R(\epsilon_{k,t}) \leq K\right\}, \tag{11}$$

where $\epsilon_{k,t} = corr(\boldsymbol{p}_k^l, \boldsymbol{p}_t^l)$. An average of the price p_{k+1} will then be used as the forecast of p_{t+1}.

$$p_{t+1}^{\epsilon} = \frac{\sum_{\{k:\boldsymbol{p}_k^l \in \boldsymbol{R}_t\}} p_{k+1}}{K} \tag{12}$$

A difficult part of the instance-based decision is to address how instances are formed in the first place. In many real-life situations, it can be hard to tell whether two instances are closely related or similar. A proposed distance or similarity measure can be sensitive to different attribute spaces. Some critical but hidden attributes could be ignored and may never be found. Nevertheless, what matters is not whether the decision maker has built her decision upon the "true" attribute space, but instead whether they actually follow instance-based reasoning to streamline their decisions. It can be argued that without such a framework, the decisions can be harder and may be less satisfactory. Accordingly, as shown in Figure 21.1, the instance-based decision making addresses the needs of a less loaded decision making process. Amartya Sen termed the situation *decisional inescapability*, in that a decision or a choice has to be made even before the completion of a judgement process (Sen, 1997). To cope in such instances, decision makers may have to learn and evolve to develop various heuristics, such as the instance-based decisions, to handle these otherwise inescapable situations. The often observed decision making based on *stereotypes* can be interpreted as an instance-based decision (Bodenhausen, 1990; Chaxel, 2015; Fabre et al., 2015). Again, here, the stereotype attached to a specific instance, say, a person, a city, a country, a gender, a culture, or a brand, etc., can be imprecise, but what matters is that this frame facilitates decision making, particularly when a reason is needed or when the time available for making the decision is severely limited. In fact, as we shall see below, evolutionary computation can allow agents to discover useful instances, which constitutes a part of the learning for agents (Figure 21.1).

K-means and self-organizing maps

The number of nearest neighbors, that is, K, obviously, is a key parameter in the KNN algorithm. The question of the optimum number of K has been addressed in the third of the above-mentioned intellectual origins of KNNs, that is, the chaotic-dynamics origin (physicist approach) (Takens, 1981).[4] In this stream of the literature, it has been shown that, based on the Takens theorem, KNNs can help forecast the chaotic time series, specifically, the deterministic chaotic time series. To do so, the parameter K is determined by the embedding dimension l (Equation 10). It has been suggested that $k = 2(l + 1)$ (Casdagli, 1991), but, under the case of stochastic non-linear systems, it also depends on the noise level: the higher the added noise level, the higher the K. Nonetheless, the above analysis is entirely from a mathematical viewpoint. From a cognitive viewpoint, a number of other considerations need to be incorporated.

First of all, how can humans actually retrieve similar instances from their memory? And how many such instances can be retrieved? Considering the brain with its limited capacity for memory, a pertinent question concerns how the brain deals with increasing information by not memorizing all of it or by forgetting some of it. How does it do the much necessary *pruning*? This is still a non-trivial issue pursued by neuroscientists today.[5] This suggests a role for redundancy-reduction behavior. Hence, similar instances, given a certain tolerance level of noise, may be combined into one instance. A large number of instances are then substantially reduced to a few representative instances. Hence, when making a new decision, the number of referred neighbors may be very low, say, close to those magic numbers which psychologists normally refer to (Miller, 1956; Mathy & Feldman, 2012).[6] The computational model of the aforementioned compression behavior is known as a clustering algorithm in computational intelligence, and the two popularly used clustering algorithms are *K-means* and *Kohonen's self-organizing maps* or SOMs (Kohonen, 1995). *K*-means clustering, developed by MacQueen (1967), is one of the widely used clustering algorithms that groups data with similar characteristics or features together. SOMs resemble *K*-means. They both involve minimizing some measure of dissimilarity, called the cost functions, in the instances within each cluster. The difference between the *K*-means and the SOM lies in their associated cost functions. Consider a series of n instances, each of which has M numeric attributes:

$$\boldsymbol{a}_1^M, \boldsymbol{a}_2^M, \ldots, \boldsymbol{a}_n^M, \boldsymbol{a}_i^M \in \boldsymbol{R}^M, \forall i = 1, 2, \ldots, n \tag{13}$$

where

$$\boldsymbol{a}_i^M \equiv \{a_{i,1}, a_{i,2}, \ldots, a_{i,m,}\}.a_{i,l} \in \boldsymbol{R}, \forall l = 1, 2, \ldots, M \tag{14}$$

The *K*-means clustering is to find a series of k clusters, the centroids of which are denoted, respectively, by

$$C_1, C_2, \ldots, C_k, C_j \in \boldsymbol{R}^M, \forall j = 1, 2, \ldots, k \tag{15}$$

such that each of the observations is assigned to one and only one of the clusters with a minimal cost, and the cost function is defined as follows:

$$C_{K-means} = \sum_{i=1}^{n} \sum_{j=1}^{k} \left\| \boldsymbol{a}_i^M, C_j \right\| \delta_{i,j}, \tag{16}$$

where $\left\| \boldsymbol{a}_i^M, C_j \right\|$ is the standard Euclidean distance between $\boldsymbol{a}_i^M$ and $\boldsymbol{C}_j$, and $\delta_{i,j}$ is the delta function:

$$\delta_{i,j} = \begin{cases} 1, & \text{if } \boldsymbol{a}_i^M \in C_j \\ 0, & \text{if } \boldsymbol{a}_i^M \notin C_j \end{cases} \tag{17}$$

To minimize the cost function (16), one can begin by initializing a set of k cluster centroids. The positions of these centroids are then adjusted iteratively by first assigning the data samples to the nearest clusters and then recomputing the centroids. Corresponding to (16), the cost function associated with SOM can be roughly treated as follows:

$$C_{SOM} = \sum_{i=1}^{n} \sum_{j=1}^{k} \left\| \boldsymbol{a}_i^M, C_j \right\| \cdot h_{\omega(\boldsymbol{a}_i^M),j} \tag{18}$$

where $h_{\omega(\boldsymbol{a}_i^M),j}$ is the neighborhood function or the neighborhood kernel, and $\omega(\boldsymbol{a}_i^M)$, the winner function, outputs the cluster whose centroid is nearest to the input $\boldsymbol{a}_i^M$. In practice, the neighborhood kernel is chosen to be wide at the beginning of the learning process to guarantee the global ordering of the map, and both its width and height decrease slowly during learning. For example, the Gaussian kernel whose variance monotonically decreases with iteration times is frequently used. By comparing Equation (16) with (18), one can see that in SOM the distance of each input from all of the centroids is weighted by the neighborhood kernel *h*, instead of just the closest one being taken into account. Through either KNNs or SOM, our experiences of the past can then be constantly processed by clustering, which provides us with *points of reference* or *anchors* upon which the subsequent decisions can be based and facilitated.

Reinforcement learning

In the context of *discrete choice*, Equation (8) is a more general version of reinforcement learning. To see this, simply impose the requirement that ε to zero, that is, only consider those perfectly identical instances, and require *g* to be a Gibbs–Boltzmann distribution with the temperature parameter λ,

$$\text{Prob}(d_t = d_t) = \frac{exp^{\lambda\pi_k(t)}}{\sum_{\{s:\boldsymbol{a}_s \in R_t\}} exp^{\lambda\pi_s(t)}}, \tag{19}$$

in which case we have a Roth–Erev version of reinforcement learning (Roth & Erev, 1995).

Reinforcement learning has already been applied to explain or predict human behavior in the context of game experiments. It is considered to be consistent with the robust properties of learning observed in the large experimental psychology literature on both human and animal learning, specifically, the *Law of Effect* (Roth & Erev, 1995).[7] The recent progress in neuroscience indicates that humans, and more generally, mammals are naturally endowed with a reinforcement learning mechanism in their brains. In fact, one of the most impressive recent results in neuroscience is the discovery of the relationship between the dopamine neural system and reinforcement learning.[8] Technically, reinforcement learning has been extended to take into account a number of psychological factors in learning, such as memory (Roth & Erev, 1995), counterfactual thinking (Camerer & Ho, 1999), aspiration (Erev & Roth, 1998) and attention (Chen & Hsieh, 2011).

The standard version of reinforcement learning only considers a fixed and finite set of alternatives, since the decision environment is homogeneous. The typical example used to illustrate this decision environment is the *multi-armed bandit problem* (Bush & Mosteller, 1955). The decision maker at each time is always offered a fixed number of bandits, and, since instances are always the same, the decision can be automated by using the stochastic choice formulation given in Equation (19). In a special case where $\lambda = 0$, the default turns out to be the one with the highest updated strength (most successful experience), or simply, the best one so far. In this special case, it is similar to the *take-the-best heuristic* (Gigerenzer, 2007), a member of *one-good-reason heuristics*.[9] The generalized version, Equation (7), simply adds a hierarchical structure to the set of rules by classifying them according to their applicability to a certain instance.[10] Hence, each instance corresponds to a specific set of rules with different strengths. The set of rules may be globally the same over different instances, but their respective weights (strengths) and hence priorities can differ from one instance to another. In behavioral economics, reinforcement learning has been proposed as a model of *low rationality* (Erev & Roth, 1998; Duffy, 2006; Chen, 2013). This original intention may lead people to misperceive it as a mere model fitting for very simple behavior in a rather recurrent decision environment.[11] However, as we shall see, this is not the case. Not only can reinforcement learning serve as a model to handle novel situations, but it can also serve as a meta-level learning model, that is, *to learn how to learn*. Vriend (2002) is the best illustration to exemplify these two features.

Vriend (2002) considers the kind of decisions which are unique and, hence, not repeated (not similar). Examples can be buying a car, buying a house, choosing a restaurant in Pinamar, and booking a hotel in Reykjavik. Hence, strictly speaking, reinforcement learning cannot be directly applied in these situations, since available alternatives (available experiences) are not transferrable (commutable) from one place to the other. Nevertheless, with such a series of novel situations, one can learn from the experiences of others, the so-called *social learning*, and there are different ways to learn from others (Nowak, 2006; Scott, 2012). Vriend considered three types of rules, namely, randomly behaving rules (throwing a coin), following what the majority did (herding), or replicating the good experiences of others.

These three types of rules can always be applicable to any novel situation, as long as the decisions made by others and their resultant experiences are available. In fact, Vriend (2002) can be read as a contribution to the economy of Web 2.0 and the agent-based study of Big Data in the following sense. First, as mentioned in Chen, Chie and Tai (Chapter 18), the essential characteristic of the Web 2.0 economy concerns the user-initiated and user-supplied content, and the online customer review is one major form of digital content. Second, while online customer review reports can help consumers acquire more information on the quality of the product, their fast accumulation can result in an overload of information for consumers. To understand how consumers make use of this digital content, the aforementioned three types of rules seem to be a reasonable beginning. The randomized rule does not require any cognitive effort from the decision maker. The second one needs only a counting of heads. The last one needs to read the reviews and to know users' experiences; hence, it may be more time-consuming. Reinforcement learning can then be applied to these three levels of learning: no learning, shallow learning, and deep learning. Reinforcement learning can then serve as a model of meta learning.

Hierarchical structure of decisions

Quite contrary to what is usually taught in economics, many of our decisions or choices are not always based on insufficient information, but on overloaded information. In behavioral economics, this conundrum is known as the *information overload hypothesis*.[12] A typical heuristic to make a decision in such a situation is not to look at all information at once; instead, information will be given a sequential or hierarchical structure so that one needs to get access to more information only when the decision cannot be made based on the "abridged" version. Because of this practical need, a tree or a hierarchical structure can play quite a crucial role in decision making or choice making.

Decision trees

The decision tree, a canonical model in computational intelligence, can be interpreted as a computational behavioral model corresponding to the hierarchical structure of decision making. Suppose that we are interested in knowing how a tennis player decides whether to play tennis. We have a sequence of observations of her past decisions,

$$(\boldsymbol{D}_T, \boldsymbol{A}_T) = \{(d_t, \boldsymbol{a}_t)\}_{t=1}^{T}$$

where d_t is a binary decision variable, either to play $d_t = 1$ or not to play $d_t = 0$. $\boldsymbol{a}_t$ can be a vector of attributes which may help define an instance; for example, outlook, humidity and wind, if she is only concerned with the weather condition.

A decision tree is constructed based on a *top-down greedy algorithm*, known as the ID3 in machine learning (Quinlan, 1986). The key idea is fairly straightforward. First, one finds the attribute a^*, say, outlook, that *best* classifies $\boldsymbol{D}_t$, and then uses this attribute as the *root* of the decision tree. The process is then repeated for each subtree. The main issue in this greedy algorithm concerns the criterion regarding the choice of the best classifying attribute. A common solution to this problem is to select the attribute with the *highest information gain*, which is defined as the expected reduction in the *entropy* of the dataset $\boldsymbol{D}_t$ caused by knowing the value of the attribute $\boldsymbol{A}_T^* = \{a_t^*\}_{t=1}^T$.

An illustration of a decision tree which is built is given in Figure 21.2. In this illustration, among a sequence of information, the tennis player will first look at the outlook, and there are three values for the outlook: sunny, overcast, and rainy. If the outlook is overcast, then the tennis player will simply disregard the unread information and will decide to play tennis. On the other hand, if it is not overcast, then the information (the second attribute) to be further examined depends on whether the outlook is sunny or rainy. The second attribute is humidity if the outlook is sunny, and wind, if the outlook is rainy. In each of these two branches, the decision can always be made without further looking into the remaining information. In other words, although each instance is defined by three attributes, at any given time at most two attributes are required in order to make a decision.

Decision tree has been considered to be a fast and frugal heuristic in behavioral economics (Gigerenzer, 2007). It might, therefore, be worth discussing the connection between machine learning and behavioral economics in their respective use of decision trees. First of all, the top-down greedy algorithm as introduced by the artificial intelligence (AI) community is applicable to the study of the real decision process; for example, in using it for analyzing the observations of human-subject experiments. In fact, the idea of decision trees has already been used as a model to analyze and understand the decision making observed in human-subject experiments, such as the prisoner's dilemma games (Axelrod, 1984), ultimatum games (Duffy & Engle-Warnick, 2002), and trust games (Rieskamp & Gigerenzer, 2002; Engle-Warnick & Slonim, 2004, 2006). The heuristics studied in these papers, such as the TIT-FOR-TAT, can be presented in the form of a decision tree heuristic. However, none of these studies has formally applied the top-down greedy

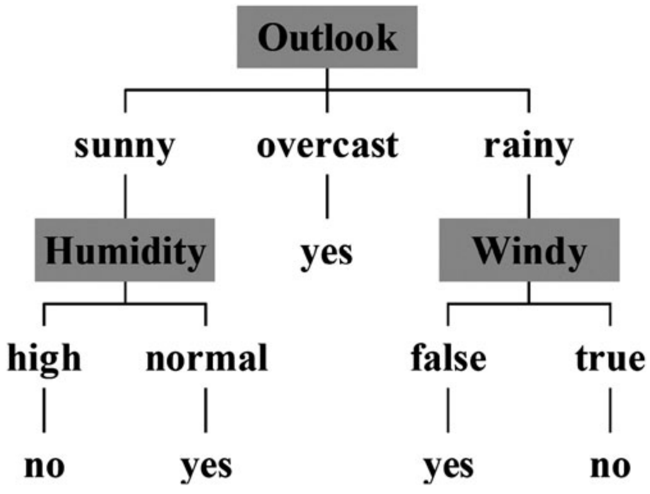

Figure 21.2 The decision tree of the *play tennis* decision

algorithms to build and formulate a decision tree heuristic; therefore, there is room for applying the decision tree model to discover the decision tree heuristics followed by human subjects in experimental or real data (Tagiew, 2012; Rosenfeld et al., 2015).

Second, while the top-down greedy algorithm can be useful for data mining and rule extraction, the algorithm per se may not provide a good description of the process of formation of these heuristics from a behavioral viewpoint. For example, humans may find the root attribute, "outlook" in Figure 21.2, based on their intuition, experience or preferences. In the case of the tennis player, putting outlook as the root attribute may be entirely due to the player's enjoyment in playing, but it may also be due to her past performance under different weather conditions. Hence, what is needed in behavioral economics is a learning (formation) process for the decision tree heuristics that are employed.

Incremental reinforcement decision tree construction[13]

The learning (formation) process includes two parts: first, the list of all relevant attributes, and, second, their ranks (positions) in the decision tree. The first issue is more complex and involves the discovery process, which we will show later on. Once a_t is determined, the second issue can be answered by reinforcement learning. Assume that decision makers begin with the *one-reason heuristic* and try to find out the best attribute, and then make a decision based on that attribute. In our tennis player example, the three attributes will compete for the attention of the tennis player at the first stage. After a while, overcast is selected through reinforcement learning as the first attribute, and the decision is:

IF ((Outlook=overcast)

THEN YES (Play Tennis))

As time goes on, the player may then discover that when the outlook is not overcast, he could still have fun playing tennis, and a competition for the second attribute is triggered again through another reinforcement learning cycle, which leads to the identification of humidity and wind as the second attribute under different branches of Figure 21.2, and the newly developed decision tree is: IF ([(Outlook=overcast)] OR [(Outlook=sunny) AND (Humidity=normal)] OR [(Outlook=rain) AND (Wind=weak)]) THEN YES (Play Tennis)).

In sum, the above proposal is to replace the original top-down greedy algorithm with incremental reinforcement learning. In this way, a learning (formation) process of the decision tree heuristic is articulated. The essence of the proposed behavioral algorithm is that it is *incremental*; basically, it decomposes the entire tree formation process into many "multi-armed bandit problems" and applies reinforcement learning to each of these bandit problems. Hence, as we have learned from Vriend's model, reinforcement learning can be applied generally to a meta-level of learning, and hence is much more powerful than it might seem.

In terms of understanding the human decision making process, decision trees can also be compared to the frequently used multivariate regression models, including the probit and logit models. First, human decisions may be fitted well by both these approaches, but multivariate regression only gives a summary of decision making, rather than a *process* of decision making. Hence, when trying to give an account of how a specific decision is made, it is easier to communicate using decision trees rather than by using multivariate regression. Second, when making a decision, multivariate regression essentially needs decision makers to pay attention *simultaneously* to multiple attributes, whereas decision trees only require them to focus on one attribute at a time.

From the viewpoint of cognitive loading, decision trees are less demanding than multivariate regression.[14]

Evolutionary computation

Autonomous agents

Evolutionary computation plays a critical role in the development of behavioral economics, in particular, the contribution to crystallizing the idea of *autonomous agents*, that is, agents who are able to discover chances or novelties without external guidance, in particular, without those "interventions" from modelers themselves. Behavioral economics has long criticized the notion of *Homo economicus* used in mainstream economics, but their proposed alternative, *Homo sapiens*, also suffers from operational emptiness. John Tomer's recent proposal on the notion of *smart persons* may not be an entirely new idea, but it clearly reveals the fact that the boundedly rational agents in behavioral economics have a blurred face (Tomer, 2015).[15] The *missing ingredient*, as Tomer calls it, in our view is exactly a notion of autonomous agents. One reason that the autonomous agents have not been well incorporated into behavioral economics is the lack of toolkits. It would probably be fair to say that the tools available for economists to build chance-discovering or novelty-discovering agents[16] with a moderate degree of autonomy were rather limited before the early 1990s.

In the early 1990s, genetic algorithms were formally introduced to economics as a tool to construct autonomous agents (Holland & Miller, 1991). The notion of autonomous agents is crucial for behavioral economics since a set of heuristics, be they biased or frugal, should not be taken as given, except those which are proved to be genetically driven and are innate. In general, the employed heuristics are constantly evolving and, as time goes on, new heuristics may be discovered. In a nutshell, heuristics should not be treated as scientific laws; instead, they can be best understood as an evolutionary process.

A good illustration of the evolution of heuristics as well as personal traits is the integration of gambling psychology in an agent-based lottery market (Chen & Chie, 2008). In their model, Chen and Chie (2008) incorporated three characteristics into their gambling decision making model; these three are the halo effects (lottomania)—related to participation ratio, conscious selection, and aversion to regret. What differentiates their model from the typical behavioral models is that these three characteristics are not imposed *exogenously*, but are probabilistic emergent properties.

A bit string, also known as a *chromosome* in genetic algorithms, is used to code the three characteristics of agents, and after decoding one can know the state of each characteristic, as shown in (20).

$$\overbrace{\underbrace{1001\ldots\ \ldots\ \ldots..0001}_{20\,bits}}^{participation\ ratio} \mid \overbrace{\underbrace{100\ldots\ \ldots\ \ldots..111}_{16\,bits}}^{conscious\ selection} \mid \overbrace{\underbrace{0\ldots\ \ldots\ \ldots\ \ldots\ \ldots\ \ldots 1}_{4\,bits}}^{degree\ of\ regret\ aversion} \tag{20}$$

The standard single-population genetic algorithm is then applied to evolve a population of these randomly-generated bit strings, characterizing the initial heterogeneities of gamblers on these characteristics. One can then observe how each of these characteristics changes over time, both at the individual level and aggregate level. From a market design perspective, Chen and Chie (2008) studied the effect of the lottery tax rate on the population size of non-gamblers (agents with a zero lottery participation rate). While the expected return of "investing" in a lottery is negative, the gamblers will not be driven out by the market selection mechanism defined by genetic algorithms. In addition, from probability theory, while conscious selection of winning numbers does not make any sense, Chen and Chie, however, showed that a rather moderate degree of conscious selection behavior will remain in the market; hence, the market also fails to drive out this "irrational" behavior. Perhaps the most intriguing part concerns their analysis of the regret aversion behavior. It was found that the attention to other gamblers' rewards (jackpots), a kind of social preference, may co-evolve with their devotion to gambling; both are co-determined by the lottery design (the lottery tax rate). Specifically, when the lottery tax rate is high, the size and the winning probability of jackpots become low and the gamblers' devotion also decreases, accompanied by their greater pleasure in being released from the possible regrets of not gambling. This exemplifies how evolutionary computation can work with behavioral economics by making the implicit selection process explicit and by providing a test for the stability of these behavioral patterns.

Hierarchical modularity

If behaviors (routines and heuristics) are not static but are constantly evolving, then one has to ask what the universal representation of the behavior of the evolution of behaviors is. In this section, inspired by Simon (1962), we propose hierarchical modularity as the fundamental representation.[17] Generally speaking, modularity refers to the idea of self-encapsulated, independently operationable, and reusable (evolvable) routines, procedures or programs. It provides us with a constructive way to think about what a decision making system is, and, in particular, how a decision maker can cope with complexity and survive in the constantly evolving environment. Take decision trees as an example. Each decision tree can be perceived as a module, and, therefore, can be used to construct a bigger decision tree if the root of the child tree is an emanating node of some mother trees. In other words, decision trees can be considered as a special case of hierarchical modularity.

In computational intelligence, the idea of modularity can be realized by genetic programming (Koza, 1992). Instead of working on finite-length strings (bits), genetic programming directly operates on the space of computation programs which are represented using the formal language theory, specifically, the context-free grammar (Linz, 2006). Starting with a finite set of alphabets (primitives) and following the given grammar (production rules), one can develop phrases, sentences, paragraphs, chapters, books, all the way up without a limit. In each stage of this development, simpler or lower-level modules are used to construct sophisticated or higher-level modules, and this process can continue without an end. To understand the meaning of a decision rule, one only needs to harness its immediate constituents (modules). Since each module is already encapsulated, there is no need to go further down to their modules, and their modules' modules, and so on. The modular structure, therefore, reduces the huge amount of information required in applying a rule or making a decision.

Genetic heterogeneity

In addition to being a tool for the computational behavioral model of searching and discovery, evolutionary computation also contributes to behavioral economics by generating agents with heterogeneous traits. Recently, there has been a growing attempt to explore the genetic influence that concerns human decision making. Some recent areas of focus in behavioral economics, such as self-control, impulsivity, addiction, patience, risk preference, and cognitive capacity, are being examined for possible heritable factors. The literature on this area continues to grow.

In 2007, Daniel Benjamin and his colleagues gave this nascent field a neologism: *genoeconomics* (Benjamin et al., 2007).

The relation between cognitive capacity and decision making has become an issue of focus in this stream of the literature. Earlier genoeconomic studies have indicated a possible pathway from genetic causes to cognitive capacity, to education and to income. Recently, the decision making capability under an uncertain environment has also been included as a part of this pathway (Beauchamp et al., 2011; Callaway, 2012; MacKillop, 2013; Ashraf & Galor, 2013). In parallel, experimental economists have also begun to design human-subject experiments to examine the possible effects of cognitive capacity on economic decisions.[18]

If cognitive capacity does affect decision making, including both processes and outcomes, then what will be the ideal computational model to take account of this factor? Recently, it has been suggested that the *population size*, a key parameter used in evolutionary computation, can be regarded as a proxy variable for cognitive capacity (Casari, 2004; Chen, Tai, & Wang, 2010). In physical terms, population size is related to space complexity in computation theory. The logistics of a complex product requires many intermediate steps and, hence, needs a large space to store and to integrate intermediate products. If the space is not large enough, then a complex product may be beyond the affordability of all available logistics. Hence, population size directly determines the capability of parallel processing of many intermediate tasks.

On the other hand, the working memory capacity of a human being is frequently tested based on the number of cognitive tasks that humans can simultaneously process (Cappelletti, Guth & Ploner, 2008). Dual tasks have been used in hundreds of psychological experiments to measure the attentional demands of different mental activities (Pashler, 1998). Hence, the population size seems to be an appropriate choice with regard to mimicking the working memory capacity of human agents; in this sense, evolutionary computation can directly control the 'cognitive capacity' of a computational behavioral model through varying population size. The heterogeneity of cognitive capacity of different human subjects can be represented by a society of artificial agents driven by genetic algorithms or genetic programming with different population sizes.

The proposed computational behavioral model of cognitive capacity, working memory capacity (WMC), has been applied to agent-based double auction markets to examine the effect of WMC on earning performance (Chen, Tai, & Wang, 2010).[19] It is found that the artificial traders with larger WMC can earn more than the artificial traders with smaller WMC. However, this dominance becomes less (statistically) significant when WMC increases further. Moreover, if we allow artificial traders with lower WMC more time to learn so that their deficiency in terms of WMC can be compensated by the longer time of learning (evolution), then the above income gap can disappear if the difference in WMC among traders is limited; otherwise, the gap can only be narrowed, but it will not disappear. Therefore, the above simulation shows that even though the double auction market is an easy environment, it can still generate persistent income inequality if the heterogeneity in the cognitive capacity of traders is significant enough.

Ant as a model of human behavior

Our next section focuses on the ant colony optimization algorithm, another computational intelligence tool that is frequently used in the context of optimization, such as the travelling salesman problem (Dorigo & Stützle, 2010). Compared to some other CI tools, such as reinforcement learning and evolutionary computation, the ant algorithm or, more generally, swarm intelligence is relatively less familiar to behavioral economists. Due to the important contributions by Alan Kirman (1991, 1993), economists have a chance to access interesting findings and puzzles related to ants' foraging behavior.

Earlier entomological experiments, cited in Kirman (1993), have shown that ants' foraging behavior over two identical equidistant food sources can demonstrate constant asymmetric distribution over the two sources; say, one source attracts the majority of ants and the other source attracts the minority of ants. Furthermore, as time goes on, the majority side and the minority side will switch without any external environmental changes. In other words, ants can collectively generate an endogenous fluctuation of their foraging distribution over the two sources of food. While this is an entomological finding, it has some significant implications for economics and other social sciences. Its possible implications have been well surveyed in Kirman (1993), including providing support for a fundamental instability in financial markets.

The underlying mechanism for this endogenous switching is known as a communication mechanism called *stigmergy*. The communication among ants is not necessarily direct, but more indirect, partially due to their poor visibility. The ants' reliance on indirect communication has been noticed by the French biologist Pierre-Paul Grasse (1895–1985), and he termed this style of communication or interaction *stigmergy* (Grosan and Abraham, 2006). He defined stigmergy as: "Stimulation of workers by the performance they have achieved." Stigmergy is a method of communication in which the individuals communicate with each other via modifying their local environment. For ants, this is achieved by the release of pheromone along their foraging trails.

However, the essence of these algorithms is to have an explicit modeling of social interactions on individual behavior. These algorithms are again built on empirical grounds, in this case, entomological experiments. Due to the nature of entomology, one would hardly argue whether these ants or locusts or other low-level swarms are consciously choosing to do anything "rational"; studies of their behavior tend to be more in the biological or neurological direction (Garnier, Gautrais & Theraulaz, 2007; Beekman, Sword & Simpson, 2010). Hence, the experimental results obtained here seem to put more focus on the effect of social interactions on emission or release of chemical materials, such as pheromone in the case of ants, or neurotransmitters, such as serotonin in the case of locusts (Paula et al., 2015).

We have known that social interactions have many channels to affect agents' decision and behavioral rules, such as social norms, social conformity, homophily, etc. In Kirman's ant model, the proposed social interaction mechanism is binary so that only a simple stochastic process, an urn process, is introduced to determine how one agent's decision can be affected by a randomly encountered agent. In computational intelligence, the behavioral algorithm is more explicitly related to the accumulated pheromone or accumulated serotonin, hence even though the decision can still be random, it is stochastic in a way related to various characteristics of social interactions, such as the degree of social polarization and the size of social network (for the concern of social conformity) (Valentini & Hamann, 2015). This type of algorithm essentially allows us to address the connection between social interactions and individual decisions through the biological and neural mechanisms. In this regard, the development of swarm intelligence stands in a unique position in computational behavioral economics in the sense that it can effectively incorporate the findings of neuroscientific experiments with these insects into the behavioral algorithms proposed for these swarms. Since entomological experiments are easier to implement, we hope that the behavioral economists can gain some useful insights, which are more difficult to glean from human fMRI experiments.

Can randomization be a heuristic?

All the heuristics reviewed up to this point correspond to some degree of learning from either one's own or others' experiences and reasoning with them. There is, however, a heuristic which requires no memory, no learning, and no reasoning. This is known as the *zero-intelligence heuristic*, to which we now turn.

The zero-intelligence (ZI) agent has been one of the widely employed characterizations of an agent in agent-based models and it has had a remarkable impact in both economics and finance (Ladley, 2012). The supposed simplicity of this kind of agent stems from their lack of strategy and their random behavior. Gode and Sunder (1993), and many since then, have employed this device to illustrate the irrelevance of a high level of sophistication in strategies and learning at the individual level in achieving market level efficiency.[20]

ZI agents or randomly behaving agents have been employed in wider contexts that range beyond a mere device to separate the effect of strategies from that of the market mechanism. The rationale for this agent design is that the individual level details become worn out in the aggregate with a large number of heterogeneous agents. Another reason advanced is the lack of precise knowledge about strategies used by different agents, at any given point in time. Hence, modeling them as if they behave in a random fashion (from a bounded set of strategies) allows one to not commit to one strategy *a priori*. Consequently, "zero intelligence agent" may be a misnomer and *entropy maximizing agents* can serve as a better term. This is because the relationship between zero intelligence, cognitive ability and the ease or the simplicity of random behavior may not be as obvious or straightforward.[21]

While there may be a case to start with entropy maximizing agents in the face of ignorance, their behavioral underpinnings ought to be scrutinized. The *entropy maximizing* role needs to be distinguished from *random behavior* as being a proxy for simplicity or naivety in terms of strategies (or a lack of them). By relating "zero intelligence" to random behavior, the implicit assumption is that random behavior is simple to execute and that it requires very little sophistication. To design artificial economic agents more like human agents, we need to examine whether the programmed actions have a psychological or behavioral foundation. Hence, the plausibility of human beings to be able to "behave" in an analogous fashion and the associated cognitive demands need to be studied. In this context, it is therefore natural to question the ability and the extent to which human agents can choose strategies randomly. More generally, we need to examine whether it is behaviorally plausible for an agent to act randomly and for the others to perceive such an action to be random.

Studies from psychology indicate that the human ability to perceive randomness and act randomly may be limited (Wagenaar, 1972). This problem can be subdivided into the ability to perceive, discriminate and generate random behavior, each of which is far from easy. In the light of limited memory, cognitive limitation (Hahn & Warren, 2009) and finiteness of data, detection and execution of random or patternless behavior seems notoriously hard (Kahneman & Tversky, 1972). This is further complicated by difficulties in the characterization of randomness when the data are finite. Even a supposedly elementary task of generating random sequences has been found to be a non-trivial, difficult exercise for human subjects in experimental environments.[22] In addition, the distinction between the perception and identifiability of randomness raises further questions about the indiscriminate use of randomly behaving agents in strategic and interactive environments that one often encounters in economics and agent-based models (Zhao, Hahn, & Osherson, 2014). If randomness is interpreted as a lack of a pattern or rule in the sequence of responses generated, then such random behavior requires the avoidance of any discernible pattern. Interpreted this way, random behavior may require far more intelligence, cognitive ability, and sophistication than otherwise assumed.

In sum, although randomization in the form of entropy maximization may be often considered as a cognitively effortless heuristic, our review indicates that this "stereotype" may not be entirely correct; hence, without relying on an external device, such as a coin, dice, or an oracle, making a truly random decision may not be that easy for the human brain.

Concluding remarks

Computational intelligence or machine learning has been developed independently of behavioral economics over a period of about three decades. Before this, and even through this period, the dominating approach regarding decision making in economics has been probability and statistics, upon which the rational expectations revolution has been built. The formulation of decision making in the mainstream economics literature is basically the application of statistical decision theory, which, in turn, is the application of von Neumann and Morgenstern's expected utility maximization framework (Ferguson, 2014). Computational intelligence is a credible alternative to this paradigm. Instead of a model driven approach, it is mainly a data-driven or an experience-based approach. Instead of being restricted to a "small world" (Savage, 1972), it mainly deals with uncertainty in a "large world" in which a proper probabilistic formulation of the world is often infeasible.

Computational intelligence relies on various heuristics to build another set of guidelines to learn from the past, to cope with complexity, and to make decisions. Some of these heuristics that are reviewed in this chapter include similarity, closeness, smoothness, reinforcement, default, automation, hierarchy, and modularity. These heuristics together help shape what is known as *behavioral AI*, to be distinguished from classical AI or symbolic AI (Wooldridge, 2009).[23] We believe that computational intelligence can consolidate and enrich the study of behavioral economics by providing the computational underpinnings of decision making processes. This direction, referred to as computational behavioral economics, will also enhance the interdisciplinary conversations between behavioral economics and other related disciplines.

Acknowledgments

The authors have benefitted from two excellent referee reports, and we are grateful to Manuel Scholz-Wäckerle and Andreas Duus Pape for their painstaking review of the paper with many valuable and constructive comments and suggestions. The first author and the third author are grateful for the research support in the form of Ministry of Science and Technology (MOST) grants, MOST 103-2410-H-004-009-MY3, and MOST 104-2811-H-004-003, respectively.

Notes

1 For a comprehensive treatment of Hayek's contribution to behavioral economics, the interested reader is referred to Frantz and Lesson (2013).

2 In the literature, it is also known as instance-based decision or instance-based reasoning; in this chapter, we shall use these terms interchangeably.

3 Descriptive complexity looks into the amount of information required to "describe" a given, finite, binary sequence, but more generally to mathematical objects. The notion stems from the work of Andrey Kolmogorov (1903–87), Ray Solomonoff (1926–2009) and Gregory Chaitin in the 1960s, and the description needs to be understood in terms of output of ideal Turing machines. This is different from computational complexity. See Li and Vitanyi (2008).

4 Takens' paper was in the context of fluid dynamics, identifying procedures to decide whether or not to attribute experimental data to the presence of strange attractors. By *the physicist approach* what we mean is those from the dynamical systems origin.

5 The same issue can interest economists as well, because it concerns the efficient use of limited space. An early study on reward-motivated memory formation by neural scientists may provide an economic foundation for the memory formation (Adcock et al., 2006). Adcock et al. (2006) reports brain-scanning studies in humans that reveal how specific reward-related brain regions trigger the brain's learning and memory regions to promote memory formation.

6 The magic number *seven*, originally proposed by Miller (1956), is a measure related to short memory capacity or working memory capacity, characterized by the number of items that an individual can

discriminate or remember over very short periods of time, say, seconds. Based on a few experiments that he reviewed, Miller concluded that most people can correctly recall about 7 ± 2 items. Different numbers have been proposed in this stream of literature.

7 Reinforcement learning has also been used to explain institutional change, more precisely, the interdependence between economic behavior of agents and institutional change. See, for example, Heinrich and Schwardt (2013).

8 See Montague (2006), chapter 4, for a vivid historical review of the research on the dopamine system and reinforcement learning.

9 By one-good-reason heuristics, agents focus on *only one good reason* or *cue* to make a decision, rather than considering all cues and weighting them. Contrary to expectations, they are not just fast, but also more accurate in a variety of environments (Snook et al., 2005; Gigerenzer and Gaissmaier, 2011).

10 While we use the term hierarchy, Equation (7) is not the *hierarchical reinforcement learning* normally formulated in the context of a Markov decision process (Barto and Mahadevan, 2003) and recently applied to computational neuroscience (Botvinick, 2012). The kind of decision considered by us in this chapter is not Markovian, but a type of reinforcement learning model frequently used by experimental economists. The usual hierarchical reinforcement learning models use the idea of subroutines, macro procedures, modularity, or the so-called abstraction states to deal with the curse of dimensionality. We shall come back to this idea in the fifth section.

11 This ideal environment is very similar to the situation depicted by the movie *Groundhog Day* as briefly mentioned in Thaler (2000).

12 Given that there are other chapters devoted to this subject, for example, Chen, Chie and Tai (chapter 18), to avoid redundancy, we shall not elaborate on this hypothesis further.

13 We are grateful to Andreas Pape for suggesting this term to replace our originally proposed term, incremental reinforcement learning.

14 This is specific when we consider some cognitive constraint, such as Miller's magic number, *seven* (Miller, 1956). The point we wish to make in this paragraph is not on which decision model is real, i.e., the actual mechanisms/processes that exist in people's heads, since both multivariate regressions and decision trees may not be real decision processes. Instead, we address the transparency of a decision model in its dynamics.

15 In a spectrum between Homo Economicus and Homo Sapiens, Tomer (2015) tries to position an agent called the *smart person*, who differs from those at both ends. To do so, Tomer identifies the missing ingredients in both ends.

16 While chance-discovering is tied to the notion of random behavior, the idea and the process of novelty-discovery does not necessarily have to be random. Also, see Witt (2009).

17 By Simon (1962), the reason that one can harness complex systems is because they tend to be near decomposable and evolving hierarchical.

18 For a survey of these experiments, the interested reader is referred to Chen (2015), chapter 17.

19 See Wäckerle, Rengs and Radax (2014) for the role of different memory sizes on social trust and institutional change analyzed within an agent-based framework.

20 For a critical discussion on the cognitive ability of the ZI agents, see Tubaro (2009).

21 See Chen (2012) for a discussion on the relationships.

22 There are studies which argue that random behavior can be learned in the presence of feedback (Neuringer, 1986). However, in the standard version of ZI, agents do not learn.

23 About behavioral AI, Wooldridge (2009) made the following remarks:

> The workers in this area were not united by any common approaches, but certain themes did occur in this work. Recurring themes were the rejection of architectures based on symbolic representations, an emphasis on a closer coupling between the agent's environment and the action it performs, and the idea that *intelligent behavior can be seen to emerge from the interaction of a number of much simpler behaviors*.
>
> *(Ibid.: 395; italics added.)*

Bibliography

Adcock, A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B. & Gabrieli, J. (2006). Reward motivated learning: Mesolimbic activation precedes memory formation. Neuron, 50(3), 507–17.

Aha, D. (1997). Lazy Learning. Kluwer.

Aha, D., Kibler, D. Marc, K. (1991). Instance-based learning algorithms. Machine Learning, 6(1), 37–66.
Ashraf, Q. & Galor, O. (2013). The "Out of Africa" hypothesis, human genetic diversity, and comparative economic development. American Economic Review, 103(1), 1–46
Axelrod, R. (1984) The Evolution of Cooperation. New York: Basic Books.
Barto, A. & Mahadevan, S. (2003). Recent advances in hierarchical reinforcement learning. Discrete Event Systems Journal, 13, 41–77.
Beauchamp, J., Cesarini, D., Johannesson, M., van der Loos, M., Koellinger, P., Groenen, P., Fowler, H., Rosenquist, J., Thurik, A., & Christakis, N. (2011). Molecular genetics and economics. Journal of Economic Perspectives, 25(4), 57–82.
Beekman, M., Sword, G. & Simpson, S. (2010). Biological foundations of swarm intelligence. In Blum, C. & Merkle, D. (eds.), Swarm Intelligence: Introduction and Application. Springer, 3–41.
Benjamin, D., Chabris, C., Glaeser, E., Gudnason, V., Harris, T., Laibson, D., Launer, L., & Purcell, S. (2007). Genoeconomics. In: Weinstein, M., Vaupel, J. & Wachter, K. (eds.), Biosocial Surveys. National Academic Press, 303–35.
Betsch, T. & Haberstroh, S. (eds.) (2014). The Routines of Decision Making. Psychology Press.
Bodenhausen, G. (1990). Stereotypes as judgmental heuristics: Evidence of circadian variations in discrimination. Psychological Science, 1(5), 319–22.
Botvinick, M. (2012). Hierarchical reinforcement learning and decision making. Current Opinion in Neurobiology, 22(6), 956–62.
Boyer, C., Brorsen, W. & Zhang, T. (2014). Common-value auction versus posted-price selling: An agent-based model approach. Journal of Economic Interaction and Coordination, 9(1), 129–49.
Bush, R. R. & Mosteller, F. (1955). Stochastic Models for Learning. New York: JohnWiley & Sons.
Callaway, E. (2012). Economics and genetics meet in uneasy union. Nature, 490(7419), 154–5.
Camerer, C. & Ho, T.-H. (1999). Experienced-weighted attraction learning in normal form games. Econometrica, 67(4), 827–74.
Cappelletti, D., Guth, W. & Ploner, M. (2008). Being of two minds: An ultimatum experiment investigating affective processes. Jena Economic Research Papers, 2008-048.
Casari, M. (2004). Can genetic algorithms explain experimental anomalies? An application to common property resources. Computational Economics, 24, 257–75.
Casdagli, M. (1991). Chaos and deterministic versus stochastic nonlinear modeling. Journal of the Royal Statistical Society B, 54, 303–28.
Chan, N., LeBaron, B., Lo, A. & Poggio, T. (1999). Information dissemination and aggregation in asset markets with simple intelligent traders, Working paper, MIT.
Chaxel, A. (2015). How do stereotypes influence choice? Psychological Science, 26(5), 641–5.
Chen, S.-H. (ed.) (2002a). Evolutionary Computation in Economics and Finance. Physica-Verlag.
Chen, S.-H. (ed.) (2002b). Genetic Algorithms and Genetic Programming in Computational Finance. Kluwer.
Chen, S.-H. (2012). Varieties of agents in agent-based computational economics: A historical and an interdisciplinary perspective. Journal of Economic Dynamics and Control, 36(1), 1–25.
Chen, S.-H. (2013). Reasoning-based artificial agents in agent-based computational economics. In: Nakamatsu, K & Jain, L. (eds.), Handbook on Reasoning-based Intelligent Systems, World Scientific, 575–602.
Chen, S.-H. (2015). Agent-Based Computational Economics: How the Idea Originated and Where it is Going. Routledge.
Chen, S.-H. & Chie, B.-T. (2008). Lottery markets design, micro-structure, and macrobehavior: An ACE approach. Journal of Economic Behavior and Organization, 67(2), 463–80.
Chen, S.-H. & Hsieh, Y.-L. (2011). Reinforcement learning in experimental asset markets. Eastern Economic Journal, 37(1), 109–33.
Chen, S.-H., Tai, C.-C. & Wang, S.-G. (2010) Does cognitive capacity matter when learning using genetic programming in double auction markets? In: Di Tosto, G. & Parunak, V. (eds), Multi-Agent-Based Simulation X. Springer, 37–48.
Cleveland, W. (1979). Robust locally weighted regression and smoothing scatterplots. Journal of the American Statistical Association, 74(368), 829–36.
Cover, T. & Hart, P. (1967). Nearest neighbor pattern classification. IEEE Transactions on Information Theory, 13(1), 21–7.
Damasio, A. (1994). Descartes' Error: Emotion, Reason and the Human Brain. Grosset: Putnam.
Dorigo, M. & Stützle, T. (2010) Ant colony optimization: Overview and recent advances. In: Gendreau, M. & Potvin, J. (eds.), Handbook of Metaheuristics. Springer, 227–63.

Duffy, J. (2006). Agent-based models and human subject experiments. In: Tesfatsion, L. & Judd, K. (eds), Handbook of Computational Economics: Agent-based Computational Economics, Vol. 2., Oxford: Elsevier, 949–1011.

Duffy, J. & Engle-Warnick, J. (2002). Using symbolic regression to infer strategies from experimental data. In Chen, S. (ed.), Evolutionary Computation in Economics and Finance, Springer, 61–82.

Engle-Warnick, J. & Slonim, R. (2004). The evolution of strategies in a repeated trust game. Journal of Economic Behavior and Organization, 55(4), 553–73.

Engle-Warnick, J. & Slonim, R. (2006). Inferring repeated-game strategies from actions: Evidence from trust game experiments. Economic Theory, 28, 603–32.

Erev, I. & Roth, A. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. American Economic Review, 88, 848–81.

Fabre, E., Causse, M., Pesciarelli, F. & Cacciari, C. (2015). Sex and the money: How gender stereotypes modulate economic decision-making: An ERP study. Neuropsychologia, 75, 221–32.

Farmer, J. & Sidorowich, J. (1987). Predicting chaotic time series. Physical Review Letters, 59(8), 845–7.

Ferguson, T. (2014). Mathematical Statistics: A Decision Theoretic Approach (Vol. 1). Academic Press.

Frantz, R. & Leeson, R. (eds.) (2013). Hayek and Behavioral Economics. Palgrave Macmillan.

Garnier, S., Gautrais, J. & Theraulaz, G. (2007). The biological principles of swarm intelligence. Swarm Intelligence, 1(1), 3–31.

Gilboa, I. & Schmeidler, D. (1995). Case-based decision theory. Quarterly Journal of Economics, 110(3), 605–39.

Gilboa, I. & Schmeidler, D. (2001). A Theory of Case-Based Decisions. Cambridge University Press.

Gigerenzer, G. (2007). Gut Feelings: The Intelligence of the Unconscious. Penguin Books.

Gigerenzer, G. & Gaissmaier, W. (2011). Heuristic decision making. Annual Review of Psychology, 62, 451–82.

Gode, D. & Sunder, S. (1993). Allocative efficiency of markets with zero intelligence traders: Market as a partial substitute for individual rationality. Journal of Political Economy, 101, 119–37.

Grosan, C. & Abraham, A. (2006). Stigmergic optimization: Inspiration, technologies and perspectives. In Abraham, A., Gorsan, C. & Ramos, V. (eds.), Stigmergic Optimization. Springer, 1–24.

Hahn, U. & Warren, P. (2009). Perceptions of randomness: Why three heads are better than four. Psychological Review, 116(2), 454–61.

Hayek, F. (1945). The use of knowledge in society. American Economic Review, 35(4), 519–30.

Heinrich, T. & Schwardt, H. (2013) Institutional inertia and institutional change in an expanding normal-form game. Games, 4, 398–425.

Holland, J. & Miller, J. (1991). Artificial adaptive agents in economic theory. American Economic Review, 81(2), 365–70.

Kahneman, D. (2011). Thinking, Fast and Slow. Farrar, Straus and Giroux.

Kahneman, D. & Tversky, A. (1972). Subjective probability: A judgment of representativeness. Cognitive Psychology, 3, 430–54.

Kirman, A. (1991). Epidemics of opinion and speculative bubbles in financial markets. In: Taylor, M. (ed.), Money and Financial Markets. Blackwell, Cambridge, 354–68.

Kirman, A. (1993). Ants, rationality, and recruitment. Quarterly Journal of Economics, 108, 137–156.

Kohonen, T. (1995). Self-organizing Maps. Springer.

Koza, J. (1992). Genetic Programming: On the Programming of Computers by Means of Natural Selection. MIT Press.

Ladley, D. (2012). Zero intelligence in economics and finance. Knowledge Engineering Review, 27(2), 273–86.

Li, M. & Vitanyi, P. (2008). An Introduction to Kolmogorov Complexity and its Applications. Springer.

Linz, P. (2006). An Introduction to Formal Languages and Automata. 4th edition. Jones & Bartlett.

MacKillop, J. (2013). Integrating behavioral economics and behavioral genetics: Delayed reward discounting as an endophenotype for addictive disorders. Journal of the Experimental Analysis of Behavior, 99(1), 14–31.

MacQueen, J. (1967). Some methods for classification and analysis of multivariate observations. Proceedings of 5th Berkeley Symposium on Mathematical Statistics and Probability. University of California Press, 1, 281–97.

Madrian, B. (2014). Applying insights from behavioral economics to policy design. Annual Review of Economics, 6, 663–88.

Mathy, F. & Feldman, J. (2012). What's magic about magic numbers? Chunking and data compression in short-term memory. Cognition, 122, 346–62.

Miller, G. (1956). The magical number seven plus or minus two: Some limits on our capacity for processing information. Psychological Review, 63, 81–97.

Montague, R. (2006). Why Choose This Book? New York: Dutton.

Mormann, M., Koch, C. & Rangel, A. (2011). Consumers can make decisions in as little as a third of a second. Judgment and Decision Making, 6(6), 520–30.

Neuringer, A. (1986). Can people behave "randomly"?: The role of feedback. Journal of Experimental Psychology: General, 115(1), 62–75.

Newell, B. & Shanks, D. (2014). Unconscious influences on decision making: A critical review. Behavioral and Brain Sciences, 37(1), 1–19.

Nowak, M. (2006). Evolutionary Dynamics: Exploring the Equations of Life. Harvard University Press.

Pape, A. & Kurtz, K. (2013). Evaluating case-based decision theory: Predicting empirical patterns of human classification learning. Games and Economic Behavior, 82(1), 5265.

Pashler, H. (1998). The Psychology of Attention. Cambridge, MA: MIT Press.

Paula, J., Messias, J., Grutter, A., Bshary, R. & Soares, M. (2015). The role of serotonin in the modulation of cooperative behavior. Behavioral Ecology, arv039.

Quinlan, R. (1986). Induction of decision trees. Machine Learning, 1(1), 81–106.

Rieskamp, J. & Gigerenzer, G. (2002). Heuristics for social interactions: How to generate trust and fairness. Paper presented at a colloquium at the Workshop in Political Theory and Policy Analysis, Indiana University, Bloomington, on April 8, 2002.

Rosenfeld, A., Bareket, Z., Goldman, C., LeBlanc, D. & Tsimhoni, O. (2015). Learning drivers, behavior to improve adaptive cruise control. Journal of Intelligent Transportation Systems, 19(1), 18–31.

Roth, A. & Erev, I. (1995). Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. Games and Economic Behavior, Special Issue: Nobel Symposium, 8, 164–212.

Sargent, T. (1993). Bounded Rationality in Macroeconomics. Oxford.

Savage, L. (1972). The Foundations of Statistics. Courier Corporation.

Scott, J. (2012). Social Network Analysis. 3rd edition, Sage Publications, London.

Sen, A. (1997). Maximization and the act of choice. Econometrica, 65(4), 745–79.

Simon, H. (1962). The architecture of complexity. Proceedings of the American Philosophical Society, 106(6), 467–82.

Snook, B., Zito, M., Bennell, C. & Taylor, P. J. (2005). On the complexity and accuracy of geographic profiling strategies. Journal of Quantitative Criminology, 21, 1–26.

Stone, C. (1977). Consistent nonparametric regression. Annals of Statistics, 5(4), 595–620.

Tagiew, R. (2012). Mining determinism in human strategic behavior. In: Tagiew, R., Ignatov, D., Neznanov, A. & Poelmans, J. (eds.), Experimental Economics in Machine Learning. Workshop co-located with the 10th International Conference on Formal Concept Analysis, May 2012, Leuven, Belgium. 84–91.

Takens, F. (1981). Detecting strange attractors in turbulence. In: Rand, D. & Young, L.-S. (eds.), Dynamical Systems and Turbulence. Springer-Verlag, 366–81.

Thaler, R. (2000). From Homo economicus to Homo sapiens. Journal of Economic Perspectives, 14(1), 133–141.

Thaler, R. & Sunstein, C. (2008). Nudge: Improving Decisions about Health, Wealth, and Happiness. Penguin Books.

Tomer, J. (2015). Smart persons and human development: The missing ingredient in behavioral economics. In: Altman M. (ed.), Behavioral Economics with Smart People, Edward Elgar.

Tubaro, P. (2009). Is individual rationality essential to market price formation? The contribution of zero-intelligence agent trading models. Journal of Economic Methodology, 16(1), 1–19.

Valentini, G. & Hamann, H. (2015). Time-variant feedback processes in collective decision making systems: Influence and effect of dynamic neighborhood sizes. Swarm Intelligence, 9(2–3), 153–76.

von Neumann, J. & Morgenstern, O. (1944). Theory of Games and Economic Behavior. Princeton University Press, Princeton.

Vriend, N. (2002). Was Hayek an ACE? Southern Economic Journal, 68(4), 811–40.

Wäckerle, M., Rengs, B. & Radax, W. (2014). An agent-based model of institutional life-cycles. Games, 5, 160–87.

Wagenaar, W. (1972). Generation of random sequences by human subjects: A critical survey of literature. Psychological Bulletin, 77(1), 65–72.
Wansink, B. & Sobal, J. (2007). Mindless eating: The 200 daily food decisions we overlook. Environment & Behavior, 39(1), 106–23.
Witt, U. (2009). Propositions about novelty. Journal of Economic Behavior & Organisation, 70, 311–20.
Wooldridge, M. (2009). An Introduction to Multi-Agent Systems, 2nd Edition. Wiley.
Zhao, J., Hahn, U. & Osherson, D. (2014). Perception and identification of random events. Journal of Experimental Psychology: Human Perception and Performance, 40(4), 1358–71.