ISSN 0819-2642
ISBN 978 0 7340 3722 0

THE UNIVERSITY OF MELBOURNE

DEPARTMENT OF ECONOMICS

RESEARCH PAPER NUMBER 1012

October 2007

Behavioural Anomalies, Bounded Rationality and Simple Heuristics

by

Suren Basov, Liam Blanckenberg
& Lata Gangadharan

Department of Economics
The University of Melbourne
Melbourne Victoria 3010
Australia.

Behavioural Anomalies, Bounded Rationality and Simple Heuristics

October, 2007

Suren Basov*

s.basov@latrobe.edu.au

La Trobe University, Bundoora, Victoria 3086, Australia

Liam Blanckenberg

bll@unimelb.edu.au

University of Melbourne, Parkville, Victoria 3010, Australia

Lata Gangadharan

latag@unimelb.edu.au

University of Melbourne, Parkville, Victoria 3010, Australia

Abstract: The use of bounded rationality in explaining economic phenomena has attracted growing attention. In spite of this, there is still considerable disagreement regarding the meaning of bounded rationality. Basov (2005) argues that when modeling boundedly rational behaviour it is desirable to start with an explicit formulation of the learning process. A complete understanding of the boundedly rational decision-making process requires development of an evolutionary-dynamic model which can give rise to such learning processes. Evolutionary dynamics implies that

*Author for correspondence

individuals use heuristics to adjust their choices in light of past experiences, moving in the direction that appears most beneficial, where these adjustment rules are assumed 'hardwired' into human cognition through the process of biological evolution. In this paper we elaborate on the latter point by building a model of evolutionary selection relevant to heuristics. We show that in addition to explaining the origin of learning rules this approach also sheds light on some well documented preference anomalies.

Keywords: Bounded Rationality, Heuristics, Replicator Dynamics

JEL Classification: C0, D7

1 Introduction

Recently there has been a growing acceptance of the use of bounded rationality models in explaining economic phenomena. For example Offerman, Schram and Sonnemans (1998) use a model of quantal response equilibrium to explain step-by-step provision of public goods, while Anderson, Goeree and Holt (1998, 2001) use similar models to explain behaviour in all-pay auctions and coordination games. On the subject of social learning Kandori, Mailath and Rob (1993) and Young (1993) use dynamic models based on gradual adjustment to explain the evolution of conventions.

Despite this growing reliance on models of boundedly rational behaviour there is still considerable disagreement regarding the precise meaning of the term. Most applied papers model bounded rationality as probabilistic choice using Luce's (1959) model and its extensions. Basov (2005) criticised this approach and argued that when modeling boundedly rational behaviour it is desirable to start with an explicit formulation of the learning process. This, however, leaves open the question: What determines the learning rule? A complete model of bounded rationality should endogenise the entire decision-making process. Basov (2006) suggests using the evolutionary-dynamic approach which states that individuals adjust their choices in light of past experiences and move in the direction that appears most beneficial. This approach therefore assumes that the adjustment rules are themselves hardwired into us through the process of biological evolution. In this paper we focus on the latter point by building a model of evolutionary selection of heuristics that gives rise to adjustment rules. In doing so we demonstrate that in addition to explaining the origin of such rules this approach also sheds light on some well documented preference anomalies.

The literature documenting preference anomalies has grown considerably over the last 30 years. We focus briefly on literature relevant to (i) the endowment effect and (ii) behaviour in public good games. For a broad collection and discussion of behavioural anomalies see Thaler (1992).

The commonly observed *endowment effect*, a pattern of behaviour first formalised by Thaler (1980), is the tendency for individuals to value a commodity (a good, a bundle of goods or a bundle of lotteries) that they *own* more than an identical commodity they could *obtain* through a transaction. The endowment effect is one of several suggested explanations behind the observed disparity between individuals' 'willingness to pay' (WTP) and 'willingness to accept' (WTA) measures of value (see Knetsch & Sinden, 1984). This anomaly is most prevalent in contingent valuation settings, with the robust result being that an individual's stated WTA is frequently considerably higher than their WTP for a commodity of identical value. For a comprehensive collection of empirical findings and associated discussion see Bateman, Kahneman, Rhodes, Starmer and Sugden (2005), and Horowitz and McConnell (2002).

It is theoretically expected that the WTA and WTP measures of value be almost equivalent (barring unrealistically large income effects). Despite controlling for a range of possible causes including income effects (Brookshire & Coursey, 1987); learning effects (Coursey, Hovis & Schulze, 1987); the effect of incentives (Kahneman, Knetsch & Thaler, 1990); and the use of Vickery auctions (Coursey *et al*, 1987), amongst others, the observed disparity persists.

Another persistent pattern of anomalous behaviour is observed in public good experiments. In such experiments a group of k people are endowed with s dollars each. They can contribute any amount between 0 and s into a common pool, the sum total of which is then multiplied by some factor $n < k$ and distributed equally among all participants. It is clear that the dominant strategy is to contribute nothing while an efficient outcome involves contributing everything. Subjects tend to invest $40 - 60\%$ of their endowment in the common pool. If subjects play this game repeatedly the contribution rates tend to fall but nevertheless remain positive. For a detailed review of such experiments see Thaler (1992).

This paper is organised as follows: In Section 2 we provide some examples of simple rule-based heuristics, while Section 3 demonstrates how two such heuristics can explain several behavioural anomalies. In Section 4 we describe a general model of heuristic evolution, while in Section 5 we assume environmental complexity changes but at a rate slower than the characteristic time scale of heuristic evolution, leading to the emergence of more sophisticated agents. Section 6 concludes with a summary of results and a focus for further research.

2 Simple Heuristics

Our contention is that explaining certain behavioural anomalies is possible provided the idea of case-by-case optimisation is relaxed and instead one assumes that individuals use simple rules of thumb, or *heuristics,* which are hardwired into our instinctive cognition through the process of biological evolution. To clarify what is meant by simple heuristics we provide the following examples:

Caution Heuristic Stick with the status quo when faced with a risky or unfamiliar decision, irrespective of the expected outcome of the risky proposition. For example, if someone offers to exchange your car for a much better looking example for a modest sum (say only $100) you might instinctively decline the offer and act cautiously.

Recognition Heuristic Start with ordering the alternatives and choose the first alternative you recognise. For example, assume an investor orders stocks in some random manner (i.e. alphabetically or by company announcements) and invests in the first stock they recognise.[1] The probability an investor will invest in asset i is determined by the ordering selected and by previous experience. If profitable firms advertise more than less profitable ones, and the ordering is random, then the probability of investment in a given stock will be increasing in its returns. Conversely, if larger firms advertise more than smaller ones then the probability of investing in a given stock will be increasing in the size of the firm and not necessarily in the return on its stock.[2] One can demonstrate that a portfolio constructed according to this principle can perform quite well (Gigerenzer & Selten, 2001). Moreover, this

[1] The recognition heuristics is thoroughly discussed by Gigerenzer and Selten (2001).

[2] We use the recognition heuristic purely as an illustrative example and so ignore the possibility of strategic interaction by firms in this environment.

heuristic is shown to have evolutionary roots, in that rats tend to select foods the smell of which they recognise over foods which they don't, presumably to minimise the chance of being poisoned (Gigerenzer & Selten, 2001).

Select the Best Assume that each individual is endowed with a set of 'keys', ordered in accordance to their ecological validity.[3] In the previous stock market example, assume that an investor is endowed with two keys: (i) whether the company name appears in the Fortune 500 and (ii) whether their friend invested in the stock of the same company.[4] Given the Fortune 500's construction as the 500 largest American corporations, our individual will abandon their first key in favour of their second key assuming more than one company is listed. Relying on the second key implies the investor consults a friend regarding their investment. If the friend invested in a Fortune 500 company our individual does the same, if not they randomly choose one of the 500 companies identified using their first key. As in the previous example this process will lead to some probability distribution over the investment in various stocks.

Note that the first three heuristics are *static* and can be used to explain initial behaviour. The following two heuristics are *dynamic* and can be used to explain behaviour across repetition.[5]

Tit-for-Tat Suppose you are confronted with a situation where repeated cooperation with another individual is required to achieve a common goal and where cooperation is personally costly (thus refusing to do so is privately optimal). In such cases you would find yourself in a standard prisoner's dilemma situation. The

[3] Ecological validity of a key is the fraction of the correct decisions made based on the value of the key.

[4] For simplicity assume the friend invests solely in one company.

[5] Thus these heuristics can give rise to social learning.

tit-for-tat strategy calls for cooperation in period one and in subsequent periods if, and only if, your partner cooperates in each period before you.

Generalised Tit-for-Tat Suppose you are confronted with a situation resembling a standard public good game (i.e. being asked to contribute to a common pool from which all players will receive an even share). The generalised tit-for-tat strategy calls for starting with a high contribution and increasing (decreasing) contributions in subsequent periods provided your contribution is below (above) the mean group contribution.

A commonality amongst the above mentioned heuristics is their optimality in very specific cases. For example, the behaviour suggested by the *Caution Heuristic* is justified provided there is a high chance of a 'lemon offer' and information asymmetries are present, while *Tit-for-Tat* is the Nash equilibrium strategy in an infinitely repeated prisoner's dilemma provided players have sufficient patience. The central idea behind the heuristic approach is that once behaviour specified by a heuristic becomes hardwired into cognition, individuals follow such rules in all situations, even when such a strategy leads to sub-optimal payoffs. In other words, the heuristic approach does not assume individuals possess any form of 'cognitive override' which might correct such erroneous behaviour. If we evolved in a world where the chance of receiving a lemon trade is high we will continue to decline risky proposals, even in situations where the chance of getting a lemon trade is negligible, due to pessimism or having 'learnt the hard way'. If *Tit-for-Tat* (or indeed *Generalised Tit-for-Tat*) is a good strategy in some social interactions, we may continue to follow such a strategy in other games with different structures, possibly relying on an obsolete or sub-optimal rule.

Observe that all five heuristics presented above are simple, i.e. they can be summarised by a 'rule of thumb' and as such are assumed to incur zero complexity costs. Conversely one might consider a more sophisticated heuristic. For example, when faced with a choice under uncertainty one may use the *Caution Heuristic* if there are three or more possible states of the world but compute the expected utility from each state if faced with a binary choice. Such a heuristic is no longer rule based, thus increasing its complexity, and as such an agent should incur some cost to develop the relevant computational abilities to make such decisions.

A reasonable question to ask is: What determines the repertoire of heuristics available to a given agent? Our answer: Evolution. To formalise this idea we abandon the notion that individuals maximise on a case-by-case basis. Instead, we specify a set of problems, Ω, an agent might face and assume there is a finite set of heuristics, H, available to that agent. Agents are assumed to select a heuristic $h_j \in H$ and apply it to all problems in Ω. This generates a payoff $\pi(h_j, x)$ where $x = (x_1, ..., x_H)$[6] and x_i is the fraction of the population that uses heuristic h_i. We assume that the evolution of x_i is described by the replicator dynamics and that the distribution of heuristics in a given population corresponds to asymptotically stable steady states of such dynamics.

3 Motivating Examples

The notion of a heuristics approach to decision-making is appealing due to its intuitive foundations, as noted by Rosenthal (1992). It is difficult to validate the conviction that individuals possess the analytical skills to construct a model, assess

[6] With the standard abuse of notation we use H for both the set of heuristics and its cardinality.

the relevant probabilities and finally maximise their utility criteria on a routine basis. A far more realistic scenario is where an individual, faced with an unfamiliar and uncertain decision, instinctively invokes a simple 'rule of thumb'. Such a rule would represent an accumulated pool of knowledge from similar past decisions, as well as an instinctive assessment of the environment currently being confronted.

Example 1 *Caution Heuristic* Consider a large but finite population of risk-neutral[7] individuals faced with a family of problems, P. Each problem $p \in P$ can be represented as a choice between two actions: safe and risky. The safe choice results in a certain payoff π_H while the risky choice is represented as a binary lottery:

$$\pi^l \text{ with probability } \gamma$$
$$\pi^h \text{ with probability } 1-\gamma$$

where $0 \leq \gamma \leq 1$ and $\pi^l \neq \pi^h$. The payoffs and probabilities vary from problem to problem and we assume that there are problems $p_1, p_2 \in P$ such that

$$\gamma(p_1)\pi^l(p_1) + (1-\gamma(p_1))\pi^h(p_1) > \pi_H(p_1) \quad (1)$$
$$\gamma(p_2)\pi^l(p_2) + (1-\gamma(p_2))\pi^h(p_2) < \pi_H(p_2) \quad (2)$$

i.e. for some problems the safe option is the optimal choice and for others the risk is worth taking. Furthermore we assume that making *always* the safe choice is better than making *always* the risky choice.

Suppose that there are two types of individuals; naifs and sophisticates.[8] Naifs

[7] Risk neutrality is assumed purely for simplicity. A similar analysis could be conducted incorporating risk aversion.

[8] We borrow this terminology from O'Donoghue and Rabin (1999) but adopt a literal interpretation in that sophisticates are more congnitively advanced than naifs.

must restrict themselves to the same choice for all problems. Given our assumptions on the payoffs, they will eventually learn to rely on the *Caution Heuristic*, i.e. will always take the safe option. The sophisticates, on the other hand, solve this problem using standard expected utility theory (EUT).[9] We assume that in order to develop mental capacities necessary to apply EUT individuals incur some cost, τ (such a cost could be biological in nature, i.e. extra metabolic requirements necessary to support advanced cognition, or simply the opportunity cost of ones time taken to learn the new technique). Once the necessary mental capacity is developed we assume sophisticates can apply this method costlessly in future situations. It is straightforward to show that in the long-run the population will develop sophistication provided

$$\tau > E(\pi_{EUT}) - E(\pi_H) \tag{3}$$

and rely on the *Caution Heuristic* otherwise. Expectations in (3) are taken over all problems in P. This model can be enriched to allow for the expected payoffs of a particular strategy to depend on the fraction of the population which already utilise that strategy. In such cases the long-run result could be a non-trivial mix of naifs and sophisticates within a given population.

The above example can be used to explain the endowment effect. We define the endowment effect as the tendency for individuals to value a commodity they *own* more than an identical one they could *obtain* through a market transaction. In a typical experiment to demonstrate the endowment effect mugs bought at a university shop are distributed randomly among participants. Those who received a mug are

[9] One should not assume that these agents possess enough sophistication to solve *any* problem using EUT - They just possess enough sophistication to apply EUT to binary choices.

asked about the minimal price at which they would be willing to part with the mug (WTA), while those who did not receive a mug are asked about the maximal price at which they would be willing to buy the mug (WTP). A consistent experimental finding is that $WTA > WTP$. To explain this result under assumptions of perfect rationality one must either assume unreasonably large income effects or postulate a direct preference for endowment.

We argue, however, that neither is necessary if one gives up the idea that choices should be rationalised on a case-by-case basis. One should rather imbed the mug problem into a class of problems in which sticking with your endowment is the sensible thing to do, on average, by invoking the *Caution Heuristic*. Such a class of decision problems might spring to mind if mugs are replaced by used cars. It is easy to justify a higher asking price by an individual selling a car versus a lower bidding price by an individual looking to purchase the same car since this is exactly the situation on the 'lemon market' for used cars (see Akerlof, 1970).

Example 2 *Generalised Tit-for-Tat* Consider a standard public good game where k subjects are given s dollars each and where they must decide on what amount $x \in [0, s]$ to contribute to a common pool. The contributions to the common pool are multiplied by some constant $n < k$ and distributed equally between all participants. It is clear that the dominant strategy in this game is to contribute nothing since for each dollar you contribute you get only $n/k < 1$ dollars back. Numerous experiments (see Thaler, 1992), however, report that individuals start with sufficiently high contributions which tend to decrease over time.

Again, under assumptions of perfect rationality, it is hard to explain such behaviour. However, the explanation is rather straightforward if we assume that individuals imbed such situations into a broad class of cooperation games and apply or adapt some rules that work well in typical cooperation environments. Many cooperation games can be modelled as a repeated prisoner's dilemma. An extremely successful strategy (in evolutionary terms) in this setting is *Tit-for-Tat*, which calls for cooperation in period one and in subsequent periods if, and only if, your partner cooperates each period before you (see Axelrod, 1984). One can adapt such a strategy for the above game to imply; start with a high contribution and increase (decrease) your choice in subsequent periods if your contribution is below (above) the population mean. Formally, the contributions of agent i who follows the generalised tit-for-tat strategy will evolve according to:

$$x_{it+1} = x_{it} + \gamma(z_t - x_{it}) \tag{4}$$

where z_t is the mean population contribution, x_{it} is the contribution of agent i at time t and $\gamma \in (0, 1)$ is some constant. If the entire population follows rule (4) then the mean population contribution will be given by:

$$z_t = \frac{1}{N} \sum_{k=1}^{N} x_{it} \tag{5}$$

where N is the number of agents. It is straightforward to see that the unique solution of (4) for given initial contributions is:

$$z_t = z_0, \;\; x_{it} = z_0 + (1-\gamma)^t (x_{i0} - z_0), \tag{6}$$

where x_{i0} and z_0 are initial contributions of agent i and the initial population mean respectively. Therefore, if everyone in the population follows the generalised tit-for-tat strategy, the average contribution remains constant while individual contributions revert to the population mean. Now assume that there is a fraction $\rho \in (0, 1)$ of sophisticates who 'figure out' the game and contribute nothing in each round. Then equation (4) will describe the evolution of contributions of the naifs and equation (5) should be modified to read

$$z_t = \frac{1-\rho}{N} \sum_{k=1}^{N} x_{it} \tag{7}$$

One can solve for the mean population contribution of the naifs in this case to obtain

$$z_t = (1-\rho)z_0 + z_0\rho(1-\gamma)^t \tag{8}$$

Note that the average contributions deteriorate over time but remain positive, which is consistent with experimental evidence. The individual contributions now follow:

$$x_{it+1} = (1-\gamma)x_{i0} + \sum_{k=0}^{t} (1-\gamma)^k z_{t-k} \tag{9}$$

The behaviour of naifs in this example is hard to explain looking at the payoffs of any particular game in isolation, however such behaviour makes perfect sense if viewed as a simple response to a generalised cooperation environment.

4 The General Model

In this section we present a general model of heuristics evolution. Consider a population of individuals who are repeatedly faced with a set of problems, P, to solve. They have a finite set of H heuristics available to use. Heuristic $h \in H$ applied to a problem $p \in P$ produce a solution $s(h, p) \in S_p$, where S_p denotes the set of all feasible solutions to problem p. The expected payoff this solution generates for an individual is $\pi(s, x)$, where $x = (x_1, ..., x_H)$ and x_h is the fraction of the population using heuristic h. We will endow P with a structure of measure space (P, Σ, μ), where Σ is a sigma-algebra of subsets of P, μ is a probability measure on Σ, and for $\forall B \in \Sigma$ we assume with probability $\mu(B)$ an individual encounters problem $p \in B$. The expected payoff to an individual from using heuristic H is then:

$$U(h, x) = \int_P \pi(s(h, p), x) d\mu(p) - c(h), \tag{10}$$

where $c(h)$ is the complexity cost associated with heuristic h. The average payoff of an individual in the population is given by:

$$U(x) = \sum_{h=1}^{H} x_h U(h, x) \tag{11}$$

We assume that the rate of change of the fraction of a population that follows heuristic $h \in H$ is proportional to the difference between the expected payoff an individual receives using heuristic h and the mean population payoff, i.e.

$$\frac{dx_h}{dt} = \eta x_h (U(h, x) - U(h)), \tag{12}$$

where $\eta > 0$.

Equation (12) has the same form as the replicator equation in evolutionary game theory.[10] Interpretation in this case is, however, different. While in evolutionary game theory one usually restricts attention to a particular game and studies the

[10] There is, however, a slight formal difference since we do not require $U(r, x)$ to be polylinear in x. Complexity costs present an obvious reason to violate such a condition.

evolution of strategies, in this case we are interested in the evolution of rules, which can be applied in a variety of different situations (i.e. different games). For example, the *Caution Heuristic* can be applied in the mug experiment and in the market for used cars while *Generalised Tit-for-Tat* can be used in a repeated prisoner's dilemma and in public good games.

It is immediate from equations (10), (11) and (12) that

$$\frac{d}{dt}\sum_{h=1}^{H} x_h = 0 \tag{13}$$

Therefore, if fractions of a population playing different strategies initially sum to one, they will continue to do so at any other time. Note also that $x_h = 0$ implies $dx_h/dt = 0$, i.e. if $x_h(0) > 0$ it will remain non-negative at all times. Consequentially system (12) confines the dynamics of x to the unit simplex.

Our next objective is to analyse the steady states of system (12) and their associated stability. Let us formally consider a non-cooperative game $G = (\{S_\alpha\}_{\alpha \in C}, \{u_\alpha\}_{\alpha \in C}, C)$ in which each player has a strategy set $S_\alpha = H$. Players have Bernoulli utility $u_\alpha : \Delta(H) \times H \to R$ where C is the set of all players and where

$$u_\alpha(h, x) = U(h, x) \tag{14}$$

A non-standard feature of this formulation is that the payoff to a player who adopts a particular strategy depends on population averages rather than on the strategy profile of opponents. For this reason it is proper to think of the interaction between players as being conducted jointly at a population-wide level in a way that does not lend itself to modelling with interaction among a randomly selected group of players.[11] A crucial observation, however, is that the von Neumann-Morgenstern utility of a player who adopts a mixed strategy y against a population with distribution of strategies x is

$$U(x,y) = \sum_{h=1}^{H} y_h U(h,x) \tag{15}$$

where $U(x,y)$ is linear in y.[12] It is now straightforward to define a Nash equilibrium as a strategy profile y^* such that

$$U(y^*,y^*) \geq U(y^*,y) \tag{16}$$

for all $y \in \Delta(H)$. The standard results of evolutionary game theory (see Weibull, 1995) can now be applied.

Result 1 Any Lyapunov stable state of system (12) is a Nash equilibrium of G.

In our case this implies that if a mix of heuristics used by the population is stable, it is impossible to outperform an average member of the population using heuristics from the same repertoire. Note, however, that this result does not imply that the heuristics repertoire that emerges in equilibrium will allow one to obtain the maximal possible payoff for any problem in P. Instead it simply implies that it is impossible to get systematically higher payoffs using heuristics from set H.

Result 2 Any strict dominant strategy is a globally attractive steady state.

Result 3 Any interior Nash equilibrium (if such exists) is reachable, i.e. there exists an initial steady state from which the system will eventually converge to. This

[11] For other interesting applications of population-wide models see Hansen and Kaarbøe (2002).

[12] We implicitly assume that C has cardinality of continuum and therefore the strategy choice of one player does not disturb the population mean.

means that the issue of equilibrium selection, familiar to game theorists, arises here.[13] Such a problem of equilibrium selection is relevant only if one insists that prediction of a theory should correspond to an equilibrium point and that dynamics are at best an aid in selecting such a point. We, on the contrary, take dynamics seriously as describing the evolution of human behaviour over time. In this case, multiplicity of equilibria can be an asset rather than liability, since it allows us to study how chosen heuristics were shaped by real evolutionary processes.

5 Changing Environments and the Evolution of Complexity

Let us now assume that time passes in discreet periods and that during each period the set of problems is P_n where $n \in N$. We will assume that $P_n \subset P_{n+1}$ and interpret problems $P_{n+1} \backslash P_n$ as more complex than those in P_n (for example, the evolution of society and technology brings forth new challenges). We also assume that the sets of heuristics $\{H_k\}_{k=1}^{p}$ are ordered in such a way that all heuristics in the same set H_k have the same complexity cost c_k and that $c_{k+1} > c_k$. Moreover, we assume that for $\forall n \in N \ \exists k(n) \in N : \exists h \in H_k(n)$ such that

$$\pi(s(h,p),x) \geq \pi(s^*,x) \text{ for } \forall s^* \in S_p, \forall p \in P_n \tag{17}$$

In words, this means that for any problem set P_n there is a sophisticated enough heuristic which can solve all problems in P_n optimally, net of complexity costs. It is

[13] Recent literature has shown that adding noise does not solve this problem, since the exact equilibrium selected depends on the way the noise is added. For population-wise games this was first demonstrated by Hansen and Kaarbøe (2002).

natural to assume that the heuristics sets $k(n)$ are increasing in n and that individuals acquire heuristics gradually, i.e. they possess repertoires of heuristics of the form:

$$H^m = \bigcup_{k=1}^{m} H_k \tag{18}$$

This amounts to saying that environments of higher order n are more complex. It is important to note that we do not assume that every problem in P_{n+1} is necessarily more complex than any problem in P_n. We only assume that a universal algorithm to find the optimal solution for *any* problem in P_{n+1} is more complex than a universal algorithm to find the optimal solution for *any* problem in P_n. Therefore, it is quite possible that most individuals are able to find the optimal solution to a complex problem in P_n (e.g. to manage running a marathon), but are unable to solve a simple problem in P_{n+1} (e.g. how much to save for retirement given time-preferences and market interest rates). Finally, we assume that each period is sufficiently long such that we can concentrate on the long-run dynamics of system (12).

Definition 1 *Given environment P_m, the repertoire of heuristics H^m is evolutionary stable if*

$$\max_{h \in H^m} \int_P \pi(s(h,p),x)d\mu(p) - c_m \geq$$

$$\max\Big(\max_{h \in H^{m-1}} \int_P \pi(s(h,p),x)d\mu(p) - c_{m-1}, \max_{h \in H^{m+1}} \int_P \pi(s(h,p),x)d\mu(p) - c_{m+1}\Big) \tag{19}$$

It is strictly stable if the above inequality is strict.

Before proceeding further let us consider a simple example.

Example 3 Suppose environment P is such that individuals are called repeatedly to choose one of three outcomes. The utilities of each outcome are independently distributed according to the uniform distribution on $[0, 1]$. Three available heuristics sets H_0, H_1, H_2 are all singletons. Set $H_0 = \{h_0\}$ is "choose the first option you encounter." The evolutionary cost associated with this set is 0. If we assume that the order in which the options are encountered is uncorrelated with the payoffs then the expected utility attained by individuals who use this heuristic is

$$E(u) = \frac{1}{3}(E(u_1) + E(u_2) + E(u_3)) = \frac{1}{2} \tag{20}$$

Set $H_1 = \{h_1\}$ is "choose a pair of options at random, compare the utilities and choose the best" (this heuristic is sophisticated enough to compare a pair of options but not sophisticated enough to store the results). The evolutionary cost associated with this set is c_1 and the expected utility attained by individuals who use this heuristic is

$$E(u) = \frac{1}{3}(E\max(u_1, u_2) + E\max(u_2, u_3) + E\max(u_1, u_3)) = \frac{2}{3} \tag{21}$$

Finally, set $H_2 = \{h_2\}$ is "choose a pair of options at random, compare the utilities, remember the best, compare with the remaining options, and choose the best," (this heuristic allows for short term memory and is sophisticated enough to always lead to the optimal choice in environment P).[14] The evolutionary cost associated with this set is c_2 and the expected utility attained by the individuals who use this heuristic is

$$E(u) = E\max(u_1, u_2, u_3) = \frac{3}{4} \tag{22}$$

[14] It is wrong, however, to identify agents that use such a heuristic as 'utility maximising'. This is because heuristic H_2 will fail to maximise utility in a choice between four or more alternatives.

If $c_1 \geq \frac{1}{6}$ then the evolutionary stable set is H^0, if $c_1 \leq \frac{1}{6}$ and $c_2 \geq \frac{1}{12}$ then the evolutionary stable set is H^1 and if $c_2 \leq \frac{1}{12}$ (which implies $c_1 < \frac{1}{12}$) then the evolutionary stable set is H^2.

Let us return to the general characterisation of evolutionary stable steady states.

Proposition 1 *Assume that $c_{k+1} > c_k$, so that for any $n \in N$*

$$\max_{h \in H^m} \int_{P_n} \pi(s(h,p),x) d\mu(p) - \max_{h \in H^{m-1}} \int_{P_n} \pi(s(h,p),x) d\mu(p) \tag{23}$$

is decreasing in m, and that their exists a $q \in N$ such that

$$\max_{h \in H^{q+1}} \int_{P_n} \pi(s(h,p),x) d\mu(p) - \max_{h \in H^q} \int_{P_n} \pi(s(h,p),x) d\mu(p) < 0 \tag{24}$$

Then for any n there are only finitely many evolutionary stable steady states.

Proof *Evolutionary stability implies*

$$\max_{h \in H^m} \int_{P_n} \pi(s(h,p),x) d\mu(p) - \max_{h \in H^{m-1}} \int_{P_n} \pi(s(h,p),x) d\mu(p) \geq c_m - c_{m-1} > 0 \tag{25}$$

Conditions (23) - (24) imply that there are only finitely many solutions to (25) ■

Note that if multiple steady states exist then those with the lowest complexity are more likely to be observed since they require a smaller number of simultaneous mutations. Our assumption regarding the gradual acquisition of heuristics repertoires implies that the evolution of set H_k to H_{k+n} requires n simultaneous mutations. The implication of this result is that a population can be trapped in a state with relatively low complexity even though a higher degree of complexity might have been beneficial.

6 Conclusions

Empirical work using experiments has uncovered numerous behavioural anomalies; two such examples being the endowment effect and public good contributions. An attempt to understand such anomalies on the basis of case-by-case optimisation inevitably leads to the postulation of non-standard preferences (e.g. preferences being expressed in changes rather than levels or kinked utility functions at the point of endowment). We argue that such modifications are unnecessary if one abandons the assumption of case-by-case optimisation and instead embraces *holistic* optimisation.

The holistic optimisation approach proposes that individuals apply simple heuristics to a wide variety of decision problems. Solutions generated by such rules need not be optimal on a case-specific basis, however they should on average dominate the solutions generated by alternative and feasible rules. The appropriate criterion for evaluation of such rules is evolutionary fitness. Given reasonable assumptions, such an environment can be represented under the expected utility framework (see Robson, 2001), net of complexity costs. Under a rule-based system of problem-solving (governed by evolutionary dynamics) no ad hoc modifications to preferences, such as those advocated by many behavioural economists, are necessary.

The ideas set forward in this paper can be further developed in two directions. Firstly, one can study more carefully real evolutionary environments in which humans have developed to shed light on the actual set of heuristics commonly used in economic decision-making. Secondly, one can use reduced-form models of human behaviour that arise from this heuristics approach to rationalise economic phenomena. As argued in Basov (2005, 2006) such reduced-form models should assume that an individual's choices follow a Markov process over some state space.

References

[1] Akerlof, G.A. (1970). The Market for "Lemons": Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, 84(3), 488-500

[2] Anderson, S., Goeree, J. and Holt, C. (1998). Rent Seeking with Bounded Rationality: An Analysis of the All-Pay Auction. *Journal of Political Economy*, 106(4), 828-853

[3] Anderson, S., Goeree, J. and Holt, C. (2001). Minimum Effort Coordination Games: Stochastic Potential and the Logit Equilibrium. *Games and Economic Behaviour*, 34(2), 177-199

[4] Axelrod, R. (1984). *The Evolution of Cooperation,* New York, USA: Basic Books

[5] Bateman, I., Kahneman, D., Munro, A., Starmer, C. and Sugden, R. (2005). Testing Competing Models of Loss Aversion: An Adversarial Collaboration. *Journal of Public Economics*, 89(8), 1561-1580

[6] Basov, S. (2005). Bounded Rationality: Static versus Dynamic Approaches. *Economic Theory*, 25(4), 871-885

[7] Basov, S. (2006). Mechanics of Choice. *Progress in Economics Research*, 10(1), 55-64. Nova Publishers, NY, Editor: Albert Tavidze

[8] Brookshire, D.S. and Coursey, D.L. (1987). Measuring the Value of a Public Good: An Empirical Comparison of Elicitation Procedures. *American Economic Review*, 77(4), 554-566

[9] Coursey, D.L., Hovis, J.L. and Schulze, W.D. (1987). The Disparity Between Willingness to Accept and Willingness to Pay Measures of Value. *Quarterly Journal of Economics*, 102(3), 679-690

[10] Gigerenzer, G. and Selten, R. (2001). *Bounded Rationality: The Adaptive Toolbox*, Rev. Ed., Cambridge, USA: MIT Press

[11] Hansen, S. and Kaarbøe, O. (2002). Equilibrium Selection in Coordination Games with Simultaneous Play. *Economic Theory*, 20(4), 793-807

[12] Horowitz, J.K. and McConnell, K.E. (2002). A Review of WTA/WTP Studies. *Journal of Environmental Economics and Management*, 44(3), 426-447

[13] Kahneman, D., Knetsch, J.L. and Thaler, R.H. (1990). Experimental Tests of the Endowment Effect and the Coase Theorem. *Journal of Political Economy*, 98(6), 1325-1348

[14] Kandori, M., Mailath, G. and Rob, R. (1993). Learning, Mutation, and Long Run Equilibria in Games. *Econometrica*, 61(1), 29-56

[15] Knetsch, J.L. and Sinden, J.A. (1984). Willingness to Pay and Compensation Demanded: Experimental Evidence of an Unexpected Disparity in Measures of Value. *Quarterly Journal of Economics*, 99(3), 507-521

[16] Luce, R.D. (1959). *Individual Choice Behaviour: A Theoretical Analysis*, Rev. Ed., New York, USA: Wiley

[17] O'Donoghue, T. and Rabin, M. (1999). Doing it Now or Later. *The American Economic Review*, 89(1), 103-124

[18] Offerman, T., Schram, A. and Sonnemans, J. (1998). Quantal Response Models in Step-Level Public Good Games. *European Journal of Political Economy,* 14(1), 89-100

[19] Robson, A.J. (2001). The Biological Basis of Economic Behaviour. *Journal of Economic Literature,* 39(1), 11-33

[20] Rosenthal, R.W. (1992). Rules of Thumb in Games. *Journal of Economic Behaviour and Organisation*, 22(1), 1-13

[21] Thaler, R. (1980). Towards a Positive Theory of Consumer Choice. *Journal of Economic Behaviour and Organisation*, 1(1), 39-60

[22] Thaler, R. (1992). *The Winner's Curse: Paradoxes and Anomalies of Economic Life*, Rev. Ed., Princeton, USA: Princeton University Press

[23] Weibull, J.W. (1995). *Evolutionary Game Theory*, Rev. Ed., Cambridge, USA: MIT Press

[24] Young, P. (1993). The Evolution of Conventions. *Econometrica*, 61(1), 57-84