

Evolutionary Selection of Modular Decision Architectures

Alberto Bisin*

Emil Iantchev[†]

New York University and NBER

Syracuse University

May 21, 2009

Abstract

We study the evolutionary properties of decision processes. In particular, we show that in the presence of harmful mutations, a population of decision makers who possess an architecture consisting of hierarchically organised decision modules will have a strictly higher asymptotic growth factor than a population of decision makers with a unitary, fully connected, decision architecture. In addition, we show that under imperfect precision of the architecture and small cyclical fluctuations in the environment, conflict among the reference policies of decision modules arises as an evolutionary equilibrium. Finally, we show that economic models of multiple decision processes can be represented as examples of the type of modular hierarchy we investigate.

1 Introduction

In this paper we investigate the evolutionary selection of two stylized architectures for decision processes: i) a *unitary (U) architecture* in which one fully connected circuit always determines behavior and ii) a *modular hierarchical (MH) architecture* in which a supervisory module assigns control to one of two separate executive decision modules.

Our analysis is motivated by the recent interest of economists in the study of decision processes in environments where rational decision making, arguably, fares poorly.¹ In particular, this is the case for intertemporal choice in self-control environments, which has lead economists to formulate decision making

*Department of Economics, NYU, New York, NY 10012. E-mail: alberto.bisin@nyu.edu

[†]Department of Economics, Syracuse University, Syracuse, NY 13244. E-mail: iantchev@syr.edu

¹See for instance the surveys in Rabin (1998), Lowenstein and Prelec (1992), and Camerer (2009). On the methodological issues involved in modeling decision processes rather than choices and preferences, see Caplin and Schotter (2008).

models characterized by the interaction of different "actors", *multiple selves*, with conflicting models of behavior.² We show that these models have a representation in terms of the modular hierarchical decision processes studied in this paper.

The main results of this paper are as follows. Under a stable environment, but in the presence of deleterious mutations, we show that the MH has a strictly higher asymptotic growth factor than the U architecture. The MH architecture will thus outperform the unitary architecture in the long run starting from any initial condition. This is because the modular architecture is more robust with respect to the harmful effects of accumulating mutations. The U architecture is fully connected: a mutation in any part of the circuit will influence its overall performance level. The MH architecture, on the other hand, is only loosely connected: it will continue to operate nearly optimally in the presence of several kinds of harmful mutations.

We also compare the performance of the two architectures in the presence of cyclical fluctuations in the environment. Because of the different levels of connectedness, the two architectures will typically face a complexity-efficiency trade-off. In particular, a unitary *decision making process* might expend more energy while being able to respond more precisely to fluctuations in the environment. Under these conditions, and in the presence of harmful mutations, the relative performance of the modular architecture will depend on its precision level. Specifically, we show that if the precision of a MH architecture is above a critical level, it will outperform a unitary architecture of the same size starting from any initial condition. This is because the gain in fitness for the unitary architecture, due to increased accuracy, is more than offset by the loss in fitness due to increased energy consumption.

In addition, we show that under imperfect precision and cyclical fluctuations in the environment, conflict among decision modules and the corresponding heterogeneity in behavior among decision making processes is an integral part of an evolutionary equilibrium. Intuitively, in the presence of undetectable fluctuations in the environmental state, internal conflict serves as a diversification device. This is because a completely redundant, correctly specified model for the current state becomes a rigid, misspecified model in the event of a small environmental perturbation.

By subjecting decision processes to evolutionary selection, we stress genetic mutations. But are mutations really that important for the evolution of cognitive and behavioral processes in humans? Rather than surveying the nature-nurture question, we simply note that an affirmative answer is consistent with the

²See Frederick, Loewenstein, and O'Donoghue (2002) for a survey of the empirical evidence. Some examples of multiple selves models are Thaler and Shefrin (1981), Laibson (1997), Bernheim and Rangel (2004), Benhabib and Bisin (2005), Fudenberg and Levine (2006), Loewenstein and O'Donoghue (2007), and Brocas and Carrillo (2008).

fact that many cognitive and behavioral disorders have a prominent heritable component and that several psychiatric disorders, in addition to being strongly heritable are also associated with explicit structural and functional changes in certain brain regions.³ The interested reader should refer to chapters 7 and 16 in Breedlove, et al. (2007) for an introduction to the subject.

We believe our results provide an evolutionary justification for the models of multiple decision processes explored by economists, and in particular for multiple selves models. Because modular hierarchies are a type of neural network and hence can serve as universal approximators, they can provide a common conceptual foundation for all of these models. This could potentially be very useful for identification and testing since variation across models corresponds only to variation in the objectives of modules within the same MH architecture.

1.1 Related literature

This paper is related to the *indirect evolutionary approach*, which subjects agents' preferences to the analysis of evolutionary selection in order to identify their fundamental characteristics. Robson and Samuelson (2009) comprehensively survey this literature. In this paper, however, we subject agents' decision processes, rather than their preferences, to the analysis of evolutionary selection. Decision processes become a natural unit of analysis when agents' rationality is relaxed to allow for behavioral decision making.⁴ In particular, as we discuss in Section 4, several of the recent decision theoretic analyses of intertemporal choice (can be interpreted to) develop explicit models of decision processes.

Two recent papers, Dasgupta and Maskin (2005) and Netzer (2009), are closely related to our study as they also concentrate on the evolution of hierarchical decision processes. In both papers, however, hierarchical decision processes arise in response to very specific choice problems that the decision maker is supposed to have encountered during its evolutionary past. On the contrary, in our paper, hierarchical decision processes arise regardless of what choice problems the decision maker might encounter, suggesting that modularity might be relatively robust across decision problems. In addition, in Dasgupta and Maskin (2005) and Netzer (2009), no internal conflicts arise in equilibrium and those which are observed are explained as a remnant of our evolutionary past. In our paper, on the other hand, internal conflicts have a positive evolutionary value because they serve as diversification devices in the presence of undetectable fluctuations in the environment.

We restrict our analysis to the comparison of unitary architectures with the simplest modular architec-

³Down, fragile X, autism, Asperger's, and Tourette's Syndrome, for instance, all have explicit genetic basis. The same is true for disorders like schizophrenia, panic, and depression.

⁴See the papers by Rubinstein and Salant (2008) and Benhabib and Bisin (2008) in Caplin and Schotter (2008).

tures, those consisting of two executive and one controlling processes. In particular, in our set-up, each executive process selects a feasible reference course of action. The controlling process then processes both feasible actions and chooses one of the two. This particular class of decision making architectures often arises in artificial life simulations, see for instance Cangelosi, Parisi, and Nolfi (1994), where a controlling unit evolves to monitor internal and external states and to assign control to one of the executive processes. Furthermore, this class of hierarchical architectures fits well with recent models and data on cognitive control developed in neuroeconomics.⁵ In particular, hierarchical architectures can represent the interaction between the prefrontal cortex and the anterior cingulate cortex. Miller and Cohen (2001) provide theory and evidence that distinct areas in the human brain are associated with the monitoring of conflict and the execution of action. The anterior cingulate cortex (ACC) appears to be involved primarily in the monitoring of cognitive conflict. However, the ACC does not appear to be directly involved in executive functions. Rather, it exerts influence by activating other regions, such as the dorsolateral prefrontal cortex, which are ultimately responsible for adjustments in behavior.⁶ In the context of intertemporal choice, evidence for decision processes based on the interaction between the prefrontal cortex and the anterior cingulate cortex is presented by McClure, Laibson, Loewenstein, and Cohen (2004).⁷

Finally, our analysis is also related to the recent literature on *kludges*⁸—see Ely (2007) and Marcus (2008). A *kludge* is a "marginal adaptation which compensates for, but does not eliminate fundamental design inefficiencies" in brain processes - from Ely (2007). Modular decision processes, in our analysis, can in fact be interpreted as a kludge.

2 Set-up

Consider a decision-making process embedded in an agent that is a member of an age-structured population with a reproductive life lasting n periods. We keep the standard practice in age-structured models and treat a period of chronological time and an age class as equivalent. For each age $i = 1, \dots, n$, let the environment in a given state be denoted by $s(i) \in S$. The vector $\varsigma = (s(1), s(2), \dots, s(n))$ denotes then the environment during the agent's life, which is an element of some metric space S^n endowed with a norm d . In any environment $s(i) \in S$, a *decision-making process* (DP) executes an action a from a given set A , endowed with the norm

⁵See Glimcher (2003) and Camerer, Loewenstein, and Prelec (2005) for general surveys of the field.

⁶See also, Kerns, Cohen, MacDonald III, Cho, Stenger, and Carter (2004).

⁷But see Glimcher, Kable, and Louie (2007) for a skeptical view of these results.

⁸Sometimes also referred to as *kluges*.

w .⁹

Let $\mathbf{a}^\zeta = (a_1^{s(1)}, \dots, a_1^{s(n)}) \in A^n$ denote the action profile with the lowest asymptotic fitness loss (relative to the maximum) in environment ζ . Let $L_\zeta(\mathbf{a}, \mathbf{a}^\zeta)$ be the asymptotic loss associated with any profile $\mathbf{a} \in A^n$.¹⁰ Our goal is to examine the asymptotic growth factors of the two decision architectures, unitary (U) and modular hierarchical (MH), in the presence of fitness reducing mutations. In the U architecture, one single choice process determines the mapping between S^n and A^n . In the MH architecture, multiple, nearly decomposable, decision modules interact to produce a map between S^n and A^n . Next, we formally describe the decision process of each architecture.

Unitary Architecture. Let $\mathbf{a}^{U,\zeta}$ denote the hereditarily acquired action of a unitary architecture in environment ζ . If profile $\mathbf{a}^{U,\zeta}$ is executed, the asymptotic fitness loss of the organism will be given by $L_\zeta(\mathbf{a}^{U,\zeta}, \mathbf{a}^\zeta)$. In the long run, DPs that have the correct reference model, \mathbf{a}^ζ , for the environment ζ , will dominate the population.

Modular Hierarchical Architecture. Let a DP consist of two stage-one modules (modules 1 and 1') and one stage-two module (module 2). Each stage-one module has its own choice profile, namely $\mathbf{a}_1^\zeta = (a_{11}^{s(1)}, \dots, a_{1n}^{s(n)})$ and $\mathbf{a}_{1'}^\zeta$. The same is true for module 2, which has a profile given by \mathbf{a}_2^ζ . For age i and state $s = s(i)$ in environment ζ , module 2 allocates control to either module 1 or 1' by solving

$$\max_{a \in \{a_{1i}^s, a_{1'i}^s\}} -w(a, a_{2i}^s).$$

Basically, module 2 allocates control to the executive module whose action is closer to the reference action of module 2. Let $\mathbf{a}^{H,\zeta}$ be the solution to such a problem. The asymptotic fitness loss will be minimal for any DP with $\mathbf{a}^{MH,\zeta} = \mathbf{a}^\zeta$. This can be achieved if:

- (i) $\mathbf{a}^\zeta = \mathbf{a}_1^\zeta = \mathbf{a}_{1'}^\zeta$, for any feasible profile \mathbf{a}_2^ζ ;
- (ii) $\mathbf{a}^\zeta = \mathbf{a}_j^\zeta \neq \mathbf{a}_k^\zeta$ for $j = 1$ or $1'$, $k \neq j$ and a_{2i}^s is closer to a_{ji}^s than to a_{ki}^s for every state s in ζ and every i .

Hence for the modular hierarchy to survive in a stable environment, we need at least one of the module 1 actions to minimize fitness loss and the action of module 2 to be closer to the fitness minimizing action than the action associated with the other executive module.

⁹Note that the unit of analysis in our context is a process rather than an agent, a decision-making process and not a decision maker.

¹⁰The function L_ζ will typically be a monotonically increasing function of the distances $w(a_i, a^{s(i)})$. However, we do not require that to be the case.

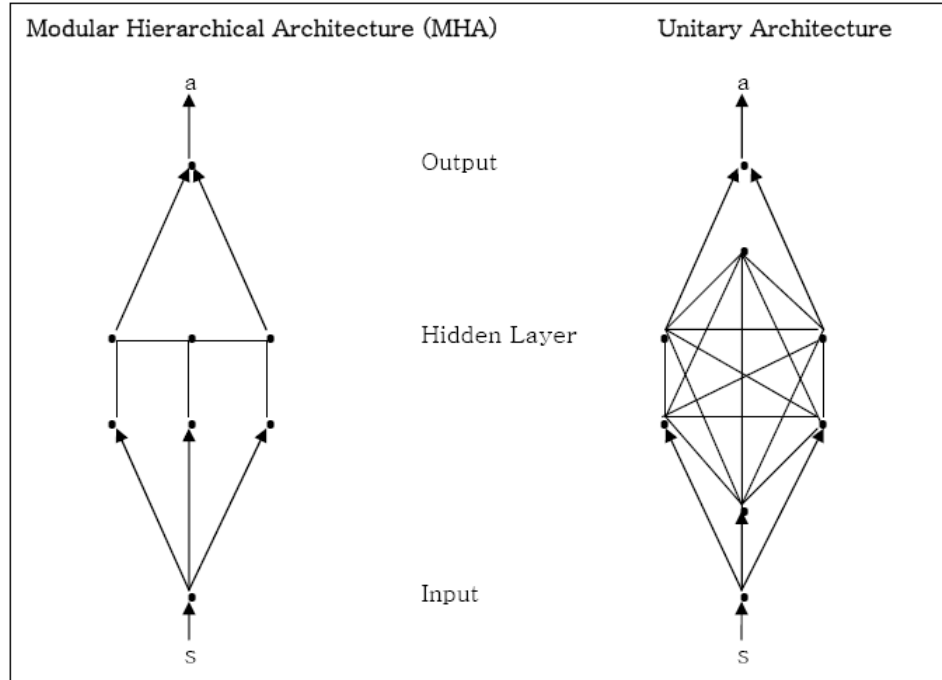


Figure 1: *Example of a unitary and a modular hierarchical decision architecture with 3 input, 2 output connections, and 6 hidden layer nodes. Undirected edges represent connections in both directions.*

Figure 1 provides an illustration of the difference between the two types of decision architectures. In the unitary architecture, all 6 hidden layer nodes, each of which could be a network in itself, are interconnected with each other. The architecture’s response to a signal is thus determined through an interaction among all nodes in the hidden layer. More importantly, a change at a particular node, *ceteris paribus*, will always influence the architecture’s overall response since that node is connected to every other node in the hidden layer.

The MH architecture on the other hand is only sparsely connected. The left and right side columns in the hidden layer correspond to the two executive modules responsible for the architecture’s eventual response. The middle column is the controller that aggregates all information—the input signals and the recommended policy by each executive module—and allocates control to one of the nodes with output connections. The difference here is that there are direct connections among nodes within a module on a given level of the hierarchy but not across modules on the same level—they interact through the higher level controlling module. The MH architecture is thus nearly decomposable in the sense that changes at a particular node, *ceteris*

paribus, need not have global implications for behavior.¹¹

Finally, note that for the same number of hidden layer nodes, there are 5 bilateral connections in the MH architectures versus 16 in the unitary. If operating and maintaining connections requires energy, the MH architecture will be much more energy efficient. On the other hand, the unitary architecture may exhibit much more sophisticated behavior since the greater number of connections allows for more nonlinear transformations of the input signals. This complexity-efficiency trade-off is at the heart of the results derived in section 3.2.

3 Evolutionary selection

In order to determine which architecture is more likely to be evolutionary successful, we subject them to mutations. The architecture which has a higher asymptotic growth factor under these conditions will thus be more likely prevail in the long run. We study evolutionary selection for unitary *vs.* modular hierarchical architectures in the presence of deleterious mutations i) under a stable environment and ii) in the presence of small cyclical fluctuations in the environment.

For each age $i = 1, \dots, n$, with the environment in a given state $s(i) = s \in S$, a *decision-making process* of type $K = U, MH$ executes an action a^K from the given set A . Let $G_i^K(s)$ denote the expected number of surviving offsprings from an agent of age i with decision process K in state s , *under no deleterious mutations*.¹² Also, let each agent of age i survive with probability $P_i^K(s)$ if $i < n$, and with probability 0 if $i = n$, again *under no deleterious mutations*. Each architecture is endowed with a measure of *precision*: we say that architecture K has *precision* $\frac{1}{\phi^K}$ if it can perceive the difference between two arbitrary states $s, s' \in S$ only when their distance is larger than ϕ^K : only when $d(s, s') \geq \phi^K$.¹³

Let q be the per period probability with which a harmful mutation occurs in a decision node. A mutation changes the reference action profile of the module in which it occurs. Since mutations are rare events, we assume that q is sufficiently small so that multiple mutations per period can be ignored.¹⁴ In addition, since the vast majority of mutations are harmful (there are many more ways to increase than decrease the fitness loss), the probability of a reversal (a mutation in the direction of the loss minimizing action), r , will typically

¹¹For a thorough discussion of nearly decomposable architectures see Simon (1996).

¹²The expectation here is taken over some distribution of idiosyncratic individual shocks which average out in a sufficiently large population.

¹³Essentially, this is a measure of how well the organism can separate the fundamental change due to a change in the environmental state from the noise generated by the idiosyncratic shocks at the individual level.

¹⁴This is consistent with the standard model of Luria and Delbrueck (1943) where mutations are governed by a Poisson process independent of the environment.

be even smaller: $r \leq q$.

The effect of deleterious mutations is as follows:

- every time a deleterious mutation occurs, the expected number of surviving offsprings, $G_i^K(s)$, resulting from executing the changed action profile is reduced by a factor $\delta_g \in (0, 1)$, for any i and any K ;
- every time a deleterious mutation occurs, the expected survival probability of the DP, $P_i^K(s)$, resulting from executing the changed action profile is also reduced by a factor $\delta_p \in (0, 1)$, for any i and any K ;¹⁵
- the offsprings of an agent carrying m mutations are born with m mutations;
- a DP can carry on at most W deleterious mutations: the accumulation (hereditary or not) of $W + 1$ deleterious mutations is lethal.

Furthermore, a complexity-efficiency trade-off will determine each architecture's relative fitness. The more energy consuming an architecture, the smaller will be, other things equal, the number of its offsprings and its survival probability. Specifically, the more energy consuming an architecture K is, the smaller will $G_i^K(s)$ and $P_i^K(s)$ be for given i and s . On the other hand, the more energy consuming an architecture is, the more precisely it will be responding to fluctuations in the environmental state $s \in S$. Specifically, the more energy consuming an architecture K is, the higher its precision $\frac{1}{\phi^K}$ will be.

As noted before, the crucial distinction between unitary and hierarchical decision-making architectures is that the U architecture is fully connected, while the MH is only sparsely connected. Our reading of the literature on energy budgeting of the human brain¹⁶ leads us to formulate the following implications of the different levels of connectedness for survival and precision.

Assumption 1. *A complexity-efficiency trade-off characterizes the U and MH architectures:*

1. *the MH architecture is more energy efficient*

$$G_i^{MH}(s) \geq G_i^U(s) \text{ and } P_i^{MH}(s) \geq P_i^U(s), \text{ for any } i, s;$$

while, on the other hand,

¹⁵All results go through without modification if the effect of mutations is additive rather than multiplicative.

¹⁶For a summary of recent studies on the energy budget of the human brain see Raichle and Gusnard (2002) and the references therein.

2. the U architecture is more precise

$$\frac{1}{\phi^U} \geq \frac{1}{\phi^{MH}}.^{17}$$

We say the the complexity-efficiency trade-off is non-trivial if the inequalities hold strictly.

3.1 Stable environment

We are now ready to study evolutionary selection of the two DP architectures. We first perform the exercise in a stable environment. We shall study cyclical variations of the environmental state in the following section.

Assumption 2-s. *The environment is stable, $s(i) = s \forall i$; that is, $s(i)$ does not vary over the lifetime $i = 1, \dots, n$ of any agent.*

We can then drop reference to the environmental state s in the notation of this section, without loss of generality. Consider first the MH architecture. Let $N_{i,t}^{j,k,l}$ denote the size at time t of the subpopulation of processes of age i characterized by a total of $j + k + l$ mutations, with j , k , and l mutations in each module. For expositional clarity, we consider the case with 2 age classes (two period life), and 2 maximum mutations ($n = W = 2$).¹⁸ In this case the population structure vector at the beginning of period t is given by $\mathbf{N}_t = [N_{1,t}^{0,0,0}, N_{2,t}^{0,0,0}, N_{1,t}^{1,0,0}, N_{2,t}^{1,0,0}, N_{1,t}^{1,1,0}, N_{2,t}^{1,1,0}, N_{1,t}^{2,0,0}, N_{2,t}^{2,0,0}]^T$. For the MH architecture, we need to consider the distribution of mutations across the decision modules. This is because a process with two deleterious mutations in the same module will continue to operate optimally regardless of which module is affected. On the other hand, a processes with one mutation along any two of the three decision modules will select a suboptimal course of action.¹⁹

The dynamics of the population of MH processes are governed by the linear system

$$\mathbf{N}_{t+1} = A^{MH} \mathbf{N}_t,$$

where, assuming each module is equally likely to mutate, A^{MH} is the 8×8 nonnegative projection matrix

¹⁷Attwell and Laughlin (2001) for instance estimate that over 3/4 of the brain's energy consumption is used for signalling among neurons. For a given number of nodes (neurons) in a network, an architecture like U with a large number of connections among those nodes will need to generate and maintain a large number of signals and will thus consume more energy than a sparsely connected architecture with the same number of nodes. Given the correct action profile and the same initial amount of resources, the unitary architecture will consume more energy, which will result in $G^U < G^{MH}$ and the same for P .

On the other hand, the unitary architecture allows for a greater number of nonlinear transformations of any input which, given the universal approximation results by Hornik, Stinchcombe, and White (1989, 1990), suggests that it should be able to approximate more precisely the loss minimizing policy with respect to small perturbations in the environmental state.

¹⁸All result are proved, however, for the general case of arbitrary n and W .

¹⁹If module 1 has the correct action while modules 1' and 2 have 1 deleterious mutation each, we have $w(a_{1i}^s, a_{2i}^s) \geq w(a_{1'i}^s, a_{2i}^s)$ for each age $i = 1, \dots, n$. In all proofs, we will bias the scenario against the success of the MH architecture by assuming that whenever $w(a_{1i}^s, a_{2i}^s) = w(a_{1'i}^s, a_{2i}^s)$, the less optimal action is executed.

than the unitary architecture starting from any non-zero initial condition. Furthermore, this is the case even if

$$G_i^{MH} = G_i^U = G_i \text{ and } P_i^{MH} = P_i^U = P_i.$$

Proof: Consider without loss of generality the case in which $G_i^{MH} = G_i^U = G_i$ and $P_i^{MH} = P_i^U = P_i$. The dynamics of the two subpopulations, U and MH , are governed, respectively by $\mathbf{N}_{t+1} = A^U \mathbf{N}_t$ and $\mathbf{N}_{t+1} = A^{MH} \mathbf{N}_t$. The Perron-Frobenius theorem for nonnegative matrices implies that the asymptotic growth factor of each population is governed by the spectral radius of, respectively the matrix A^U , and A^{MH} . The spectral radius, ρ^K , is given by the dominant root of the associated characteristic equation. So in principle, if we have two competing populations, only the one with the higher asymptotic growth factor will survive in the long run. We show in the appendix that the spectral radius of A^{MH} is larger than the spectral radius of A^U . ■

The result in Proposition 1 states that the MH architecture does strictly better than the U architecture even if the hierarchical decision process is not more energy efficient than the unitary process, and hence even if MH has no ex-ante evolutionary advantage over U in terms of the number of surviving offsprings, G_i^K , and survival probability, P_i^K .

The intuition behind Proposition 1 is quite straightforward. Comparing A^{MH} and A^U we can see that every entry in A^{MH} is greater than or equal to the corresponding entry in A^U . This is not surprising as every mutation reduces fitness in the U architecture, which is not the case for the MH architecture. The unitary architecture is fully connected. Hence the accumulation of mutations in any part of the decision architecture will influence the DP's overall course of action. This is not true for the modular architecture. The low level modules, for instance, are not connected to each other so a change in policy for just one of those will not affect the DP's course of action. Thus for significant fitness loss to occur, we need multiple changes in the MH process to take place simultaneously, an *essentially* zero probability event which can be ignored. The MH architecture thus survives under a stable environment because it is more robust to the accumulation of deleterious mutations over a DP's evolutionary dynamics.

3.2 Cyclically fluctuating environment

In a stable environment, through the gradual process of mutation, both unitary and hierarchical architectures will eventually stumble upon the optimal action profile \mathbf{a}^s . In this section, we consider small cyclical fluctuations in the environment.

Assumption 2-f. *The environment fluctuates cyclically, that is, the environmental state switches between s and s' after every $P \geq 1$ periods, where $\phi_- \leq d(s, s') \leq \phi_+$. Furthermore, fluctuations are non-trivial, that is, $a_i^{s'} \neq a_i^s$ for at least one age class i .*

We shall consider a U architecture of precision $\frac{1}{\phi^U} \geq \frac{1}{\phi_-}$ and a MH architecture of precision $\frac{1}{\phi^{MH}} \leq \frac{1}{\phi_+}$. Effectively we are assuming that the U architecture can perceive the fluctuating environment while the MH cannot. Consistently with the complexity-efficiency trade-off (Assumption 1), we assume that continuing to execute the loss minimizing action under state s when the environmental state is s' (and vice versa) reduces offsprings and survival. Let $G_i^K(s | s')$ denote the number of offsprings from an agent of age i with decision process K in state s , when the state s' is perceived. Also, let $P_i^K(s | s')$ denote the survival probability of an agent of age i with decision process K in state s , when the state s' is perceived. More precisely, a modular DP executing action a^s under state s' will feature

$$G_i^{MH}(s' | s) = \gamma_G \left(\frac{1}{\phi^{MH}} \right) G_i^{MH}(s') \text{ and } P_i^{MH}(s' | s) = \gamma_P \left(\frac{1}{\phi^{MH}} \right) P_i^{MH}(s'),$$

where $\gamma_G(\cdot)$ and $\gamma_P(\cdot)$ take values in the interval $[0, 1]$ with $\gamma_j(0) = 0$ and $\lim_{x \downarrow \phi_-} \gamma_j(1/x) = 1$ for $j = G, P$. We also assume that $\gamma_G \left(\frac{1}{\phi^{MH}} \right)$ and $\gamma_P \left(\frac{1}{\phi^{MH}} \right)$ increase with ϕ^{MH} so that loss of offsprings and survival depend negatively on precision.

Because of its greater complexity, a unitary DP suffers no reduction in fitness as a result of the environmental fluctuations, while a MH does. The unitary DP, however, consumes more energy. This is a *strict* version of the complexity-efficiency trade-off introduced in Assumption 1. In this environment it is ambiguous which of the two architectures is selected. It turns out however that we can still classify the comparative advantage of each architecture.

Proposition 2. *In the presence of a strict complexity-efficiency trade-off (Assumption 1) and a cyclically fluctuating environment (Assumption 2-f), if the precision $\frac{1}{\phi^{MH}}$ of the MH architecture is greater than a critical value, then the MH architecture has a strictly higher asymptotic population growth factor than the corresponding unitary architecture under cycles of any frequency and starting from any non-zero initial condition. On the other hand, if the MH architecture is sufficiently imprecise, $\frac{1}{\phi^{MH}} \rightarrow 0$, the conclusion is reversed and the U architecture has a higher asymptotic population growth factor.*

Proof: The proof is a corollary of Proposition 1. However, when the environmental state changes from s to s' , the entire matrix A^{MH} changes. With different precisions, it may well be the case that the spectral radius of $A^{MH}(s)$ is greater than that of $A^U(s)$, while the spectral radius of $A^{MH}(s')$ is smaller than that

of $A^U(s')$. When $P = 1$, for instance, the asymptotic behavior of the system is determined by the spectral radii of the product matrices and we need to show that $A^{MH}(s')A^{MH}(s)$ and $A^{MH}(s)A^{MH}(s')$ have larger spectral radii than $A^U(s')A^U(s)$ and $A^U(s)A^U(s')$ respectively. See the appendix for the formal argument.

■

Intuitively, for any given loss minimizing action profile \mathbf{a}^s , the unitary architecture is less energy efficient, $G_i^U < G_i^{MH}$ and $P_i^U < P_i^{MH}$. There is thus some space for the MH architecture to be less efficient than the unitary on the precision dimension while still winning the evolutionary race. Therefore, for relatively precise MH architectures, the marginal benefit of increasing precision by moving to a unitary architecture is outweighed by the marginal cost of increased energy consumption. For relatively imprecise MH architectures, the opposite obtains.²⁰

3.2.1 Cyclical fluctuations and equilibrium heterogeneity in action profiles

Under a stable environment, the population will be dominated by DPs that minimize the loss of fitness. Evolutionary selection, therefore, will not lead to the heterogeneity in actual behavior we arguably observe. Hence our next step is to investigate a setup which might induce heterogeneity in behavior as an evolutionary equilibrium.

Propositions 1 and 2 ignore the possibility of mutations being beneficial. In the cyclically fluctuating environment, under our assumptions, however, the MH architecture chooses the action which is optimal for state s even when the environmental state is s' . Beneficial mutations therefore exist for this architecture, which push the actions of a module in the direction of $a_i^{s'}$ and reduce the costs associated with the lack of precision of the MH architecture. We expect such mutations to occur with a probability less than or equal to q . The fitness of a DP will then depend on the distribution of mutations over its modules.

Proposition 3. *In the presence of a strict complexity-efficiency trade-off (Assumption 1), in a cyclically fluctuating environment (Assumption 2-f), and under natural assumptions about the effects of beneficial mutations (see the Appendix for details), evolutionary selection induces (i) a non-trivial fraction of the population with decision modules that have conflicting models of optimal behavior and (ii) a non-trivial*

²⁰Proposition 2 suggests that modular decision architectures should be observed later in evolutionary history, assuming encephalization increases with time. Empirically, this conjecture is in line with studies in evolutionary biology. For example the cortex, which in humans is posited to contain decision processes in addition to those of the reptilian brain, appears later in evolutionary history (in mammals) and matures later in life, particularly in humans. The distinction need not be primarily structural, however, but may instead be functional. Rats and humans for instance have the same mammalian brain structures. Despite this structural similarity, Proposition 2 implies that rats should behave more in line with the unitary model since they are much less encephalized than humans. And indeed, laboratory experiments, (Gardner and David 1999, Berridge 1999), show that to be the case.

number of DPs that choose actions which are suboptimal with respect to the current environmental state.

The intuition behind Proposition 3 is that when precision is imperfect, conflict among the reference policies of modules has a positive value from an evolutionary standpoint. This is because a redundant, correctly specified model for the current state becomes a rigid, misspecified model in the event of a small environmental perturbation. We illustrate the importance of such internal conflicts with the following simple example.

Numerical example Consider a modular DP with precision $1/\phi$ whose reproductive life lasts one period. We start in environmental state s but after every generation, the state switches between s and s' where $d(s, s') < \phi$. Let a^s denote the optimal action in environment s and let the number of surviving offsprings associated with this action be $G(s) = g = 1.03$ (and, symmetrically, $G(s') = g$). Assume that under environment s' , however, we have $G(s' | s) = \xi^2 g$ with $\xi = 0.98$ (and symmetrically, $G(s | s') = \xi^2 g$).

Suppose that there is a mutation which pushes the action in the direction of $a^{s'}$. One push results in the action a^I which has values for offspring of $G^{a^I}(s) = G^{a^I}(s') = \xi g$. Two pushes are sufficient to change the action from a^s to $a^{s'}$ and vice versa. Assume that $w(a^I, a^{s'}) < w(a^I, a^s)$ so that in action space, a^I is closer to $a^{s'}$ than to a^s . In a modular DP, two pushes require at least four mutations, two at any two of the three modules. Assume a mutation happens with probability $r = 0.0001$ per generation. Finally, assume that a DP can carry on at most 4 of these mutations. A DP with 5 mutations leaves no offspring in either state of nature.²¹

We have the following law of motion, where a cycle lasts two generations,

$$\mathbf{N}_{2i} = [A(s')A(s)]\mathbf{N}_{2(i-1)}.$$

The population structure vector is of dimension 13 since we have to take into account how the distribution of mutations across the modules of a given type of DP influences its current action and the transition probabilities to other types. Similarly, the projection matrices are 13×13 . A more detailed description of the setup and results can be found in the appendix.

It is well known (e.g., Seneta (1981), Theorem 1.2) that the asymptotic distribution of the population for a projection matrix A is given by the positive, normalized left eigenvector u associated with the spectral radius $\rho(A)$, that is $u'A = \rho(A)u'$ with $\sum_{i=1}^{13} u_i = 1$. Figure 2 shows the asymptotic distribution of the

²¹All of these assumptions are made for tractability, since the complexity of the problem increases very quickly with n and W .




Figure 2: *Asymptotic distribution of the population over the possible actions $\{a^s, a^I, a^{s'}\}$ under three environments.*

population over the possible actions $\{a^s, a^I, a^{s'}\}$ under three scenarios: (i) a stable environment at s , (ii) a stable environment at s' , and (iii) a generational cycle between s and s' . As expected, under scenarios (i) and (ii), the asymptotic distribution is totally dominated by types that select the optimal course of action for the stable environment. Even though there is heterogeneity in reference policies among modules for a large fraction of the population, that conflict does not translate into behavioral heterogeneity.

Under a fluctuating environment, on the other hand, we have heterogeneity in both reference models and actual behavior.²² In every generation, between 45 and 90 percent of the population choose a suboptimal course of action. Even more interestingly, there are DPs—such as the type with three total mutations distributed across two modules—which possess the correct policies for each state, yet the internal conflict leads to the suboptimal course of action a^I being chosen under both s and s' . Intuitively, under imperfect precision and fluctuating environment, it pays for the population to diversify by having types which face internal conflict or exhibit suboptimal behavior. This is because it will be easier for the population to respond to the fluctuations—the types with internal conflict or suboptimal behavior will either already be behaving optimally or will require fewer mutations to start doing so when the state changes.

²²The distribution is asymmetric because of our assumption that a DP can carry on at most 4 mutations. In order to have an architecture with a completely redundant reference policy $a^{s'}$ we require 6 mutations, which is ruled out a priori. Hence our example is inherently biased against the presence of types that execute the action $a^{s'}$.

4 Hierarchical decision processes in economics

Our results provide an evolutionary justification for the models of *multiple selves* recently explored by economists. For instance, our result (Proposition 3) that evolutionary selection would favor hierarchical modules characterized by conflicting models of optimal behavior, and even by choices which are suboptimal with respect to the current environmental state, could explain the "puzzles" observed in laboratory and field studies of intertemporal choice in self-control environments. Frederick, Loewenstein, and O'Donoghue (2002) comprehensively survey this evidence. To illustrate this point, we translate into our framework the first multiple selves model, developed by Thaler and Shefrin (1981), that studies self-control and attempts to explain these intertemporal choice "puzzles".²³

This model is formulated in terms of the strategic interaction between two "actors", a farsighted "planner" and a myopic "doer". The two actors are explicitly interconnected in the sense that actions by one agent directly influence the objectives of the other. In this sense, it does not immediately appear that a modular hierarchical decision process could represent an individual decision maker in Thaler and Shefrin. It turns out however that when the objective functions are properly defined, the policy function of a given MH architecture corresponds to a discrete choice approximation of the policy function of their model.

In Thaler-Shefrin's model, the farsighted planner and the myopic doer are distinct in the utility functions with which they evaluate consumption plans. Specifically, the doer cares only about current consumption, which is represented by some concave utility function $Z_t(c_t, \theta_t)$, with $Z_t(c_t, 0)$ strictly increasing, and where $\theta_t \geq 0$ is a parameter which influences the value of c_t at which Z_t reaches a maximum. The planner, on the other hand, cares about lifetime utility represented by some strictly increasing function $V(Z_1, Z_2, \dots, Z_T)$.

Given that the preferences of the two actors differ, there will typically be a conflict between them. Thaler-Shefrin assume that the planner can modify the doer's behavior at time t by adjusting the preference parameter θ_t . The value of θ_t thus represents the modification the planner exerts on the doer at time t , where $\theta_t = 0$ means no modification. Influencing the doer is costly, however, in the sense that $\partial Z_t / \partial \theta_t < 0$. In addition, it is assumed that $d\theta_t / dc_t(Z_t = Z) < 0$: the lower the desired value of c_t , the bigger the modification and hence the cost that is required to adjust behavior.

Let $c_t(\theta_t)$ denote the policy of a doer with utility function $Z_t(c_t, \theta_t)$. Given a lifetime income stream with present value y , an individual consisting of a planner and a doer will behave as if solving

²³See also Fudenberg and Levine (2006) and Loewenstein and O'Donoghue (2007) for recent extensions.

$$\max_{(\theta_1, \dots, \theta_T) \geq \mathbf{0}} V(Z_1(c_1(\theta_1), \theta_1), \dots, Z_T(c_T(\theta_T), \theta_T))$$

$$\text{s.t.} \quad \sum_t c_t(\theta_t) \leq y.$$

We now map a two-period version of this model into our setup and show that the modular hierarchy considered in our paper provides a discrete choice approximation to the above framework. Consider an individual with endowment y who lives for two periods, that is, who has to choose how to allocate consumption between two periods. Suppose the individual behaves according to a modular hierarchical decision process in which each of the two level 1 modules are defined by the vector $(\theta_i, 0)$, for $i = 1, 1'$.²⁴ That is, for each executive module, the objective $Z(c, 0)$ is strictly increasing in consumption in the second (last) period in life. Module i then recommends consumption for period 1 given by

$$c(\theta_i) = \arg \max_{c \in [0, y]} Z(c, \theta_i).$$

On the other hand, the reference policy of the level 2 module (the controller), is given by

$$c_2 = \arg \max_{c(\theta) \in [0, y]} V[Z(c(\theta), \theta), Z(y - c(\theta), 0)],$$

where $c(\theta)$ is the policy associated with the objective $Z(c, \theta)$. Then the DM will behave according to c^* given by

$$c^* = \arg \max_{c \in \{c(\theta_1), c(\theta_{1'})\}} - [c - c_2]^2.$$

Note that if the parameter space in Thaler and Shefrin (1981) is reduced to the set $\{0, \theta_1, \theta_{1'}\}$, the policy of the MH architecture described above will correspond exactly to the optimal policy of the decision maker in their model. Hence the MH architecture produces a discrete choice approximation to the behavior described by Thaler and Shefrin's self-control model.

Recently, Bernheim and Rangel (2004) and Benhabib and Bisin (2005) have developed models of multiple decision processes that have more explicit neuroscientific basis. They fit into the MH architecture exactly rather than as approximations. The crucial difference is that in these models there are multiple level 1

²⁴A natural interpretation of this formulation of the level 1 modules is that they represent different self-control levels.

modules with independent objectives. Each level 1 module recommends a course of action and control is allocated to one of the modules based on the particular circumstances. This last step corresponds precisely to the level 2 of a MH architecture where one of the policies recommended by the first stage modules is executed. A detailed MH description for each of these models is available from the authors by request.

5 Conclusion

We have shown that modular hierarchical architectures of decision processes can be evolutionary selected. This is because a MH architecture saves energy while providing robustness against the effects of harmful mutations. In addition, we have shown that under a cyclically fluctuating environment, conflict among decision modules appears to have a positive evolutionary value, as it is a source of diversification against the influences of the environmental fluctuations.

It should be noted that we have analyzed only one potential benefit of modular hierarchical architectures. In fact, MH architectures may confer many other benefits to the decision makers that possess them. Some promising directions include the possibility for greater adaptability in response to observable changes in the environment, the ability to process information in parallel fashion, and the ability for certain modules to specialize in the analysis and execution of certain behaviors. We believe exploring any of these will improve our understanding of decision processes and hence of the behavior of decision makers who possess them.

6 Appendix

Proof of Proposition 1:

In the general case with n age groups and W maximum total mutations, the population structure vector will enumerate (i) all types with different total mutation loads and (ii) for a given total number of mutations, all possible distributions of mutations across the nodes in the given decision architecture. We can then compare the two projection matrices A^U and A^{MH} entry by entry. Consider the matrix row associated with type $N_i^{j,k,l}$, where we need $j + k + l \leq W$ and $i \leq n$. Each element in that row will be a product of the transition probability between the two particular types and the value for the number of surviving offspring or the survival probability of type $N_i^{j,k,l}$. The transition probabilities are the same in both matrices since the mutation process is exactly the same in both decision architectures.

Now consider the values for surviving offspring and survival probability for type $N_i^{j,k,l}$. For the unitary

architecture, regardless of the distribution of mutations across the circuit, these values are $(\delta_g)^{j+k+l}G_i$ and $(\delta_p)^{j+k+l}P_i$ respectively. On the other hand, in the MH architecture, we know that these values will be greater than or equal to $(\delta_g)^{\max\{j,k,l\}}G_i$ and $(\delta_p)^{\max\{j,k,l\}}P_i$. Specifically, the equality will hold in the situation where $\max\{j,k,l\}$ mutations occur in an executive module and the number of mutations in the controlling module brings its reference policy closer to the action associated with $\max\{j,k,l\}$ mutations than to the reference action of the other executive module. In all other cases, $N_i^{j,k,l}$ will have a strictly greater number of offspring and survival probability. But then, given that $\delta_g < 1$ and $\delta_p < 1$, the projection matrix A^{MH} can be obtained from A^U by increasing each of its entries by either 0 or some positive amount. But since A^U is a nonnegative irreducible matrix, the Perron-Frobenius theorem [see for instance Thm 2.7. in Varga (2000)] implies that we must have $\rho^{MH} > \rho^U$. ■

Proof of Proposition 2:

Consider the threshold value ϕ^* defined by

$$\phi^* = \arg \min_{x \geq \phi_-} \{(\gamma_G(1/x)G_i^{MH}(s'))_{i=1}^n, (\gamma_P(1/x)P_i^{MH}(s'))_{i=1}^n\}$$

$$\begin{aligned} \text{s.t. } \quad \gamma_G(1/x)G_i^{MH}(s') &\geq G_i^U(s') & i = 1, \dots, n \\ \gamma_P(1/x)P_i^{MH}(s') &\geq P_i^U(s') & i = 1, \dots, n. \end{aligned}$$

Clearly, given that $\lim_{x \downarrow \phi_-} \gamma_j(1/x) = 1$ and $\gamma_j(0) = 0$, a solution to the above program will always exist. Moreover, given that $G_i^{MH}(s') > G_i^U(s')$ and $P_i^{MH}(s') > P_i^U(s')$, we know that $\phi^* > \phi_-$. But then for a MH architecture of precision $\phi \in [\phi_-, \phi^*)$,

$$\gamma_G(1/\phi)G_i^{MH}(s') > \gamma_G(1/\phi^*)G_i^{MH}(s') \geq G_i^U(s')$$

and

$$\gamma_P(1/\phi)P_i^{MH}(s') > \gamma_P(1/\phi^*)P_i^{MH}(s') \geq P_i^U(s').$$

The above inequalities say that for precision levels of the modular architecture above $1/\phi^*$, the benefit of increasing precision by going to a unitary architecture is outweighed by the increased energy consumption

of that architecture.

So suppose we start in environment s . Following the reasoning of Proposition 1, we know that $A^{MH}(s)$ is an increasing transformation of $A^U(s)$ and hence $\rho^{MH}(s) > \rho^U(s)$. Now suppose the environment changes permanently to s' . The above inequalities imply that $A^{MH}(s')$ is again an increasing transformation of $A^U(s')$ which implies $\rho^{MH}(s') > \rho^U(s')$. Note that when the environmental state changes, the unitary DP switches to the newly optimal profile. On the other hand, the modular DP continues to operate according to the old profile. Even though the modular DP behaves suboptimally, for ϕ small, its superior energy efficiency allows it to outperform the unitary DP even under the new environment.

Under short cycles, $P = 1$, the asymptotic behavior of the system will be determined by the spectral radii of the products $A(s')A(s)$ and $A(s)A(s')$. To see this, note that starting in environment s , after i such cycles we will have

$$\mathbf{N}_{2:i} = [A(s')A(s)]\mathbf{N}_{2:(i-1)}.$$

But we know that $A^{MH}(s)$ and $A^{MH}(s')$ are increasing transformations of $A^U(s)$ and $A^U(s')$. And since all of these matrices are nonnegative, the products $A^{MH}(s)A^{MH}(s')$ and $A^{MH}(s')A^{MH}(s)$ must be increasing transformations of $A^U(s)A^U(s')$ and $A^U(s')A^U(s)$. This is because each element of the product matrices is a dot product of two nonnegative vectors. For instance, let $A_{ij}^K(s, s')$ be the row i -column j element of $A^K(s)A^K(s')$. Clearly, $A_{ij}^K(s, s') = \sum_j A_{ij}^K(s)A_{ji}^K(s')$. But since $A_{ij}^{MH}(s) \geq A_{ij}^U(s)$ and $A_{ij}^{MH}(s') \geq A_{ij}^U(s')$ for any i and j and all matrices are nonnegative, we must have $A_{ij}^{MH}(s, s') \geq A_{ij}^U(s, s')$ with the inequality strict for at least one pair ij . But then $\rho^{MH}(s, s') > \rho^U(s, s')$ and using similar reasoning, $\rho^{MH}(s', s) > \rho^U(s', s)$.

For $P > 1$ the reasoning is exactly the same, except that now we have to determine the spectral radii of $[A(s)]^P[A(s')]^P$ and $[A(s')]^P[A(s)]^P$. But since $A^{MH}(s)$ is an increasing transformation of $A^U(s)$, $[A^{MH}(s)]^P$ will be an increasing transformation of $[A^U(s)]^P$ and the same for state s' . But then the same result as with $P = 1$ obtains. Hence as long as MHA's precision is sufficiently high, $1/\phi > 1/\phi^*$, the modular DP will outperform asymptotically a unitary DP of the same size under environmental cycles of any frequency.

For the final statement in the proposition, note that $\lim_{(1/x) \rightarrow 0} \gamma_j(1/x)J_i^{MH}(s') = 0 < J_i^U(s')$ for $J = G, P$. Hence $\lim_{(1/x) \rightarrow 0} A^{MH}(s') = \mathbf{0}$ and any unitary DP with positive precision will have a higher asymptotic growth factor under any type of cycle. ■

Proof of Proposition 3:

For instance, let a DM with m mutations in each of two modules in the direction of $a_i^{s'}$ be characterized by

$$G_i(m, s) = \xi_G^m G_i(s) \text{ and } P_i(m, s) = \xi_P^m P_i(s),$$

where $\xi_j < 1$ is the fractional reduction in J as a result of moving in the direction of $a_i^{s'}$ under s or in the direction of a_i^s under s' . Recall that $G_i(s | s')/G_i(s) = \gamma_G(1/\phi)$ and $P_i(s | s')/P_i(s) = \gamma_P(1/\phi)$. Let k be such that $\xi_j^k = \gamma_j(1/\phi)$, with $k \ll W/3$. That is, k is the number of steps required in order to completely push the action to $a_i^{s'}$ at every age i . Any further mutations will then be harmful under either environment and can be analyzed in the standard fashion.

When the environmental state switches to s' , however, the values for the above DP with $m \leq k$ mutations in each of two modules will be $J_i(m, s') = \xi_j^{(k-m)} J_i(s')$ for $J = G, P$. Hence $J_i(m, s') > \xi_j^k J_i(s') = J_i(a_i^s, s')$ for $m > 0$, where clearly $J_i(k, s') = J_i(s')$. The combination of imperfect precision and small environmental cycles thus opens the door for different action profiles to coexist in equilibrium and for DPs undertaking suboptimal actions to be present in every age group. For instance, three organisms with 0, $k/2$, and k mutations along two modules in the direction of s' only will have values of $(J_i(s), \xi_j^k J_i(s'))$, $(\xi_j^{k/2} J_i(s), \xi_j^{k/2} J_i(s'))$, and $(\xi_j^k J_i(s), J_i(s'))$ under s and s' respectively. The asymptotic growth factor of the population will again be determined by the matrices $A^j(s)$ and $A^j(s')$, $j = MH, U$; where the population structure vector now enumerates all unordered sequences of possible combinations of mutations, be they beneficial or not. ■

Numerical Example:

Under the specified assumptions, we have to distinguish among 13 types of DPs: (1) no mutations, (2) one mutation, (3) two mutations in the same module, (4) two mutations in different modules, (5) three mutations in the same module, (6) three mutations in two modules with either zero or two mutations in the controller, (7) three mutations in two modules with one mutation in the controller, (8) three mutations in three modules, (9) four mutations in the same module, (10) four mutations occurring two by two, (11) four mutations occurring three by one with either zero or three mutations in the controller, (12) four mutations occurring three by one with one mutation in the controller, and (13) four mutations in three modules. The types differ according to the action that they execute and the probability with which they transition to other types. Specifically, types 1, 2, 3, 5, 9, and 12 always execute action a^s , types 4, 6, 8, 11, and 13 always execute action a^I , and types 7 and 10 always execute action $a^{s'}$.

The projection matrices $A(s)$ and $A(s')$ are simple but tedious to represent fully. Since the population

structure vector \mathbf{N}_t is of dimension 13, the projection matrices $A(s)$ and $A(s')$ are 13×13 . To illustrate, the 3rd rows of $A(s)$ and $A(s')$ represent the inflow of DPs into a DP with two mutations in the same module. Specifically, we have

$$A_3(s) = \begin{bmatrix} 0 & rg/3 & (1-2r)g & 0 & rg & rg/3 & rg/3 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

and

$$A_3(s') = \begin{bmatrix} 0 & rg\xi^2/3 & (1-2r)\xi^2g & 0 & rg\xi^2 & rg\xi^2/3 & rg\xi^2/3 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

where we have assumed that mutations are independent, occur with the same probability, and occur in each module with probability $1/3$.

The asymptotic distribution for a stable environment at s turns out to be given by

$$u(s)^T = \begin{bmatrix} 0.367 & 0.252 & 0.173 & 0.001 & 0.119 & 0 & 0 & 0 & 0.07 & 0 & 0 & 0.018 & 0 \end{bmatrix},$$

where the row may not sum to 1 due to rounding off. The population is clearly dominated by DPs executing action a^s with only 0.001 of the DPs behaving according to a^I . Even though there is conflict among modules for over 70 percent of the population, that conflict does not translate into behavioral heterogeneity.

If the environmental state changes permanently to s' , we have

$$u(s')^T = \begin{bmatrix} 0 & 0 & 0 & 0.001 & 0 & 0.001 & 0.413 & 0 & 0 & 0.583 & 0.001 & 0.001 & 0 \end{bmatrix}.$$

Naturally, the population now is completely dominated by DP types which choose the optimal action $a^{s'}$. Again, there is no heterogeneity in behavior. Finally, under generational cycles between s and s' , we have the projection matrix $A(s')A(s)$ with associated left eigenvector,

$$u(s', s)^T = \begin{bmatrix} 0.143 & 0.131 & 0.109 & 0.108 & 0.078 & 0.077 & 0.076 & 0.077 & 0.041 & 0.04 & 0.04 & 0.041 & 0.04 \end{bmatrix},$$

which clearly demonstrates heterogeneity in both reference models and behavior.

References

- [1] Attwell, D. and S.B. Laughlin (2001): "An Energy Budget for Signaling in the Grey Matter of the Brain," *Journal of Cerebral Blood Flow & Metabolism*, 21, 1133–1145.
- [2] Bechara A. (2005): "Decision making, impulse control and loss of willpower to resist drugs: a neurocognitive perspective," *Nature Neuroscience* 8: 1458-1463.
- [3] Benhabib J. and A. Bisin (2005): "Modeling internal commitment mechanisms and self-control: a neuroeconomic approach to consumption-saving decisions," *Games and Economic Behavior*, 52, 460-490.
- [4] Benhabib J. and A. Bisin (2008): "Choice and process: Theory ahead of measurement," in *The Foundations of Positive and Normative Economics*, Caplin A. and A. Schotter (eds.), New York, Oxford University Press.
- [5] Bernheim B.D. and A. Rangel (2004): "Addiction and cue-triggered decision processes," *American Economic Review* 94, 1558-1590.
- [6] Berridge K. (1999): "Pleasure, pain, desire, and dread: hidden core processes of emotion," in *Well-being: The Foundations of Hedonic Psychology*, D. Kahneman, E. Diener, and N. Schwarz (eds.), New York, Russell Sage Foundation.
- [7] Breedlove S.M., M. Rosenzweig, and N. Watson (2007): *Biological Psychology*. Sunderland, Sinauer Associates.
- [8] Brocas I. and J. Carrillo (2008): "The brain as a hierarchical organization," *American Economic Review*, 98, 1312-1346.
- [9] Cangelosi A., D. Parisi, and S. Nolfi (1994): "Cell division and migration in a "genotype" for neural networks," *Network* 5, 497-515.
- [10] Camerer, C. (2009): "Behavioral economics," World Congress of the Econometric Society, forthcoming.
- [11] Camerer C., G. Loewenstein, and D. Prelec (2005): "Neuroeconomics: How neuroscience can inform economics", *Journal of Economic Literature*, XLIII, 9-64.
- [12] Caplin A. and A. Schotter (2008): *The Foundations of Positive and Normative Economics*, New York, Oxford University Press.

- [13] Dasgupta P. and E. Maskin (2005): "Uncertainty and hyperbolic discounting," *American Economic Review* 95, 1290-1299.
- [14] Frederick S., G. Loewenstein, E. O'Donoghue (2002): "Time discounting and time preference: A critical review," *Journal of Economic Literature*, XL, 351-401.
- [15] Ely, J. (2009): "Kludged," mimeo, Northwestern University
- [16] Fudenberg D. and D. Levine (2006): "A dual self model of impulse control," *American Economic Review*, 96, 1449-1476.
- [17] Gardner E. and J. David (1999): "The neurobiology of chemical addiction," in *Getting Hooked: Rationality and Addiction*, J. Elster and O.J. Skog (eds.), Cambridge, Cambridge University Press.
- [18] Glimcher P. (2003): *Decisions, Uncertainty, and the Brain: The Science of Neuroeconomics*, Cambridge, Ma, MIT Press.
- [19] Glimcher P.W., J.W. Kable, and K. Louie (2007): Neuroeconomic studies of impulsivity: Now or just as soon as possible?, *American Economic Review* 97(2), 142-147.
- [20] Hornik K., M. Stinchcombe, and H. White (1989): "Multilayer Feedforward Networks are Universal Approximators," *Neural Networks*, 2, 359-366.
- [21] Hornik K., M. Stinchcombe, and H. White (1990): "Universal Approximation of an Unknown Mapping and Its Derivatives Using Multilayer Feedforward Networks," *Neural Networks*, 3, 551-560.
- [22] Kerns J., J.D. Cohen, A.W. MacDonald III, R.Y. Cho, V.A. Stenger, and C.S. Carter (2004): "Anterior cingulate conflict monitoring and adjustments in control," *Science* 303, 1023-1026.
- [23] Laibson, D. (1997): "Golden eggs and hyperbolic discounting," *Quarterly Journal of Economics*, 62, 443-77.
- [24] Loewenstein G. and E. O'Donoghue (2007): "Animal spirits: affective and deliberative processes in economic behavior," mimeo, Cornell University.
- [25] Loewenstein G. and D. Prelec (1992): "Anomalies in intertemporal choice: Evidence and interpretation," *Quarterly Journal of Economics*, 107(2), 573-97.

- [26] Luria S.E. and M. Delbrueck (1943): "Mutations of bacteria from virus sensitivity to virus resistance," *Genetics* 28, 491-511.
- [27] McClure, S.M., D. Laibson, G. Loewenstein and J.D. Cohen (2004): "Separate Neural Systems Value Immediate and Delayed Monetary Rewards," *Science* 306, October 15.
- [28] Marcus G. (2008): *Kluge: The haphazard construction of the human mind*, New York, Houghton Mifflin.
- [29] Miller, E. and J.D. Cohen (2001): "An integrative theory of prefrontal cortex function," *Annual Review of Neuroscience* 24, 167-202.
- [30] Netzer, N. (2009): "Evolution of time preferences and attitudes towards risk," *American Economic Review* , forthcoming.
- [31] Rabin, M. (1998): "Psychology and economics," *Journal of Economic Literature*, 36(1), 11-46.
- [32] Raichle M.E. and D.A. Gusnard (2002): "Appraising the brain's energy budget," *Proceedings of the National Academy of the Sciences*, August 6, 99(16), 10237-10239.
- [33] Robson A., L. Samuelson (2009): "The evolutionary foundations of preferences," in *The Social Economics Handbook*, J. Benhabib, A. Bisin, and M. Jackson (eds.), Elsevier Press, forthcoming.
- [34] Seneta E. (1981): *Non-Negative Matrices and Markov Chains*, New York, Springer Verlag.
- [35] Rubinstein A. and Y. Salant (2008): "Some thoughts on the principle of revealed preference," in *The Foundations of Positive and Normative Economics*, Caplin A. and A. Schotter (eds.), New York, Oxford University Press.
- [36] Simon H.A. (1996): *The Sciences of the Artificial*. Cambridge, MA, MIT Press.
- [37] Thaler R. and H. Shefrin (1981): "An economic theory of self-control," *Journal of Political Economy*, 89, 392-406.
- [38] Varga R. (2000): *Matrix Iterative Analysis*, New York, Springer Verlag.