

Evolutionary Selection of Modular Decision Architectures

Alberto Bisin*

Emil Iantchev[†]

New York University and NBER

Syracuse University

May 19, 2010

Abstract

We study the evolutionary properties of decision processes. We show that a population of agents possessing decision architectures with hierarchically organized modules will have a strictly higher asymptotic growth factor than a population of agents with unitary decision architectures in which the modules are fully connected. Furthermore, we show that internal conflict within agents and behavioral heterogeneity across agents are properties of evolutionary equilibrium. We interpret these results as supporting economic models of multiple decision processes like e.g., *planner-doer* models.

1 Introduction

In this paper we investigate the evolutionary properties of two stylized architectures for decision processes: i) a *modular hierarchical (MH) architecture* in which a *motivational* module discriminates between the action recommendations of two separate *advisory* modules

*Department of Economics, NYU and NBER. E-mail: alberto.bisin@nyu.edu

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

and ii) a *unitary (U) architecture* in which all three modules are directly connected with each other so that one circuit always determines behavior.

Modular hierarchical decision processes are of interest as formal representations of multiple decision processes, like those recently studied by economists in choice 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 models characterized by the interaction of *multiple selves*, with conflicting models of behavior.²

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 module 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. And even though internal conflict will be present in equilibrium in a large part of the population, it will not lead to behavioral heterogeneity except for a trivial fraction of the agents.

We also compare the performance of the two architectures in the presence of small random 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 architecture might expend more energy while being able to respond more precisely

¹See for instance the surveys in Rabin (1998), Loewenstein 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).

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).

to fluctuations in the environment. Under these conditions, and in the presence of harmful mutations, the relative performance of the hierarchical 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 with the same number of modules 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 random fluctuations in the environment, both internal conflict among advisory modules and heterogeneity in behavior among agents in the population are integral parts of evolutionary equilibrium. Intuitively, in the presence of undetectable fluctuations in the environmental state, the behavioral heterogeneity induced by 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 an undetected environmental perturbation.

We interpret our results as providing an evolutionary justification for the models of multiple decision processes explored by economists, and in particular for hierarchical models of *multiple selves* such as in the work of Thaler and Shefrin (1981) and, more recently, Brocas and Carrillo (2008) and others.³ Because the modular hierarchies we study are a type of neural network and can serve as universal approximators, they can provide a common conceptual foundation for all multiple selves models, providing much needed theoretical restrictions *across* decision problems to be identified with decision theoretic and/or neurobiological data.

1.1 Related literature

This paper is related to *indirect evolutionary theory*, whose aim is to identify the fundamental characteristics of agents' preferences as the result of evolutionary selection. Robson and Samuelson (2009) comprehensively survey this literature. In this paper, however, we subject

³For a comprehensive review of the unitary vs. multiple decision systems debate within neuroeconomics see Rustichini (2008).

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 5, several of the recent decision theoretic analyses of intertemporal choice (can be interpreted to) develop explicit models of decision processes.

Three recent papers, Dasgupta and Maskin (2005), Samuelson and Swinkels (2006), and Netzer (2009), are closely related to ours in that they also concentrate on the evolution of multiple selves models. In Dasgupta and Maskin (2005) and Netzer (2009) multiple selves arise in response to specific choice problems that are believed to have been encountered by human populations during their evolutionary past. On the contrary, in our paper, hierarchical decision processes arise regardless of the choice problems the decision maker might have encountered, suggesting that the properties of decision architectures might be relatively robust across decision problems. In addition, in those two papers, no internal conflicts or behavioral heterogeneity arise in equilibrium so those which are observed in modern times are explained as a remnant of our evolutionary past. In our paper, on the other hand, internal conflicts and behavioral heterogeneity have a positive evolutionary value because they serve as diversification devices in the presence of undetectable fluctuations in the environment.

Internal conflicts also turn out to have a positive evolutionary value in Samuelson and Swinkels (2006). However, Samuelson and Swinkels (2006)'s analysis crucially depends on the imperfect observability of the environmental state due to cognitive limitations. In our paper, on the other hand, internal conflicts might be present in equilibrium even if the evolutionary environment is stable and thus perfectly observable (Proposition 1). Furthermore, the internal conflicts derived by Samuelson and Swinkels do not lead to behavioral heterogeneity. In their environment, conditionally on being in the same situation, all agents will make exactly the same choice.⁵ In our paper, instead, heterogeneity in behavioral responses

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

⁵One can argue that response heterogeneity can be introduced in the model by endowing agents with different information-processing rules, say due to differences in cognition. But then why wouldn't only the

is a crucial aspect of evolutionary selection (Proposition 3). In fact, internal conflicts have a positive evolutionary value under imperfect observability precisely because they facilitate behavioral heterogeneity, which in turn leads to demographic diversification.

Our decision model is supported by available evidence regarding the neurophysiology of the brain and is based on biologically plausible mutation processes. However, it is easier to discuss the biological foundations of our analysis, and to provide the appropriate references, after having developed it. We report on biological foundations in Section 4.

2 Decision-making architectures

We consider decision-making architectures as embedded in an age-structured population of agents with reproductive lives 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. Let $\mathbf{s} = [s(i)]_{i=1}^n \in S$ denote the age profile of the states of the evolutionary environment. The evolutionary environment could be deterministic or random, in which case \mathbf{s} is a stochastic process.⁶

A *decision-making architecture* (DA) executes an action \mathbf{a} from a given set A , for any state \mathbf{s} .⁷ Formally, we think of an action as a map $\mathbf{a} : S \rightarrow A$.

Let \mathbf{a}^s denote the action profile with the lowest asymptotic fitness loss (relative to the maximum), given \mathbf{s} . 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 and A for each age. In the MH architecture, multiple, nearly decomposable, decision modules interact to produce a map between S and A .

agents with the best information-processing rule survive natural selection?

⁶Saving on notation, we simply mention that \mathbf{s} is, formally, an element of some metric space S endowed with a norm d .

⁷Note that the unit of analysis in our context is an architecture rather than an agent, a decision-making architecture and not a decision maker.




Figure 1: *Example of a modular hierarchy (MHA) vs unitary (UNITARY) decision architecture. MHA can be obtained from UNITARY by deleting the direct connection $3 \longleftrightarrow 7$ between the modules $\{1,2,3,4\}$ and $\{5,6,7\}$.*

Figure 1 provides an illustration of the difference between the two types of decision architectures. In both decision architectures there are three modules, defined as subgraphs with densely connected nodes, which are sparsely interconnected.⁸ In the unitary architecture, every module is connected to every other module and all three interact with each other in order to produce a recommended response. More importantly, a change in a particular module, *ceteris paribus*, will always influence the architecture’s overall response since that module is connected to every other module. On the other hand, in the MH architecture modules $\{1,2,3,4\}$ and $\{5,6,7\}$ are not directly connected with each other. Hence each of the two is able to produce an independent response recommendation. Module $\{8,9,10,11,12\}$, however, is directly connected with both advisory modules. It can thus serve as a motivational unit that aggregates all information–input signals and the recommended policy by each advisory module—and initiates the execution of one of the recommendations. The difference here is that there are direct connections among nodes within a module on a given level of the hi-

⁸For the formal definition of a module in particular, based on local efficiency in information transmission (clustering), and "economic small-world" networks in general, see the paper by Latora and Marchiori (2003).

erarchy but not across modules on the same level—they interact through the higher level motivational module. The *MH* architecture is thus nearly decomposable in the sense that rewiring within a particular module, ceteris paribus, need not have global implications for behavior.⁹

Next, we formally describe the decision process of each architecture.

Unitary Architecture. Let $\mathbf{a}^U(\mathbf{s})$ denote the hereditarily acquired action of a unitary architecture in environment \mathbf{s} . In the long run, DAs that have the correct reference model, \mathbf{a}^s , for the environmental state \mathbf{s} , will dominate the population.

Modular Hierarchical Architecture. Let a DA consist of two stage-one (advisory) modules (modules 1 and 1') and one stage-two (motivational) module (module 2). Each stage-one module has its own reference choice profile, namely $\mathbf{a}_1(\mathbf{s})$ and $\mathbf{a}_{1'}(\mathbf{s})$. The same is true for module 2, which has a reference profile given by $\mathbf{a}_2(\mathbf{s})$.

Module 2 aggregates the information from the two advisory modules and executes the recommendation of either module 1 or 1'. We postulate that module 2 decides in favor of the advisory module whose recommendation is closer to its own reference action. Formally, let $\mathbf{w} : A \times A \rightarrow \mathbb{R}$ denote a distance map between action profiles. The action profile resulting from the modular hierarchical architecture, denoted $\mathbf{a}^{MH}(\mathbf{s})$, is defined by:

$$\mathbf{a}^{MH}(\mathbf{s}) \in \arg \max_{\mathbf{a} \in \{\mathbf{a}_1(\mathbf{s}), \mathbf{a}_{1'}(\mathbf{s})\}} -\mathbf{w}(\mathbf{a}, \mathbf{a}_2(\mathbf{s})).$$

The asymptotic fitness loss will be minimal for any DA with $\mathbf{a}^{MH}(\mathbf{s}) = \mathbf{a}^s$. This is achieved if:

- (i) $\mathbf{a}^s = \mathbf{a}_1(\mathbf{s}) = \mathbf{a}_{1'}(\mathbf{s})$, for any feasible profile $\mathbf{a}_2(\mathbf{s})$;
- (ii) $\mathbf{a}^s = \mathbf{a}_j(\mathbf{s})$ and $\mathbf{w}(\mathbf{a}_j(\mathbf{s}), \mathbf{a}_2(\mathbf{s})) \leq \mathbf{w}(\mathbf{a}_k(\mathbf{s}), \mathbf{a}_2(\mathbf{s}))$, for $j = 1$ or $1'$ and $k \neq j$.

Hence for the modular hierarchy to survive in a stable environment, it is needed that at least one of the module 1 actions and the module 2 action be close to the fitness minimizing

⁹For a thorough discussion of nearly decomposable architectures see Simon (1996).

action.

Any decision architecture $K = MH, U$ might implicitly impose specific restrictions on the set of action maps it can support. We shall consider in particular measurability restrictions, so that a decision architecture $K = MH, U$ is characterized by a measure of precision of the actions it executes. Formally, for any two environmental states \mathbf{s}, \mathbf{s}' , let $d(\mathbf{s}, \mathbf{s}')$ denote a measure of the distance between them. We then say that architecture K 's precision is $\frac{1}{\phi^K}$ if it can support an action map \mathbf{a}^K such that $\mathbf{a}^K(\mathbf{s}) \neq \mathbf{a}^K(\mathbf{s}')$ for any $\mathbf{s}, \mathbf{s}' \in S$ such that $d(\mathbf{s}, \mathbf{s}') \geq \phi^K$.¹⁰

3 Evolutionary selection

In order to determine which architecture is more likely to be evolutionary successful, we subject them to mutations. A mutation is a change in the connection weights of a node and thus changes the reference action profile of the module in which it occurs. By construction, any architecture subjected to no mutation executes the loss minimizing (fitness maximizing) action \mathbf{a}^s . But all nodes in an architecture are subject to a harmful mutation in each period i with probability $q > 0$ and to a reversal (a mutation in the direction of the loss minimizing action), with probability $r > 0$. Since mutations are rare events, we assume that q is sufficiently small so that multiple mutations per period can be ignored. In addition, since the vast majority of mutations are harmful (there are many more ways to increase than decrease the fitness loss), we assume $r \leq q$. Mutations are inherited and are generally associated with a node, not a Decision Architecture. Let $m = [m_1, m_{1'}, m_2]$ denote the vector of *harmful mutations accumulated by DA in the groups of nodes associated with MH modules 1, 1', and 2*. This distinction is important for the MH architecture but is irrelevant for the U architecture where only the total number of mutations $m_1 + m_{1'} + m_2$ matters. For ease

¹⁰Essentially, 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.

We spare the reader the precise definition in terms of the measurability of the map \mathbf{a}^K .

of comparability, however, we will also keep track of the distribution of mutations over the groups in the U architecture, even though they are fully interconnected.

Some *mutations* are *influential*, i.e., they change the architecture's action, while other are *neutral*. The number of influential mutations will depend on the decision architecture. In the U case, for instance, every mutation will be influential since modules are interconnected. This will not be the case for the MH architecture, where the number of influential mutations must be less than or equal to $\max\{m_1, m_{1'}, m_2\}$.

In this context we study the evolutionary selection of unitary *vs.* modular hierarchical architectures i) under a stable environment and ii) in the presence of small random fluctuations in the environment. The architecture which has higher fitness, that is higher asymptotic growth factor, under these conditions will be more likely to prevail in the long run. Each architecture's asymptotic growth factor will depend on the agent's associated survival probability and fertility rate. An architecture's asymptotic growth factor will also depend on its *precision*, that is, its ability to recognize and react to different environmental states.

More precisely, consider an agent of age i with decision architecture $K = U, MH$ in environmental state $\mathbf{s} = [s(i)]_{i=1}^n$. We can now define the fundamental concepts of the evolutionary analysis:

- i) $\mu^K(m)$ is an influence function that determines the number of mutations in the vector m that are influential (NOT neutral) for the DA's survival and procreation.
- ii) $P_i^K(\mu^K(m), \mathbf{s})$ denotes the survival probability of an agent of age i with decision architecture $K = U, MH$ with m accumulated harmful mutations in state \mathbf{s} . Consistently, we require that $P_n^K(\mu^K(m), \mathbf{s}) = 0$. We also assume a DA can accumulate at most W harmful mutations, as $W + 1$ mutations are lethal: $P_i^K(\mu^K(m), \mathbf{s}) = 0$, for $m_1 + m_{1'} + m_2 \geq W + 1$.
- iii) $G_i^K(\mu^K(m), \mathbf{s})$ denotes the expected number of surviving offsprings from an agent of age

i with decision architecture $K = U, MH$ with m accumulated harmful mutations in state \mathbf{s} .¹¹

Furthermore, recall that

- iv) $\frac{1}{\phi^K}$ denotes the precision of architecture K , that is, its ability to support an action map \mathbf{a}^K such that $\mathbf{a}^K(\mathbf{s}) \neq \mathbf{a}^K(\mathbf{s}')$ for any $\mathbf{s}, \mathbf{s}' \in S$ such that $d(\mathbf{s}, \mathbf{s}') \geq \phi^K$.

Assumption 1. *Influential mutations are harmful:*

$$P_i^K(\mu, s) \text{ and } G_i^K(\mu, s) \text{ are strictly decreasing in } \mu,$$

while

$$\mu^K(m) \text{ is weakly increasing in (each element of) } m.$$

Furthermore, the following complexity-efficiency trade-off determines each architecture's relative fitness properties: the MH architecture is more energy efficient

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

while, on the other hand, the U architecture is more precise

$$\frac{1}{\phi^U} \geq \frac{1}{\phi^{MH}}.^{12}$$

To ignore the complicating effects of sexual reproduction, we assume:

Assumption 2. *Mating is assortative in the sense that like types match with each other in order to reproduce.*¹³

¹¹The expectation here is taken over some distribution of idiosyncratic individual shocks which average out in a sufficiently large population.

¹²We say the the complexity-efficiency trade-off is non-trivial if the inequalities hold strictly.

¹³Assortative mating may be violated, especially among types which do not exhibit any variation in behavior. In the appendix, we consider the opposite extreme assumption: random mating.

3.1 Stable environment

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

Assumption 3-s. *The evolutionary environment is stable, that is $s(i)$ does not vary over the lifetime $i = 1, \dots, n$ of any agent:*

$$s(i) = s, \forall i.$$

We can then drop reference to the environmental state \mathbf{s} in the notation of this section, without loss of generality. Consider first the *MH* architecture. Recall that such an architecture is composed of three modules, 1, 1', and 2. 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, respectively, in module 1, 1', and 2. The distribution of mutations across modules is important: for example, a mutation in module 1 or 1' might have no effect on the action executed. The age and mutations profile of the population at the beginning of period t is given by

$$\mathbf{N}_t = \left[N_{i,t}^{j,k,l} \right]_{i,j,k,l}.^{14}$$

Let

$$\mathbf{G}^K(j, k, l) = [G_1^K(\mu^K(j, k, l)), \dots, G_n^K(\mu^K(j, k, l))]$$

and define $P_i^K(j, k, l) = P_i^K(\mu^K(j, k, l))$. The dynamics of the population of K processes are governed by the linear system

¹⁴E.g., in the case of two age classes (two period life), and two maximum mutations ($n = W = 2$):

$$\mathbf{N}_t = [N_{1,t}^{0,0,0}, N_{2,t}^{0,0,0}, N_{1,t}^{1,0,0}, N_{2,t}^{1,0,0}, N_{1,t}^{1,1,0}, N_{2,t}^{1,1,0}, N_{1,t}^{2,0,0}, N_{2,t}^{2,0,0}]^T$$

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

where the generic entry of A^K is constructed as follows:

For $m = (j, k, l) \neq 0$ and $j + k + l < W$, the inflow from N_t into $N_{1,t+1}^{j,k,l}$ is given by the following row of A^K

$$[0 \dots 0, q\mathbf{G}^K(j-1, k, l), q\mathbf{G}^K(j, k-1, l), q\mathbf{G}^K(j, k, l-1), (1-q-r)\mathbf{G}^K(j, k, l), \\ r\mathbf{G}^K(j+1, k, l), r\mathbf{G}^K(j, k+1, l), r\mathbf{G}^K(j, k, l+1), 0 \dots 0],$$

while for $i > 1$, the inflow from N_t into $N_{i,t+1}^{j,k,l}$ is given by the row

$$[0 \dots 0, qP_{i-1}^K(j-1, k, l), qP_{i-1}^K(j, k-1, l), qP_{i-1}^K(j, k, l-1), (1-q-r)P_i^K(j, k, l), \\ rP_{i-1}^K(j+1, k, l), rP_{i-1}^K(j, k+1, l), rP_{i-1}^K(j, k, l+1), 0 \dots 0].$$

For $m = \mathbf{0}$ and $j + k + l = W$, the entries are slightly different but the basic idea is the same. The interested reader is referred to the Appendix where the entire matrices A^{MH} and A^U are explicitly constructed for the case $n = W = 2$. Note that in a stable environment, the advantage of the U architecture in terms of precision (Assumption 1) has no effect on fitness and hence on the dynamics of \mathbf{N}_t . The MH architecture, therefore, cannot do any worse than the U architecture. In fact it does strictly better, as summarized in Proposition 1.

Proposition 1. *In a stable environment (Assumption 3-s), the MH architecture has a strictly higher asymptotic population growth factor than the U architecture, starting from any non-zero initial condition. Furthermore, this is the case even if MH has no energy advantage*

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

Sketch of the proof: Consider without loss of generality the case in which $G_i^{MH}(\mu) =$

$G_i^U(\mu) = G_i(\mu)$ and $P_i^{MH}(\mu) = P_i^U(\mu) = P_i(\mu)$, where again μ denotes the number of influential mutations. The dynamics of the two populations, 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 non-negative 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. We show in the Appendix that the spectral radius of A^{MH} is larger than the spectral radius of A^U . ■

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 interconnected. Hence the accumulation of mutations in any part of the decision architecture will influence the DA's overall course of action. This is not true for the modular hierarchy. 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 DA's course of action. Thus for significant fitness loss to occur, we need multiple changes in the MH architecture 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 DA's evolutionary dynamics.

3.2 Randomly fluctuating environment

In this section, we consider small but non-trivial fluctuations in the environment. The only interesting environment is one where the precisions of the $K = U, MH$ architectures, $1/\phi^K$, are restricted to be such that the U architecture can adjust its action profile in response to the fluctuating environment while the MH cannot. We proceed with studying this case.

Assumption 3-f. *The environmental state switches from s to s' and from s' to s according to some ergodic process, where $\phi^U \leq d(s, s') = \phi^{MH} - \epsilon$. Furthermore, these*

fluctuations are non-trivial in the sense that $G_i^K(\mu, \mathbf{s})$ and $P_i^K(\mu, \mathbf{s})$ are not constant in \mathbf{s} for some i and m .

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 $G_i^K(\mu, s' | s)$ denote the number of offsprings as a result of an action recommended by a process with μ influential mutations relative to the optimal action under s , when the actual state is s' . Also, let $P_i^K(\mu, s' | s)$ denote the survival probability as a result of such an action. More precisely, a modular DA executing action $a(s)$ under state s' will feature

$$G_i^{MH}(\mu, s' | s) = \gamma_\mu^G \left(\frac{1}{\phi^{MH}} \right) G_i^{MH}(\mu, s') \text{ and } P_i^{MH}(\mu, s' | s) = \gamma_\mu^P \left(\frac{1}{\phi^{MH}} \right) P_i^{MH}(\mu, s'),$$

where $P_i^{MH}(\mu, s')$, for instance, is the survival probability as a result of a recommendation by a process with μ influential mutations relative to the optimum under s' . The functions $\gamma_\mu^J(1/x)$ are continuous, non-increasing in x and bounded in the interval $[0, 1]$ with $\gamma_\mu^J(0) = 0$ and $\lim_{x \downarrow d(s, s')} \gamma_\mu^J(1/x) = 1$ for $J = G, P$ and $\mu = 1, 2, \dots, W$. The condition that $\gamma_\mu^J(1/x)$ are non-increasing in x corresponds to the assumption that ignoring larger fluctuations in the environment leads to proportionately larger reductions in offsprings and survival. For instance, as $1/\phi^{MH} \rightarrow 0$, the MH architecture will fail to adjust its behavior to arbitrarily large fluctuations in the environment, which we assume will have disastrous consequences for offsprings and survival.

Because of its greater complexity, a Unitary DA suffers no reduction in fitness as a result of the environmental fluctuations, while a MH does. The unitary DA, 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 a randomly fluctuating environment (Assumption 3-f), if the preci-*

sion $\frac{1}{\phi^{MH}}$ 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 any ergodic process 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: In a random environment, the asymptotic growth of a dynamical system is determined by its dominant Lyapunov exponent. Instead of spectral radii we thus have to compare the Lyapunov exponents of the two choice architectures. The proof uses the strong law of large numbers for products of non-negative random matrices derived by Hennion (1997). Hennion shows that the sequence of spectral radii generated by the product matrices of a sequence of states converges, in the limit, to the dominant Lyapunov exponent almost surely. We then show that if $1/\phi^{MH}$ is above a critical value, for any sequence of environmental states, the associated spectral radii for the products of *MH* projection matrices are greater than the spectral radii for the products of the *U* projection matrices. Hence the *MH* architecture must have a larger Lyapunov exponent than the *U* system. 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 further increasing precision by moving to a unitary architecture is outweighed by the marginal cost of increased energy consumption and increased susceptibility to the effects of mutations. For relatively imprecise *MH* architectures, the opposite obtains.¹⁵

¹⁵Proposition 2 suggests that modular decision architectures should be observed later in evolutionary history, assuming encephalization increases with time. This is consistent with evidence from neurobiology. For example, Allman et al. (2002) and Hof et al. (2001) report the presence of two distinct types of large neurons in the anterior cingulate cortex that are unique to apes and humans. These neurons facilitate information transfer and their volume is positively correlated with encephalization. This supports the notion that the ACC is a recent evolutionary specialization of the cortex that probably originated in the common ancestor of humans and great apes to monitor and facilitate the interaction between emotion and cognition.

Remark 1 *As part of the proof of Proposition 2 we show that the same result as in Proposition 2 also applies to a deterministic but cyclically fluctuating environment. More precisely, it applies to an environment in which the state switches from s to s' and from s' to s after every $P \geq 1$ and $P' \geq 1$ periods respectively.*

3.2.1 Random fluctuations and equilibrium heterogeneity in action profiles

Under a stable environment, the population will be dominated by DAs 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 such outcomes as an evolutionary equilibrium.

Propositions 1 and 2 ignore the possibility of mutations being beneficial. In a randomly fluctuating environment however, under our assumptions the MH architecture chooses the action which is optimal for state s even when the environmental state is s' . Beneficial mutations can in principle exist for this architecture, which will reduce the costs associated with the lack of precision. We expect such mutations to occur with a non-zero probability less than or equal to q .

Assumption 4. *Suppose for some i and $W \geq \mu^* > 0$, we have $G_i^{MH}(\mu^*, s' | s) > G_i^{MH}(0, s' | s)$ and $P_i^{MH}(\mu^*, s' | s) > P_i^{MH}(0, s' | s)$, while $G_i^{MH}(\mu^*, s) < G_i^{MH}(0, s)$ and $P_i^{MH}(\mu^*, s) < P_i^{MH}(0, s)$.*

The fitness of a DA will then depend on the distribution of mutations over its modules.

Proposition 3. *In a randomly fluctuating environment (Assumption 3-f), and in the presence of beneficial mutations (Assumption 4), evolutionary selection induces (i) a non-trivial fraction of the population with advisory modules that have conflicting models of optimal behavior and (ii) a non-trivial number of DAs that choose actions which are suboptimal with respect to the current environmental state.*

Proof: Clearly now the relative number of DAs with advisory processes with μ^* mutations relative to the optimum under state s will grow when the environment switches to

s' . On the other hand, when the environment switches back to s , the relative number of DAs with advisory processes that recommend the optimal action under s will grow. Thus as long as the environment continues to fluctuate, no single type will completely dominate the asymptotic population distribution of MHAs. ■

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, mis-specified model in the event of an undetected environmental perturbation. It turns out that the degree of behavioral heterogeneity present in equilibrium depends on the persistence of the process determining the environmental state. The more persistent the environmental state is, the smaller the role that behavioral heterogeneity plays in an evolutionary equilibrium. We illustrate the importance of the environmental process with the following simple example.

Numerical example Consider a *MH* architecture with precision $1/\phi$ whose reproductive life lasts one period. We consider two extreme scenarios as far as the persistence of the environmental state is concerned. In the first scenario, the environment start in state s and stays there forever. In the other scenario, we consider the case of no persistence. Specifically, we start in environmental state s but after every generation, the state switches between s and s' and vice versa with probability 1. Let $d(s, s') < \phi$. Let a^s denote the optimal action in environmental state s and suppose a DA can carry on at most 4 deleterious mutations relative to a^s . A DA with 5 mutations leaves no offspring in either state of nature¹⁶. Assume that a mutation or its reversal happen with probability $q = r = 0.0001$ per generation. The set of possible actions is given by $\{a^s, a^1, a^{s'}, a^3\}$, where we have ordered the sequence by the number of influential deleterious mutations relative to $a^s - 0, 1, 2,$ and 3 respectively. That is, mutations which occur in the module whose recommendation is executed.

¹⁶All of the the assumptions in the example are made for tractability, since the complexity of the problem increases very quickly with n and W .

Let $w(a, a')$ denote the distance map across actions.¹⁷ Assume that: (i) $w(a^I, a^{s'}) < w(a^I, a^s) < w(a^I, a^3) < w(a^s, a^3)$; (ii) $G(\mu, s) = \xi^\mu g$, while $G(\mu, s' | s) = \xi^{2-\mu} g$ for $\mu = 0, 1, 2$ with $\xi = 0.98$ and $g = 1.03$. Condition (i) insures that a^3 will never be executed and a supervisor with reference action a^I or $a^{s'}$ will execute a recommended action of a^I or $a^{s'}$ rather than a^s . Condition (ii) says that mutations initially push the action in the direction of $a^{s'}$, which is the optimal action under 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. Hence in a modular DA, two pushes require at least four mutations, two at any two of the three modules.

For the cyclically fluctuating environment, we have the following law of motion, where the unit of account c is a cycle that lasts two generations,

$$\mathbf{N}_{c+1} = [A(s')A(s)]\mathbf{N}_c.$$

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 DA 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 population over the possible actions $\{a^s, a^I, a^{s'}\}$ under the two scenarios: (i) a stable environment at s and (ii) a generational cycle between s and s' . As expected, under scenario (i), 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 almost 2/3 of population, such internal conflicts do not translate into behavioral heterogeneity. This is because only

¹⁷That is, a component of $\mathbf{w}(\mathbf{a}, \mathbf{a}')$, the distance map across action profiles.

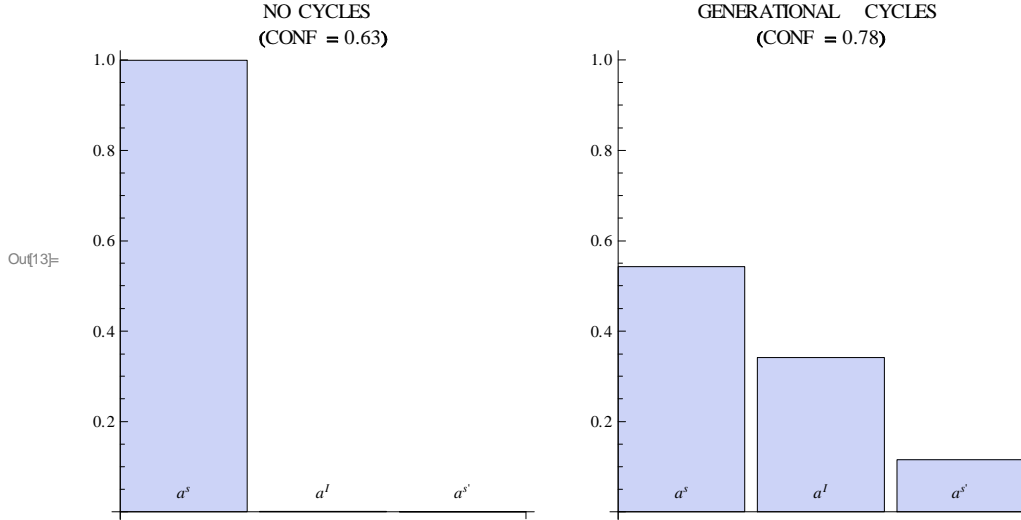


Figure 2: *Asymptotic distribution of the population over the possible actions $\{a^s, a^I, a^{s'}\}$ under two environments. $CONF$ gives the equilibrium fraction of the population that faces internal conflicts.*

mutations that are fitness neutral are allowed to accumulate in the long run.

Under a fluctuating environment, on the other hand, we have heterogeneity in both reference models and actual behavior.¹⁸ In every generation, between 45 and 90 percent of the population choose a suboptimal course of action. Even more interestingly, there are DAs—such as the type with three total mutations distributed across two modules—which possess the correct reference 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 environmental state changes.

¹⁸The 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 at $a^{s'}$, 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 Neurobiological foundations

Our model has a structure that is explicitly based on the neurophysiology of the brain and a mutation process that is biologically plausible. By subjecting decision processes to evolutionary selection, we emphasize local effects of heritable 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 fact that many cognitive and behavioral disorders have a prominent heritable component and that several disorders, in addition to being strongly heritable, are also associated with explicit structural and functional changes in certain brain regions.¹⁹ The interested reader should refer to chapters 7 and 16 in Breedlove, et al. (2007) for an introduction to the subject.

In addition, it is important for our analysis that there exist heritable mutations whose effects are localized only in certain regions of the brain. This is consistent with recent work in neurogenetics. First, there are heritable conditions such as synesthesia²⁰ that manifest in increased connectivity only in certain regions of the brain, such as the fusiform gyrus for grapheme-color synesthesia, as opposed to globally. Second, there is growing evidence that gene expression—the translation of the genetic blueprint into a functional product—varies significantly across brain regions and is controlled by other genes²¹. Our mutation process can thus be interpreted as modeling heritable mutations in the genes regulating the expression and hence the contribution to the reference action profile of the genes located

¹⁹Down, 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.

²⁰Synesthesia is a condition where stimulation of one sensory or cognitive pathway leads to automatic, involuntary experiences in a second sensory or cognitive pathway. In the most common type, grapheme-color synesthesia, individuals perceive letters and numbers as inherently colored. For an investigation into the neural basis of synesthesia, see Ramachandran and Hubbard (2001).

²¹Nadler et al. (2006) show that there are large differences in gene expression among different brain regions in various strains of inbred mice, while Monks et al. (2004) show that there is significant genetic inheritance of gene expression in human cell lines. More importantly, Meyer-Lindenberg (2009) shows that variation in genes responsible for serotonin transportation and serotonin catabolism influence the risk of anxiety/depression and impulsive behavior respectively, through the influence these genes have on the neural connectivity among three particular regions—the ventromedial prefrontal cortex, the cingulate cortex, and the amygdala.

in the various modules of the brain. Finally, our assumption that multiple mutations per period can be ignored is consistent with the standard model of Luria and Delbrueck (1943) where mutations are spontaneous rare events independent of the environment.

Our interpretation of the MH architecture as a model of the brain is rooted in theory and evidence from the neurosciences,²² including the view that the human brain is a modular hierarchy organized—using the terminology of Latora and Marchiori (2003)—as an "economic small-world" network.²³ For example, modular hierarchies can represent the interaction between decision units that are based on emotion versus cognition and a motivational unit that monitors and discriminates among their recommendations. In a series of articles—Allman et al (2001, 2002)—Allman, Hakeem, and colleagues have argued that the anterior cingulate cortex (ACC) may play such a monitoring and motivational role as the primary interface between emotion and cognition. Citing evidence from EEG, PET, fMRI, and lesion studies, they argue that the ACC is a recent evolutionary specialization crucial for many aspects of intelligent behavior, including but not limited to aggregating information, monitoring cognitive conflict, focusing attention, discriminating among conflicting informational cues, and initiating the desire to act.²⁴ 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).²⁵

In particular, we model modular architectures as nearly decomposable, and restrict our analysis to the simplest modular hierarchy consisting of two advisory and one motivational

²²See Glimcher (2003) and Camerer, Loewenstein, and Prelec (2005) for general surveys of the field.

²³These are networks which are energy efficient (relative to the ideal, fully connected, small-world network) yet have relatively high local and global efficiency in transmitting information. See Bassett and Bullmore (2006) for a recent review of "small-world" brain networks. Anatomical evidence on the "small-world" modular hierarchical organization of the human brain is presented in He, Chen, and Evans (2007), Chen et al. (2008), and Bassett et al. (2008).

A formal definition of a module, based on local efficiency in information transmission (clustering), is in Latora and Marchiori (2003).

²⁴Bush, Luu, and Posner (2000) provide a comprehensive review, circa 2000, of the anatomical and functional properties of ACC. Specifically, for the striking role ACC plays in volition see Damasio and Van Hoesen (1983), while for the case of monitoring cognitive conflict see Kerns, Cohen, MacDonald III, Cho, Stenger, and Carter (2004).

²⁵But see Glimcher, Kable, and Louie (2007) for a skeptical view of these results.

process. In our set-up, each advisory process recommends a feasible reference course of action. The motivational module 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 motivational unit evolves to monitor internal and external states and execute one of the actions recommended by the advisory processes. The results in that literature are almost exclusively numerical in nature as they are derived through simulations. We thus contribute to that literature by providing an analytical foundation for why this type of choice architecture is likely to survive.

Finally, our assumptions regarding the implications of the different levels of connectedness for survival and precision are consistent with the evidence in the literature on energy budgeting of the human brain.²⁶ Consider for instance our assumption that the unitary architecture is more precise. This is consistent with the universal approximation results by Hornik, Stinchcombe, and White (1989, 1990), since, by construction, the unitary architecture allows for a greater number of nonlinear transformations of any input and hence has the ability to approximate more precisely the loss minimizing policy with respect to small perturbations in the environmental state. On the other hand, the assumption that, given the same action profile and the same initial amount of resources, the unitary architecture will consume more energy is consistent with Attwell and Laughlin (2001), who estimate that over 3/4 of the brain's energy consumption is used for signalling among neurons. For a given number of modules in a network, an architecture like U with a large number of interconnections among those modules 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 modules.

²⁶For a summary of recent studies on the energy budget of the human brain see Raichle and Gusnard (2002) and the references therein.

5 Hierarchical decision processes in economics

Our results provide an evolutionary justification for the models of *multiple selves* recently explored by economists and neuroeconomists. 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".²⁷

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 architecture 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 and 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 and 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

²⁷See also Fudenberg and Levine (2006) and Loewenstein and O'Donoghue (2007) for recent extensions.

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

$$\begin{aligned} & \max_{(\theta_1, \dots, \theta_T) \geq \mathbf{0}} V(Z_1(c_1(\theta_1), \theta_1), \dots, Z_T(c_T(\theta_T), \theta_T)) \\ & \text{s.t.} \quad \sum_t c_t(\theta_t) \leq y. \end{aligned}$$

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 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 $i = 1, 1'$.²⁸ That is, for each advisory 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 supervisor), is given by

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

²⁸A natural interpretation of this formulation of the level 1 modules is that they represent different self-control levels.

where $c(\theta)$ is the policy associated with the objective $Z(c, \theta)$. Then the DA 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), Benhabib and Bisin (2005), and Brocas and Carrillo (2008) have developed models of multiple decision processes that have 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 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 upon request.

6 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 randomly fluctuating environment, conflict among decision modules appears to have a positive evolutionary value, as it is a source of diversification against the influences of undetectable fluctuations in the environment.

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 of certain behaviors. We believe exploring any of these will improve our understanding of decision processes and the behavior of decision makers who possess them.

References

- [1] Allman, J.M., Hakeem, A., Erwin, J.M., Nimchinsky, E., and Hof, P. (2001). "The Anterior Cingulate Cortex: The Evolution of an Interface between Emotion and Cognition", *Annals of the New York Academy of Sciences*. 935:107-17.
- [2] Allman, J.M., Hakeem, A., and Watson, K. (2002) "Two phylogenetic specializations in the human brain." *Neuroscientist* 8(4):335-46.
- [3] 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.
- [4] Bassett, D and Bullmore, E. (2006). "Small-world brain networks," *Neuroscientist* 12: 512-524.
- [5] Bassett et al. (2008). "Hierarchical Organization of Human Cortical Networks in Health and Schizophrenia," *Journal of Neuroscience* 28: 9239-9248.
- [6] Bechara A. (2005): "Decision making, impulse control and loss of willpower to resist drugs: a neurocognitive perspective," *Nature Neuroscience* 8: 1458-1463.
- [7] 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.
- [8] 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.
- [9] Bernheim B.D. and A. Rangel (2004): "Addiction and cue-triggered decision processes," *American Economic Review* 94, 1558-1590.
- [10] Breedlove S.M., M. Rosenzweig, and N. Watson (2007): *Biological Psychology*. Sunderland, Sinauer Associates.

- [11] Brocas I. and J. Carrillo (2008): "The brain as a hierarchical organization," *American Economic Review*, 98, 1312-1346.
- [12] Bush, G., P. Luu, and M. Posner. (2000). "Cognitive and emotional influences in anterior cingulate cortex," *Trends in Cognitive Sciences* 4: 215-222.
- [13] Camerer C., G. Loewenstein, and D. Prelec (2005): "Neuroeconomics: How neuroscience can inform economics", *Journal of Economic Literature*, XLIII, 9-64.
- [14] Camerer, C. (2009): "Behavioral economics," World Congress of the Econometric Society, forthcoming.
- [15] Cangelosi A., D. Parisi, and S. Nolfi (1994): "Cell division and migration in a "genotype" for neural networks," *Network* 5, 497-515.
- [16] Caplin A. and A. Schotter (2008): *The Foundations of Positive and Normative Economics*, New York, Oxford University Press.
- [17] Chen, Z. et al. (2008). "Revealing Modular Architecture of Human Brain Structural Networks by Using Cortical Thickness from MRI," *Cerebral Cortex* 8(10):2374-2381.
- [18] Damasio A, Van Hoesen GW. (1983). "Emotional disturbances associated with focal lesions of the limbic frontal lobe." In Heilman KM, Satz P, editors. *Neuropsychology of human emotion*. New York: Guilford. p 85–110.
- [19] Dasgupta P. and E. Maskin (2005): "Uncertainty and hyperbolic discounting," *American Economic Review* 95, 1290-1299.
- [20] Frederick S., G. Loewenstein, E. O'Donoghue (2002): "Time discounting and time preference: A critical review," *Journal of Economic Literature*, XL, 351-401.
- [21] Fudenberg D. and D. Levine (2006): "A dual self model of impulse control," *American Economic Review*, 96, 1449-1476.

- [22] Glimcher P. (2003): *Decisions, Uncertainty, and the Brain: The Science of Neuroeconomics*, Cambridge, Ma, MIT Press.
- [23] 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.
- [24] He, Y., Z. Chen, and A. Evans. (2007). "Small-World Anatomical Networks in the Human Brain Revealed by Cortical Thickness from MRI," *Cerebral Cortex* 17(10):2407-2419.
- [25] Hennion, H. (1997). "Limit theorems for products of positive random matrices," *Annals of Probability* 25: 1545-1587.
- [26] Hof PR, Nimchinsky EA, Perl DP, Erwin JM. (2001). "An unusual population of pyramidal neurons in the anterior cingulate cortex of hominids contains the calcium-binding protein calretinin." *Neurosci. Lett.* 307:139-42.
- [27] Hornik K., M. Stinchcombe, and H. White (1989): "Multilayer Feedforward Networks are Universal Approximators," *Neural Networks*, 2, 359-366.
- [28] 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.
- [29] 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.
- [30] Laibson, D. (1997): "Golden eggs and hyperbolic discounting," *Quarterly Journal of Economics*, 62, 443-77.
- [31] Latora V., and M. Marchiori. (2003). "Economic small-world behavior in weighted networks," *European Physical Journal B* 32: 249-263.

- [32] Loewenstein G. and E. O'Donoghue (2007): "Animal spirits: affective and deliberative processes in economic behavior," mimeo, Cornell University.
- [33] Loewenstein G. and D. Prelec (1992): "Anomalies in intertemporal choice: Evidence and interpretation," *Quarterly Journal of Economics*, 107(2), 573-97.
- [34] Luria S.E. and M. Delbrueck (1943): "Mutations of bacteria from virus sensitivity to virus resistance," *Genetics* 28, 491-511.
- [35] 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.
- [36] Meyer-Lindenberg, A. (2009). "Neural connectivity as an intermediate phenotype: Brain networks under genetic control," *Human Brain Mapping* 30: 1938-1946.
- [37] Monks, S.A., et al. (2004). "Genetic inheritance of gene expression in human cell lines," *American Journal of Human Genetics* 75: 1094-1105.
- [38] Nadler, J., et al. (2006). "Large-scale gene expression differences across brain regions and inbred strains correlate with a behavioral phenotype," *Genetics* 174(3):1229-1236.
- [39] Netzer, N. (2009): "Evolution of time preferences and attitudes towards risk," *American Economic Review*, forthcoming.
- [40] Rabin, M. (1998): "Psychology and economics," *Journal of Economic Literature*, 36(1), 11-46.
- [41] 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.
- [42] Ramachandran VS, and Hubbard, EM. (2001). "Psychophysical investigations into the neural basis of synesthesia". *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 268: 979-983.

- [43] 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.
- [44] Rustichini, A. (2008). "Dual or unitary system? Two alternative models of decision making," *Cognitive, Affective, and Behavioral Neuroscience* 8: 355-362.
- [45] Samuelson, L., and J. Swinkels. (2006). "Information, evolution, and utility," *Theoretical Economics* 1: 119-142.
- [46] Seneta E. (1981): *Non-Negative Matrices and Markov Chains*, New York, Springer Verlag.
- [47] 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.
- [48] Simon H.A. (1996): *The Sciences of the Artificial*. Cambridge, MA, MIT Press.
- [49] Thaler R. and H. Shefrin (1981): "An economic theory of self-control," *Journal of Political Economy*, 89, 392-406.
- [50] Varga R. (2000): *Matrix Iterative Analysis*, New York, Springer Verlag.

Appendix

The matrices A^{MH} and A^U when $n = W = 2$:

The dynamics of the population of $K = MH, U$ processes are governed by the linear system

$$\mathbf{N}_{t+1} = A^K \mathbf{N}_t.$$

Assuming each module is equally likely to mutate, in this case, A^{MH} is the 8×8 nonnegative projection matrix given by

$$A^{MH} = \begin{bmatrix} (1-q)G_1(0) & (1-q)G_2(0) & rG_1(0) & rG_2(0) \\ (1-q)P_1(0) & 0 & rP_1(0) & 0 \\ qG_1(0) & qG_2(0) & (1-q-r)G_1(0) & (1-q-r)G_2(0) \\ qP_1(0) & 0 & (1-q-r)P_1(0) & 0 \\ 0 & 0 & (\frac{2q}{3})G_1(1) & (\frac{2q}{3})G_2(1) \\ 0 & 0 & (\frac{2q}{3})P_1(1) & 0 \\ 0 & 0 & (\frac{q}{3})G_1(0) & (\frac{q}{3})G_2(0) \\ 0 & 0 & (\frac{q}{3})P_1(0) & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ rG_1(0) & rG_2(0) & rG_1(0) & rG_2(0) \\ rP_1(0) & 0 & rP_1(0) & 0 \\ (1-q-r)G_1(1) & (1-q-r)G_2(1) & 0 & 0 \\ (1-q-r)P_1(1) & 0 & 0 & 0 \\ 0 & 0 & (1-q-r)G_1(0) & (1-q-r)G_1(0) \\ 0 & 0 & (1-q-r)P_1(0) & 0 \end{bmatrix}.$$

where the apex $K = MH$ is dropped for notational simplicity and the value of μ in $G_i(\mu)$

and $P_i(\mu)$ reflects the number of influential mutations.

Consider for example the elements of row 5, which represent the inflow during period t into the type $N_1^{1,1,0}$. A fraction q of the processes of type $N_i^{1,0,0}$ mutate during period t . For 2/3 of these, the second mutation happens in a different module. As a result, these processes select a suboptimal action and leave surviving offsprings of $G_i(1)$. Next consider processes of type $N_i^{1,1,0}$. During period t , a fraction q develop a third harmful mutation and die without leaving any offspring. In addition, for a fraction r , one of the harmful mutations is reversed and they become type $N_i^{1,0,0}$. The remainder, $(1 - q - r)N_i^{1,1,0}$, do not undergo any change and produce a number of descendents given by $G_i(1)$. Finally, since the probability of more than one mutation per period is essentially zero, there is no direct inflow from any other type of process.

The projection matrix, A^U , assuming the values under the optimal action are still denoted by $G_i(0)$ and $P_i(0)$, is given by

$$A^U = \begin{bmatrix} (1-q)G_1(0) & (1-q)G_2(0) & rG_1(0) & rG_2(0) \\ (1-q)P_1(0) & 0 & rP_1(0) & 0 \\ qG_1(1) & qG_2(1) & (1-q-r)G_1(1) & (1-q-r)G_2(1) \\ qP_1(1) & 0 & (1-q-r)P_1(1) & 0 \\ 0 & 0 & (\frac{2q}{3})G_1(2) & (\frac{2q}{3})G_2(2) \\ 0 & 0 & (\frac{2q}{3})P_1(2) & 0 \\ 0 & 0 & (\frac{q}{3})G_1(2) & (\frac{q}{3})G_2(2) \\ 0 & 0 & (\frac{q}{3})P_1(2) & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ rG_1(1) & rG_2(1) & rG_1(1) & rG_2(1) \\ rP_1(1) & 0 & rP_1(1) & 0 \\ (1-q-r)G_1(2) & (1-q-r)G_2(2) & 0 & 0 \\ (1-q-r)P_1(2) & 0 & 0 & 0 \\ 0 & 0 & (1-q-r)G_1(2) & (1-q-r)G_1(2) \\ 0 & 0 & (1-q-r)P_1(2) & 0 \end{bmatrix}.$$

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 have $m = [j, k, l]$ with $M = j + k + l \leq W$ 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, all mutations affect the executive process. Hence we have $\mu^U(m) = M$. On the other hand, in the MH architecture, we know that $\mu^{MH}(m) \leq \max\{j, k, l\} \leq M$, where the first inequality is strict if $\max\{j, k, l\} = M$. Specifically, the first equality will hold in the situation where $\max\{j, k, l\} < M$, the number of mutations occurring in an advisory module is $\max\{j, k, l\}$, and the number of mutations in the motivational module brings its reference policy closer to the action associated with $\max\{j, k, l\}$ mutations than to the reference action of the other advisory module. We thus have $G_i(\mu^{MH}(m)) > G_i(\mu^U(m))$ and $P_i(\mu^{MH}(m)) > P_i(\mu^U(m))$. But then, 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

$$\phi^* = \arg \max_{x \geq d(s,s')} \{x\}$$

$$\begin{aligned} \text{s.t.} \quad \gamma_\mu^G(1/x)G_i^{MH}(\mu^{MH}(m), s') &\geq G_i^U(\mu^U(m), s') \quad \forall(i, m), \\ \gamma_\mu^P(1/x)P_i^{MH}(\mu^{MH}(m), s') &\geq P_i^U(\mu^U(m), s') \quad \forall(i, m). \end{aligned}$$

Clearly, given that $\lim_{x \downarrow d(s,s')} \gamma_\mu^J(1/x) = 1$ and $\gamma_\mu^J(0) = 0$, a solution to the above program will always exist. Moreover, given that $G_i^{MH}(\mu, s') > G_i^U(\mu, s')$ and $P_i^{MH}(\mu, s') > P_i^U(\mu, s')$,

we know that $\phi^* > d(s, s')$. But then for a MH architecture of coarseness $\phi^{MH} \in [d(s, s'), \phi^*)$,

$$G_i^{MH}(\mu^{MH}(m), s' | s) > \gamma_\mu^G(1/\phi^*)G_i^{MH}(\mu^{MH}(m), s') \geq G_i^U(\mu^U(m), s')$$

and

$$P_i^{MH}(\mu^{MH}(m), s' | s) > \gamma_\mu^P(1/\phi^*)P_i^{MH}(\mu^{MH}(m), s') \geq P_i^U(\mu^U(m), s').$$

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

Now consider a sequence $S_{(n)}(\omega)$ of $n \geq 1$ realizations of the random process taking values in $\{s, s'\}$. Following the reasoning of Proposition 1, we know that $A^{MH}(s)$ is an increasing transformation of $A^U(s)$ and hence $\rho^{MH}(s) > \rho^U(s)$. 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 DA switches to the newly optimal profile. On the other hand, the modular DA continues to operate according to the old profile. Even though the modular DA behaves suboptimally, for ϕ small, its superior energy efficiency and mutation resistance allow it to outperform the unitary DA even under the new environment.

Now let $A_{(n)}^K(\omega)$ denote the random product of projection matrices for type K associated with the environmental sequence $S_{(n)}(\omega)$. We are particularly interested in the associated spectral radii $\rho_{(n)}^K(\omega)$ since by a theorem of Hennion²⁹,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \ln \rho_{(n)}^K = \Lambda^K \quad \text{almost surely,}$$

where Λ^K is the dominant Lyapunov exponent governing the asymptotic long-run growth of the system of type K . We will now show that $\rho_{(n)}^{MH}(\omega) > \rho_{(n)}^U(\omega)$ for any n , so the MH will

²⁹Theorem 2 in Hennion (1997).

asymptotically outperform the U architecture starting from any non-zero initial condition.

The proof is done by induction. We already know that $\rho^{MH}(s') > \rho^U(s')$ and $\rho^{MH}(s) > \rho^U(s)$. But we also 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 $A_{(2)}^{MH}(\omega)$ must be increasing transformations of $A_{(2)}^U(\omega)$. 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, by Thm 2.7. in Varga (2000), $\rho_{(2)}^{MH}(s, s') > \rho_{(2)}^U(s, s')$.

Proceeding in the same fashion, we know that $A_{(n-1)}^{MH}(\omega) \geq A_{(n-1)}^U(\omega)$ and with the inequality strict for at least one element ij . But $A_{(n)}^K(\omega) = A^K(s_n)A_{(n-1)}^K(\omega)$ and since $A^{MH}(s_n)$ is an increasing transformation of $A^U(s_n)$ for s_n equal to s or s' , $A_{(n)}^{MH}(\omega)$ must be an increasing transformation of $A_{(n)}^U(\omega)$. Thus $\rho_{(n)}^{MH}(\omega) > \rho_{(n)}^U(\omega)$ for any $n \geq 1$ and any ω . But then using the above law of large numbers, we obtain $\Lambda^{MH} > \Lambda^U$.

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

Sexual Reproduction with Random Mating:

When mating is random, the projection matrix for each sex will reflect the flows across types for that particular sex. Moreover, the entries in each projection matrix will depend on the frequency distribution of types from the other sex. When agents from different types mate, only the type of one of the parents will be inherited. Hence the daughter of a particular female agent will be of the same type m as her mother only with probability $h^{m,m}$. Nevertheless, for any steady state MH population distribution, we can provide a sufficient condition under which all of our results hold locally around that steady state for random mating and any steady state U distribution.

Proposition 4: Let $(f_j^0)_{j=1}^q$ be a steady state vector of frequencies of MH males or reproductive age j with 0 total deleterious mutations, so that $\sum_{j,m} f_j^m = 1$. If $(\sum_j f_j^0)G_i^{MH}(\mu^{MH}(m), s) > G_i^U(\mu^U(m), s)$ for any i, m , and s , Propositions 1, 2, and 3 hold locally around this steady state.

Proof: To see why the above condition is sufficient, note that the survival probabilities do not depend on the mating process so nothing changes as far as the rows of the projection matrices determined by P_i are concerned. On the other hand, the expected number of offsprings of each type does depend on the mating process. The best case scenario for the U population is when in a steady state it is dominated by males with 0 deleterious mutations. In that case any off-diagonal entry of the female projection matrix can have a positive entry. In particular, $h^{m,0}$ of the daughters produced by a mother of a given type m will have the type $(0, 0, 0)$ of their father. On the other hand, for the MH population that will be true only for $h^{m,0} \cdot (\sum_j f_j^0) < h^{m,0}$ of the daughters. Hence as long as $(\sum_j f_j^0)G_i^{MH}(\mu^{MH}(m), s) > G_i^U(\mu^U(m), s)$, any entry in the first row of A^{MH} will be greater than the corresponding entry in A^U . Moreover, since the U population is dominated by males of type $(0, 0, 0)$ while the MH population is not, the entries for any other row in A^{MH} will be greater than or equal to the corresponding entries in A^U since there will be no flows, except for mutations, across other types in the U population. But then A^{MH} is a nondecreasing transformation of A^U . We thus must have $\rho^{MH} > \rho^U$ and all of our previous conclusions carry through. ■

Numerical Example:

Under the specified assumptions, we have to distinguish among 13 types of DAs: (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 motivational unit, (7) three mutations in two modules with one mutation in the motivational unit, (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

motivational unit, (12) four mutations occurring three by one with one mutation in the motivational unit, 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 DAs into a DA 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 DAs executing action a^s with only 0.001 of the DAs behaving according to a^I . Even though there is conflict among modules for over 60 percent of the population, that conflict does not translate into behavioral heterogeneity.

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.