Chapter 3

Cognitive Aging and Strategy Application

Cognitive aging has been described at different levels. At the behavioral level, researchers have identified age differences in learning rate and asymptotic performance (Baltes & Kliegl, 1992), interference susceptibility (Zacks & Hasher, 1997), complexity cost (Lair, Moon, & Kausler, 1969), and intra- and interindividual variability (Lindenberger & Baltes, 1997). At the neurological level, neuroanatomical (Bertoni-Freddari, Fattoretti, Casoli, Meier-Runge, & Ulrich, 1990) as well as neurochemical changes due to aging have been reported (e.g., Arnsten, 1998). Finally, at the information processing level, it has been proposed that there are age related reductions in general processing resources (Salthouse, 1996), working memory (Grady & Craik, 2000), and attentional mechanisms (McDowd & Shaw, 2000). But how do these different levels of description come together as a whole? How do they relate to decision behavior and, in particular, the application of decision strategies?

As reviewed in Chapter 1, some studies suggest that older adults look up less information and take longer to make a choice than younger adults (e.g., Johnson, 1990). Additionally, as reported in Chapter 2, which investigated the impact of cognitive aging on strategy selection, older adults tend to use simpler strategies compared to younger adults. Hence, some preliminary evidence exists supporting the idea that age-related cognitive decline impacts strategy use. However, past work did not focus on age differences in strategy application, namely, age differences in the tendency to make application errors when using decision strategies. Thus, it is not known how efficiency in the application of decision strategies changes as a function of age.

A principled way of investigating this issue is to consider already detailed models of decision making, associate them with a theory of aging and, subsequently, design empirical studies to test the models' predictions against data. This is the strategy taken here. Accordingly, the present chapter integrates the neurological, and information processing levels of description by presenting a neurocomputational account of aging effects in the efficiency of strategy use. Additionally, at the behavioral level, a test of the predictions derived from the approach is reported.

Using a connectionist framework to model the effect of aging on decision making capitalizes on the age-related decline in neural functioning being naturally modeled with subsymbolic systems (see Polk, Simen, Lewis, & Freedman, 2002, for the same argument; Li & Lindenberger, 1999). Hence, the approach adopted here connects the connectionist and the adaptive toolbox frameworks to understand age differences in strategy application.

Combining the Connectionist and Adaptive Toolbox Frameworks

A considerable number of decision making approaches, including the adaptive toolbox one, have defined decision strategies as production rules (if preconditions 1 and 2 are met, take action A; Svenson, 1979; Gigerenzer et al., 1999; Payne et al., 1993). Thus, they have endorsed the view that the mind manipulates symbols, executing operations over variables to produce behavior (Newell, 1980). However, the emergence of connectionism has lead some to question the extent to which the mind can be seen as a symbol manipulator (cf., Marcus, 2001; Mills, 1992).

The Eliminative View of Connectionism

McCulloch and Pitts's (1943) seminal work showed that a network of computational units resembling neurons (as they were believed to function at the time) — simple threshold units with binary states, performing summation of excitatory or inhibitory inputs — could compute various logical functions (e.g., AND, OR). Since then, researchers have investigated the computational power of artificial networks in the hope of understanding how the brain may perform various cognitive processes. Connectionist models are loosely based on the principles of neural information processing, and subscribe to the idea that cognitive processes involve the basic computations being performed in parallel by a large number of densely interconnected neurons. Perhaps the most significant publication supporting this view has been *Parallel Distributed Processing* by Rumelhart, McClelland, and the PDP Research Group (1986) which applied the connectionist framework to various cognitive domains, from motor planning to language acquisition (see Elman, Bates, Johnson, Karmiloff-Smith, Parisi, & Plunkett, 1996, and McLeod, Plunkett, & Rolls, 1998, for introductions to connectionist modeling).

Some connectionist enthusiasts have questioned the ontological status of production rules and portrayed neural network models as markedly opposed to the symbol-manipulation perspective (e.g., Churchland, 1995). According to this *eliminative* view of connectionism, the basic components of cognition have to do with distributed processing of individual units whose global behavior cannot be said to represent a rule.

The field of decision making research has been no exception in witnessing the proliferation of connectionist accounts of cognitive processes. Several connectionist models of decision making have been proposed (e.g., Chater, Oaksford, Nakisa, & Redington, 2003; Roe, Busemeyer, & Townsend, 2001; Usher & McClelland, 2004; see Busemeyer & Johnson, in press, for a review) and shown to handle complex processes of information integration. This contrasts with the adaptive toolbox and adaptive decision maker approaches' attempts to define decision mechanisms as production rules (Payne et al., 1993; Schooler & Hertwig, 2005). As a consequence, it is understandable that attempts have been made to pit connectionist models against heuristics (Chater et al., 2003). Nevertheless, as I argue below, this opposition is not warranted.

The Implementationist View of Connectionism

An alternative implementationist perspective to connectionism is well established in cognitive science (Marcus, 2001; Pinker & Prince, 1988; Smolensky, 1988). According to this position, the functioning of connectionist networks can be interpreted as implementing aspects of symbol manipulation. Consequently, one can implement production rule type of processes using connectionist architectures. For example, one can model decision strategies, such as TTB, using neural networks. The modeling efforts reported below can be seen as an existence proof of the possibility of understanding the workings of prototypical decision strategies, such as compensatory and noncompensatory rules, as the combined functioning of units in an artificial neural network. Thus, it is argued that these networks implement aspects of decision strategies usually defined as production rules, thus, supporting an implementationist view of connectionism.

Combining the adaptive toolbox and the connectionism frameworks provides clear benefits. First, the models proposed by the adaptive toolbox approach can inform a connectionist agenda in the domain of inference by supplying plausible algorithms to be modeled. Second, the synergy can help connect different levels of explanation, from the behavioral, to the information processing, to the neurological levels. Finally, as proposed here, it can help tackle issues such as ontogenetic change in efficiency of strategy use due to age-related cognitive decline.

In the next sections, I introduce a formal account of aging and present how connectionist implementations of decision strategies can be associated with an aging theory to obtain predictions regarding age differences in strategy application.

Formal Modeling of Aging

Theories of aging have worked at cross-level unification (Li, 2002) by relating three major levels of description: the behavioral, the information processing, and the neurological. The neurological level has been rich in providing evidence of how increased age is related to neuroanatomical and neuromodulatory changes. For example, age-related changes have been reported in synapse morphology and decrements in number of synapses, although not in number of neurons (Bertoni-Freddari et al., 1990; Geinisman, DeToledo-Morrell, Morrell, & Heller, 1995). In addition, some have proposed that milder cognitive deficits during normal aging may be caused by neurochemical changes in relatively intact neural circuits (cf. Morrison & Hof, 1997). Li and collaborators (Li, 2002; Li, Lindenberger, & Frensch, 2000; Li, Lindenberger, & Sikström, 2001) have argued that age-related decline in cathecolaminergic function, in particular dopaminergic modulation, is responsible for major age differences found between younger and healthy older adults at the behavioral level. Support for this view comes from studies showing that aging-related attenuation of particular dopamine receptor binding mechanisms is statistically associated with age differences in processing speed and episodic memory (Bäckman, Ginovart, Dixon, Robins-Wahlin, Wahlin, Halldin, & Farde, 2000).

At the information processing level, there have been attempts to link neurological facts to computational models of cognitive processes, such as the role of dopamine in cognitive control and executive functions (Cohen, Braver, & Brown, 2002; Montague, Hyman, & Cohen, 2004; see Fellous & Linster, 1998, for an overview of computational models of neuromodulation). One approach particularly concerned with aging was adopted by Li and colleagues (e.g., Li et al., 2000, 2001) who conceptualized deficits in neuromodulatory efficiency due to aging as noisy information processing and the existence of less distinct neural representations. Specifically, Li et al. (2001) proposed that deficits in catecholaminergic activation in the prefrontal cortex (Arnsten, 1998) can be modeled by adjusting the gain (*G*) parameter of the sigmoidal activation function of neural networks (see Cohen & Servan-Schreiber, 1992, Servan-Schreiber, Printz, & Cohen, 1990, for a similar approach). According to this approach, the activation function of neural units takes the form:

$$(1) \qquad Activation_{it} = \frac{1}{1+e^{-(G_{it} \times Input_{it} + bias)}}$$

where *i* refers to a unit, *t* indicates the processing step, *Input* refers to the input supplied to the unit ($Input = \Sigma\ W_{ij} I_j$), and *bias* refers to a negative bias unrelated to aging (usually, *bias* = -1;

Servan-Schreiber et al., 1990). Finally, *G* is a random number sampled from a uniform distribution ($G_i, \in [G_{min}, G_{max}]$, $G_{min} > 0$; see Figure 3.1 for the values of *G* used by Li & Lindeberger, 1999).

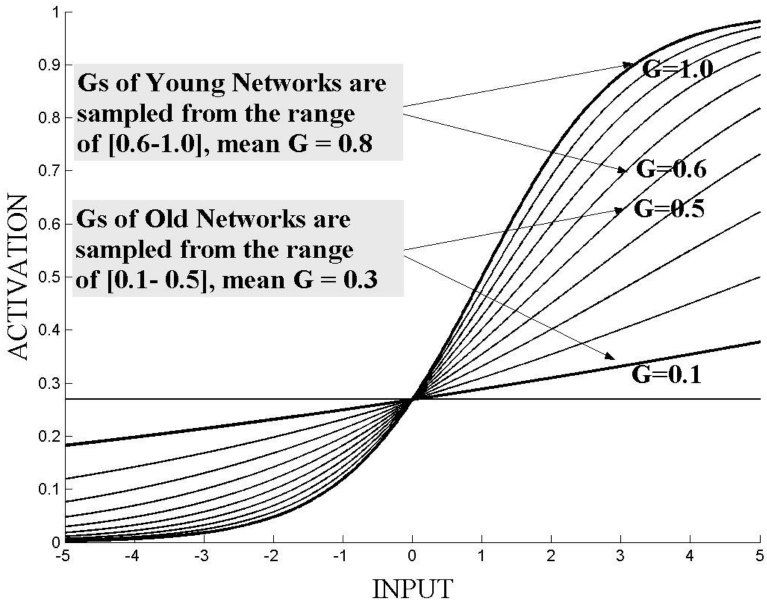

Figure 3.1: The S-shaped logistic activation function at different values of *G* (adapted from Li et al., 2001)

Reducing mean *G* flattens the logistic activation function such that a unit's average activation responsivity is reduced (see Figure 3.1). As a consequence, such as unit becomes less discriminative in responding to differences in inputs. In addition, reducing mean *G* increases the variability in the unit's activation.

Providing an encouraging example in the effort of cross-level unification, Li and colleagues have shown that manipulation of the *G* parameter allowed modeling major results stemming from aging research, specifically, age differences in mean performance, age by task difficulty effects (i.e., complexity cost), and increases in performance variability with increased age (Li & Lindenberger, 1999; Li & Sikström, 2002; Li et al., 2000, 2001; Li, Naveh-Benjamin, & Lindenberger, 2005).

The next section reports efforts to implement decision strategies using connectionist networks and an attempt to use the gain manipulation method to model age-related changes in efficiency of strategy use.

A Neurocomputational Account of Age-related Changes in Strategy Application

I used simple recurrent networks (SRN; Elman, 1990) to implement two decision strategies: TTB, and an evidence accumulation strategy (EAS), which can be thought of as a version of Franklin's Rule that implies sequential information search. Simple recurrent networks have been used extensively in the modeling of various aspects of cognition involving sequential processing of information, such as language and spatial orientation (see Elman, 2005). SRN are well suited to deal with sequential processing of information because internal states are fed back at every step, which supplies such networks with a memory, and allows them to process information sequentially over time. The feature of sequential processing is also a common requirement in decision making tasks like pair-comparison ones which are the focus of the reported modeling efforts. Hence, SRN are well-matched to the modeling efforts reported below.

Decision strategies such as TTB and EAS are composed of different building blocks, namely, a search rule, which determines the order in which cues are looked up, a stopping rule, which determines when search for information ends, and a decision rule, which determines which object is chosen. In the work reported below I focused only on the decision rules of TTB and EAS because I was particularly interested in how differences between input-output mappings' of different strategies determine age differences in performance.

Age-related decline in neuromodulation and its effect on strategy efficiency was modeled by varying *G* stochastically (cf. Li et al., 2001) in the activation function of SRN' artificial neural units. In general, sampling *G* from a distribution with a lower mean at the hidden unit level should produce less pronounced activations (cf. Li et al., 2000), which should in turn be reflected in the output units' activations. Additionally, the gain manipulation should influence variability in units' activations because it makes use of the property of squashing functions (see Figure 3.1), by which changing their slope influences activation variability. In general, it was expected that changing the networks' G parameter would increase the probability of choosing an alternative not compatible with that prescribed by the strategy being implemented.

Another issue being investigated was whether information intensive strategies were more susceptible to *G* manipulations, revealing a complexity cost effect. One reason to predict such an effect is that the code underlying the use of TTB is sparser than that of EAS: The input-output mapping required to implement TTB is simpler (all inputs lead either to [0,0], [1,0], or [0,1] outputs) than that for EAS (inputs can lead to several different outputs, such as

[0,0], [.125,0], [0,.125], [.125,.125], [.25,.125], and so on). Because *G* manipulation should produce networks with restricted computational power (Li et al., 2000) one would predict these to have more difficulties with a more complex input-output mapping.

Method

The general method underlying the experimental procedure was to train a number of SRN with the goal of finding a set of networks that would implement the decision strategies correctly and age only this successfully trained subset. The SRNs used to implement the two decision strategies had 2 input, 20 hidden, 20 context, and 2 output units (see Figure 3.2). The decision problem that the networks faced was a paired comparison between objects on the basis of 8 cue values. One-hundred networks with different initial random weights were trained using both TTB-congruent and EAS-congruent target activations. Thus, each network was trained once to implement TTB and once to implement EAS.

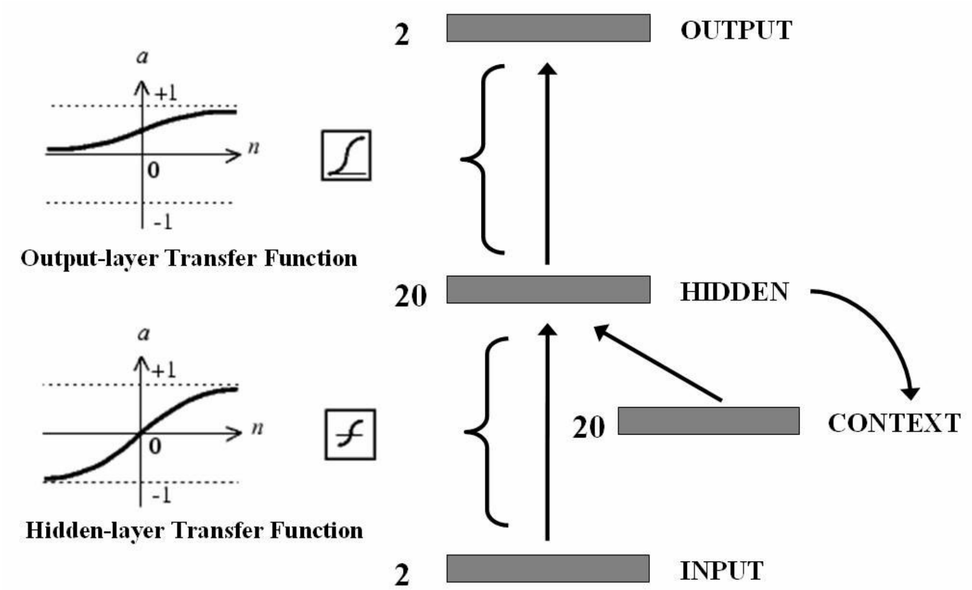

Figure 3.2: Architecture and transfer functions of the simple recurrent networks.

The number of possible input patterns for training, given all combinations of 8 binary cues for two alternatives, is 65536. However, because it was impractical to train networks using the complete set of input patterns, a random sample of 500 input patterns was selected. In addition, to insure the networks correctly implemented TTB and EAS, an additional set of 100 trials was randomly selected to be used after training in a generalization test. The total number of input patterns used to train the networks was, therefore, 500. Each training epoch consisted of supplying as input the same sequence of 500 sets of 16 vectors with 1 cue value each, corresponding to 500 decisions between 2 different objects based on 8 cues. Note that the hidden layer received at each step a vector with 1 cue value as well as vector with its own activation pattern from the previous time-step, whose values were reset at the beginning of each decision.

At the end of each epoch the network weights were updated using a gradient descent backpropagation algorithm (Rumelhart, Hinton, & Williams, 1986) with adaptive learning rate and momentum. The networks were trained for 500 epochs. Target activations were constrained to range between 0 and 1. The two versions, TTB and EAS, differed in terms of the inputs and respective targets provided. The crucial difference between the inputs reflects the fact that TTB and EAS algorithms look up information differently. While TTB searches for information in a cue-wise fashion, EAS searches information in an alternative-wise manner. The targets also reflected this input scheme. In the TTB implementation the first object having a positive value on a discriminating cue had an activation of 1, while for the EAS each positive cue value led to a proportion of the maximum possible activation[1]. For example, given 8 cues with equal validities, the contribution of each cue to an object's activation is .125. Thus, EAS targets reflected the principle of information accumulation (see Tables 3.1 and 3.2).

Table 3.1: Example of Mapping between Input and Activation of Alternatives for TTB

	TTB Inputs		TTB Targets	
	A	B	A	B
Cue 1	1	-	0	0
	-	1	0	0
Cue 2	0	-	0	0
	-	0	0	0
Cue 3	0	-	0	0
	-	0	0	0
Cue 4	1	-	0	0
	-	0	1	0
Cue 5	1	-	1	0
	-	1	1	0
Cue 6	1	-	1	0
	-	1	1	0
Cue 7	1	-	1	0
	-	0	1	0
Cue 8	1	-	1	0
	-	0	1	0

[1] For EAS, the object target activation corresponding to each input cue value was defined by: $value_i \times v_i/\Sigma v$, where $value_i$ represents the object's value on cue i, v_i the validity of that cue, and Σv the sum of all cue validities.

Table 3.2: Example of Mapping between Input and Activation of Alternatives for EAS

	EAS Inputs		EAS Targets	
	A	B	A	B
Cue 1	1	-	.125	0
Cue 2	0	-	.125	0
Cue 3	0	-	.125	0
Cue 4	1	-	.250	0
Cue 5	1	-	.375	0
Cue 6	1	-	.500	0
Cue 7	1	-	.625	0
Cue 8	1	-	.750	0
Cue 1	-	1	.750	.125
Cue 2	-	0	.750	.125
Cue 3	-	0	.750	.125
Cue 4	-	0	.750	.125
Cue 5	-	1	.750	.250
Cue 6	-	1	.750	.375
Cue 7	-	0	.750	.375
Cue 8	-	0	.750	.375

Age-related decline in neuromodulation and its effect on strategy efficiency was modeled by varying *G* stochastically (cf. Li et al., 2001) in the activation function of the hidden layer of the trained SRNs. All networks were tested using 6000 trials (i.e., each trial consisting of a set of 16 vectors corresponding to the 8 cue values for each option) which corresponded to 10 runs through the set of 600 trials used to train and test the generalization performance of the SRN. Decisions were made using a difference choice rule: the difference between activations in the output units was computed and the object with the highest value was chosen.

Results

The results section is structured in the following way; first, the results of the training procedure are described. Second, the outcomes of the gain manipulation are presented generally. Afterwards, the latter results are quantified and linked to the major and most robust findings in aging research, namely, age differences in mean performance, complexity cost, and performance variability.

Training Results

After 500 training epochs the mean MSE was considerably small (MSE_{TTB} = .0007; MSE_{EAS} = .0051). However, performance was not perfect: only 20 networks showed perfect performance in the sense that the outputs closely matched the target activations and the difference between activations was in the prescribed direction both when trained with TTB and EAS targets. This subset of networks was further analyzed concerning the generalization set and it was found that only 10 generalized perfectly to the new 100 trials (MSE < .003). No systematic variation of the architecture (e.g., number of hidden units) or learning parameters (e.g., learning algorithm) was attempted because the focus of the simulation was to create a set of networks that performed TTB and EAS input-output mappings successfully, rather than to provide a thorough explanation of how these parameters influenced the training procedure.

G Manipulation

The effects in performance of both the TTB and EAS implementations of the systematic variation of *G* can be observed in Figure 3.3. This figure shows that both the implementations of TTB and EAS show decline in performance with decreasing values of *G*. However, there are differences between the strategies' performance; while TTB and EAS seem to differ little at extreme levels of G, differences between the two are accentuated when middle values of *G* are considered. Furthermore, as suggested by differences in the size of the error bars (SD) in Figure 3.3, TTB and EAS show different variability levels in performance as a function of *G*. In sum, TTB is a *more robust strategy* compared to EAS when signal-to-noise ratio of neural units decreases.

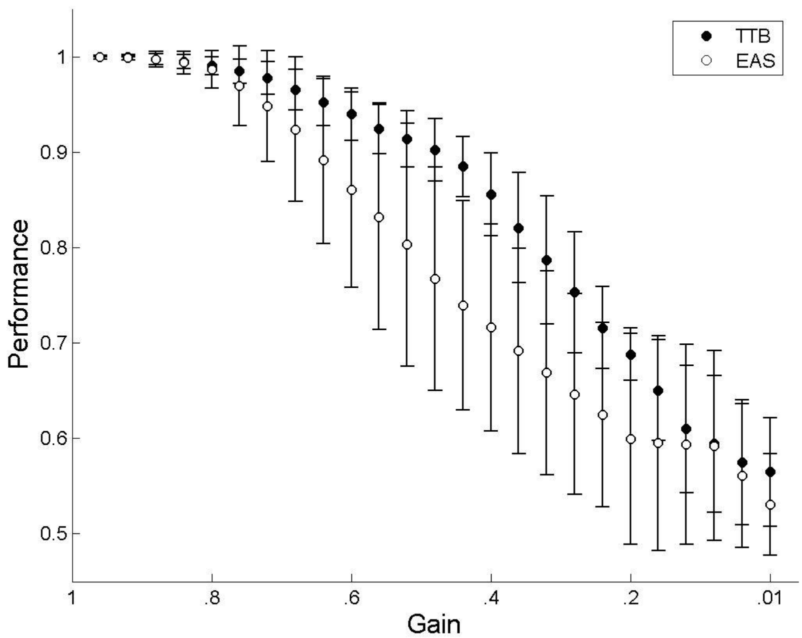

Figure 3.3: Accuracy in choosing the right object when implementing TTB and EAS given different values of *G*. Error bars correspond to intra-network variability (SD) on different runs through the testing set.

The following paragraphs summarize these results which can be viewed as predictions regarding the effects of age-related cognitive decline on the application of decision strategies. The predictions made are based on the *G* values used by Li and Lindenberger (1999).

Mean performance. Young and older adults usually differ in terms of mean performance even after considerable periods of training (Baltes & Kliegl, 1992). Also, research on strategy use in mental arithmetic and memory domains suggests that older adults use strategies more poorly than younger adults (e.g., Dunlosky & Hertzog, 1998; Lemaire et al., 2004). The results from the simulations mirror these effects: networks with low values of *G* (.1 to .5) show worse performance than those with high values of *G* (.6 to 1). The difference between the performance of young and old networks represents a large effect (97% vs. 71%; *d* = 1.09). In sum, as expected, young networks perform on average better than the old, leading to the prediction that older adults will more often make mistakes when using a decision strategy than their younger counterparts.

Complexity cost. Another robust empirical finding in the aging literature is an age by task complexity effect, that is, an increase in the difference between the performance of young and older adults with increasing processing demands or task difficulty (i.e., complexity cost; e.g., Lair et al., 1969). Several researchers have claimed that lexicographic rules are less computationally demanding than other more information-intensive decision strategies (Dosher & Russo, 1983; Gigerenzer et al., 1999; Payne et al., 1993). In accordance with these claims, an age by complexity effect was observed in the simulation results: The difference between the performance of TTB and EAS for young networks was considerably smaller than the difference between the two strategies when the performance of old networks is considered (2% vs. 9%). One interpretation of this result is that older networks deal better with sparser codes (cf. Li et al., 2000) and that TTB has an advantage in this respect.

Performance variability. Behavioral data points to an increase in intra-individual variability in performance with increased age (e.g., Fozard, Thomas, & Waugh, 1976). To evaluate whether the simulations also showed this effect, each network was run ten times on the testing set. The standard deviation in overall performance of each network was then used as a measure of network intra-variability. As expected, the variability of old networks (*SD* = .09) was greater than that of young networks (*SD* = .03). However, there may exist a bottom effect, whereby decreasing performance to chance level necessarily diminishes the variability in performance; this effect is visible in Figure 3.3 for extremely low values of *G*.

Discussion

A neurocomputational approach was presented which builds on previous accounts of decision making (Gigerenzer et al., 1999) and formal modeling of aging (Li et al., 2000, 2001) to predict effects of age-related cognitive decline on the efficiency of decision strategies. The underlying rationale was that age-related changes in neuromodulatory processes impact cognition at the level of its basic information processing components, such as working memory and processing speed. In turn, these age-related changes produce age differences at the behavioral level, including differences in aspects of higher-order cognition, namely, decision making. The remainder of the discussion addresses potential extensions to the approach, and, additionally, how the approach relates to integrative theories and models of cognition.

Extensions to the Approach

Two main extensions to the modeling approach should be attempted. First, one should consider information search and make predictions regarding age differences in this respect. As it stands, only TTB's decision rule, but not its stopping rule, was modeled (cf. Gigerenzer et al., 1999). However, one could add an output node to the SRNs to signal stopping search and train networks appropriately. Extending the approach in this manner would allow making predictions about age differences in information search behavior.

Second, the underlying dynamics of SRNs implementing the decision strategies should be investigated. Simple recurrent networks are discrete-time dynamical systems (Beer, 2000), which can thus be described by mathematical models (a set of equations) that represent the time dependence of units' activations in a geometrical space. Hence, the tools of dynamical systems theory can be used to investigate how, first, these networks arrive upon a solution to these simple inference problems (see Elman, 2005, and Rodriguez, Wiles, & Elman, 1999, for such an approach) and, secondly, the effects of the aging manipulations on the underlying dynamics.

Integration, Integration, Integration

An integrative theory is one that attempts to connect different levels of description. This approach is the result of such an attempt: Neurological and information processing levels were connected to provide a picture of the relation between age-related cognitive decline and decision making.

A different sort of integration can be done at the level of models. The adaptive toolbox approach is based on the assumption that people possess a repertoire of strategies. However, some concern has been expressed concerning the proliferation of strategies and the resulting lack of parsimony (Newell, 2004). Consequently, it has been argued that general, parameterized models should be tested in competition with simple heuristics (Lee & Cummins, 2004). The work put forward here is compatible with the perspective that a general model may account for individuals' behavior because the architectures used for the two different decision strategies were the same. However, the mapping between inputs and outputs, in the form of connections between units, differed considerably, making the networks conceptually distinct. As a consequence, this approach is also well-matched to the idea that people have at their disposal a set of tools adapted to specific environments. Nevertheless, there is some room for bringing closer the concerns of the adaptive toolbox and the general process perspectives. The potential intersection resides in the need to specify a theory of how weights of connections in networks get adjusted at different time scales, including evolutionary time, the lifecourse, or through learning during a short and novel task. One interesting avenue of inquiry is to try to understand how decision makers set off from a general model, such as SRN with random weights, to distinct mechanisms. It would be particularly important to question how one can model these changes over the different timescales.

Summary of Predictions

The simulation results provided a clear set of predictions to be tested. First, older adults should generally have more difficulties in strategy application, which should be reflected in the mean number of application errors when applying decision strategies (mean performance). Second, older adults should show a larger decrement in performance with more difficult strategies when compared to younger adults, leading to an age by strategy difficulty effect (complexity cost). Third, older adults should show larger intra-individual variability in performance (performance variability). These empirically testable predictions are examined in Study 4.

Study 4: Age Differences in Application of Decision Strategies

Study 4 investigates younger and older adults' application of decision strategies. In particular, the study tests the predictions derived from the neurocomputational approach adopted, which include (1) age differences in mean performance, (2) an age by strategy-difficulty effect, with the performance of older adults showing a larger decrement with more cognitively demanding strategies compared to younger adults, and (3) larger intra-individual variability in performance with increased age.

In addition to testing these main predictions, further support for the idea that success in strategy application is related to basic components of cognition was sought. This was investigated by, first, fitting the models' *G* parameter to mimic individual participants' behavior and, second, relating the fitted values to measures of the cognitive constructs of interest, such as working memory capacity and speed of processing. Finding a relation between fitted parameter values and individuals' performance in independent cognitive tasks would provide some support for the thesis that individual differences in cognitive capacity, in particular those due to aging, play a major role in the efficiency of strategy use.

In most decision making experiments participants are usually classified as users of a particular strategy by, first, evaluating the fit between their search and decision behavior to that prescribed by different strategies, and, subsequently, selecting the strategy with highest fit as representing the participant's overall behavior. One difficulty with such an approach is trying to determine whether participants' deviations from the search and choice prescribed by particular strategies reflect participants' use of a different, not previously considered strategy or, alternatively, an application error. Therefore, it is not possible to reliably estimate application errors on the basis of data from usual decision making experiments. One alternative is to teach participants how to use different decision strategies and later ask them to perform each one independently (see Bettman et al., 1990, for such a design in the decision making domain, and Siegler & Lemaire, 1997, for one in the arithmetic computation domain). Consequently, our study made use of trained participants to investigate potential age differences in the application of decision strategies.

The study involved teaching participants two decision strategies, TTB and EAS, in the context of a paired comparison task. After a training period, young and older participants were asked to use these strategies in a series of trials. I have argued above that age differences in strategy use will be evident in the accuracy of strategy application, therefore, Study 4 assessed whether participants correctly chose the options recommended by TTB and EAS. In addition, participants were asked to rate the cognitive demand of the two strategies on 7-point Likert scales. Past research has shown that decision makers' estimates of cognitive demand of strategies match closely the theoretical predictions of researchers (Bettman et al. 1990; Chu & Spires, 2003), thus, these variables were expected to corroborate the results found in terms of performance accuracy.

Finally, measures of cognitive capacity were used to assess working memory and processing speed constructs.

Method

Participants

A total of 22 adults participated in the study, 12 young adults (8 female, 4 male, *M* age = 25.58 years, *SD* = 3.68), and 10 older adults (2 female, 8 male, *M* age = 71.00, *SD* = 3.94). The study took about 1 1/2 hrs. Most young adults were students in various departments of the Free University of Berlin (80%), and older adults were healthy community dwellers. Payment was contingent on participant's performance: for each correct choice participants received 0.05 euro; on average, a total of about 3.5 euro each was paid plus a basic participation payment of 15 euro, making a total average of 18.5 euro per participant.

Design

The independent variable in this study was the strategy used in a particular block of trials (TTB vs. EAS; within subjects design). Each participant observed two blocks of 24 pair-comparisons each concerning 2 objects and up to 8 cue values on each object. To allow testing for age differences in performance variability, each block was composed of two presentations of the same set of 12 trials. On the second presentation the order of the objects was reversed on each trial and trials were placed in a different random order to prevent participants memorizing the decisions.

Measures

Two components of the mechanics of cognition (Baltes et al., 1999) were measured using standard working memory and speed of processing measures. In addition, participants' ratings of the cognitive demands of strategies were assessed.

Operation span. The stimuli of Hamm (unpublished) were used for this task. The version used included all items with 5 words (see Engle et al. 1999); thus, the possible score ranged from 0 to 48. Participants saw individual operation-word strings (e.g., IS (8/4)-1=1? bear). They had to solve the math problems, each of which was followed by a lowercase word, which was to be read aloud. On hearing the word "bear" the experimenter would press a key that would cause the presentation of the next string. After a set of these operation-word strings, participants recalled the words. The dependent measure was the cumulative number of words recalled from perfectly recalled trials.

Digit symbol substitution. Subjects had to write as many symbols as possible within 90s. The regular paper-and-pencil format was used (Lindenberger, Mayr, & Kliegl, 1993).

Estimates of cognitive demand of strategies. Participants' effort estimates were elicited using two 7-point Likert scales concerning each of the two strategies ("1 = not cognitively demanding, 7 = extremely cognitively demanding").

Procedure

Initially participants were familiarized with the paired comparison task, the concept of cue validity was explained, and the validity of the binary cues, as well as their direction was introduced. To match the simulation procedure described above all cues had equal validities, that is, they had equal predictive power. Participants then received extensive written instructions concerning two decision strategies, TTB and EAS, which included 3 detailed examples relative to each strategy. These examples encompassed both the correct sequence of acquisition and the final choices appropriate for each strategy.

Regarding TTB, participants were instructed to search for information in a cue-wise fashion and to make a decision as soon as a discriminating piece of information was found. In contrast, participants were instructed to apply EAS by performing an alternative-wise search, gathering information using weighted cue values, and performing a decision based on the final difference between scores. Note that because cues had equal validities the final decision of EAS was equivalent to both Franklin's Rule (FR) and Tallying. However, while FR implies weighting cue values and adding them, Tallying dispenses with cue weighting altogether. Tallying simply keeps track of positive cue values for each object, and makes a decision based on the difference between the two final tallies. Nevertheless, both versions of EAS require more information to be processes compared to TTB being, consequently, more cognitively demanding than the latter.

After learning how to apply the decision strategies, the participants were subsequently introduced to the computerized inference task. The task consisted of deciding, based on a set of cues, which of two diamonds was more expensive. Participants performed all decisions based on information search in a computerized display and were told to use a particular strategy in a given experimental block. An experimental block consisted of one practice trial, in which the experimenter made sure that the participant understood how to use the prescribed strategy. This measure ensured that all participants were able to perform the strategy equally well at the beginning of each set. Afterwards, 24 experimental trials followed in which the participants used the same strategy consecutively. The order of blocks was counterbalanced within age-group.

Although participants were instructed to use particular strategies their final decisions were otherwise unconstrained. Participants had at each step the choice between acquiring one piece of information about one of the two diamonds, up to a total of 8 cue values per diamond, or making a decision. Each cue value was presented briefly (2 seconds) and only once. The cues available to the participants were the following: size, overall proportions of the diamond, crown proportions, pavilion proportions, size of table, color, clarity, and certification laboratory. All cues had binary values (e.g., big vs. small diamond, colored vs. uncolored diamond). Both cue validities and cue labels were the same for all individuals. Participants had to touch appropriate buttons on a touchscreen to ask for information and make a decision. The order of information acquisition was partially constrained, with participants having to follow a predetermined cue order. However, participants had the possibility of choosing which alternative they wanted to find out more about at a particular time, and they were able to alternate between objects. After performing the decision task, participants' working memory and speed of processing were assessed, as well as their estimates of cognitive demands of strategies.

Results

In the results section, I first consider the main predictions outlined above. Second, I present results concerning information search in the decision task more generally. Finally, the relation between cognitive capacity and strategy application is assessed.

Age Differences in Mean Performance

The first prediction stated that older adults' mean performance should be worse than that of young adults, with older adults making on average more application errors than younger adults. Because a within-subjects design was used, and thus the performances using the two strategies were not independent, I conducted a multivariate analysis of variance (MANOVA) to test this hypothesis. In this analysis, the proportion of correct decisions when using TTB and EAS were used as dependent variables and age group (younger vs. older adults) as the independent variable. The main effect of age group was significant, $F(2, 19) = 4.04$, $p = .034$, partial $\eta^2 = .30$. As can be seen in Figure 3.4, young and older adults' performance does seem to differ at the level of mean performance. Nevertheless, note that older adult's performance was on average still considerably high ($M = 96.04$, $SD = .04$) compared to a close to perfect performance of younger adults ($M = 99.65$, $SD = .01$).

Age Differences as a Function of Strategy Difficulty

A second prediction was an age by strategy difficulty effect: It was predicted that the performance of older adults would show a larger decrement with more cognitively demanding strategies compared to younger adults.

Proportion correct decisions. Figure 3.4 shows that that the difference between young and older adults' performance when using EAS is not larger than when using TTB. To quantify these differences, I calculated contrasts between young and older adult's proportion of correct decision when using TTB and when using EAS. The difference between young and older adults' performance when using TTB was -.042, $F(1, 20) = 4.96$, $p = .030$, partial $\eta^2 = .20$, while when using EAS it was -.031, $F(1, 20) = 4.03$, $p = .058$, partial $\eta^2 = .17$. Also, please note that the differences in performance between younger and older adults represented rather small effects (partial $\eta^2 = .20$; partial $\eta^2 = .17$; TTB and EAS, respectively).

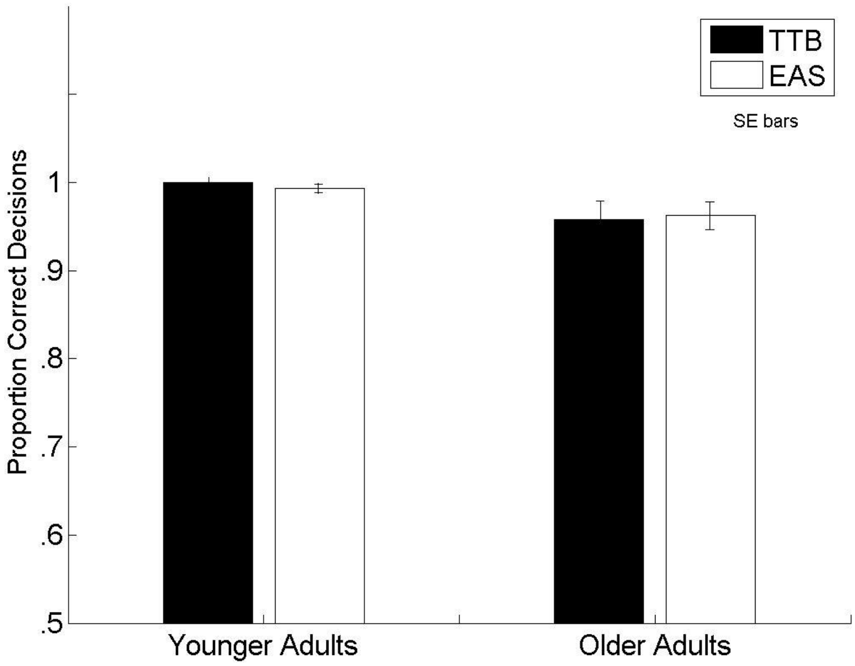

Figure 3.4: Proportion of Correct Decisions of Younger and Older Adults as a Function of Strategy Used. Error bars represent mean standard errors.

Estimates of cognitive demand of strategies. Although an interaction between strategy and age group did not emerge at the outcome level, individual estimates of effort could reveal that older adults found EAS a more cognitively demanding strategy compared to younger adults. Table 3.3 summarizes the estimates of cognitive demand of TTB and EAS for younger and older adults. An inspection of this table shows that this was not the case. Older and younger adults showed very similar estimations of effort when TTB is considered. However, concerning EAS, older adults on average reported this strategy as being less demanding in comparison with younger adults. One possibility is that younger and older adults used strategies differently. While younger adults made use of EAS with weighted cue values, older adults may have computed simple tallies to arrive at a decision.

Table 3.3: Means (and SDs) of Participants' Estimates of Cognitive Demand of Strategies and Tests of Significant Differences Between Younger and Older Participants

	TTB		EAS	
	M	*SD*	*M*	*SD*
Younger Adults (n=12)	1.08	.29	3.00	1.13
Older Adults (n=10)	1.30	.48	2.00	1.63
t tests	$t(20) = 1.30$, $p = .207$		$t(20) = 1.69$, $p = .106$	

Age Differences in Performance Variability

A third prediction concerned the existence of larger intra-individual variability in performance with increased age. Please recall that participants performed each decision twice, thus, it was possible to investigate this issue by calculating the difference between performance in a first set of trials compared to a second set. Comparing the two age groups, the difference in performance was significant, $t(20) = 3.60$, $p = .002$, $d = 1.5$, with older adults varying more between sets ($M = .10$, $SD = .09$) than younger adults ($M = .01$, $SD = .02$).

Age Differences in Search Behavior

Although the adopted modeling approach did not make predictions concerning search behavior, it was of interest to consider potential age differences in this respect. One robust finding in the aging literature is that older adults show significant increases in reaction time when performing cognitive tasks (e.g., Salthouse, 1996) and this result has also been recognized in previous research on the use of decision strategies (e.g., Johnson, 1990). To investigate the existence of such an effect, a MANOVA was conducted with age group as independent variable (younger vs. older adults) and the decision time when using TTB and EAS as the two dependent variables. A main effect of age group emerged, F (2, 19) = 5.28, p =.015, partial η^2 =.357, suggesting that overall older adults took longer to arrive at a decision. However, follow up univariate tests showed that this pattern held when TTB (M_{young} = 13.49,

SD_{young} = 3.84; M_{old} = 43.06, SD_{old} = 37.41) but not when EAS was considered (M_{young} = 16.03, SD_{young} = 6.24; M_{old} = 93.90, SD_{old} = 158.45).

Another important aspect was whether participants searched for the appropriate information. Assuming that older adults have more problems with processing large amounts of information one would expect that they show less information-intensive search behavior compared to younger adults. To analyze this issue I computed the number of cue values searched by young and older adults when using both TTB and EAS. As can be seen in Figure 3.5 both young and older adults searched close to exactly the information prescribed by the TTB strategy, and thus no age difference was found, $t(20) = .78$, $p = .446$.

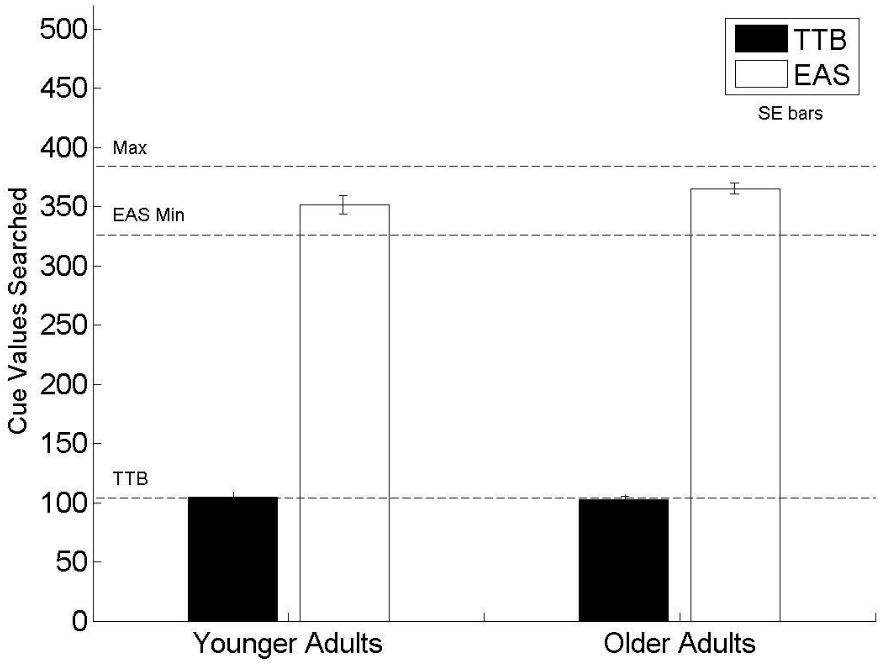

Figure 3.5: Cue Values Searched by Younger and Older adults as a Function of Strategy Used. The horizontal lines represent (respectively, from top to bottom) 1) the maximum number of cue values searchable, 2) the minimum cue values required to use EAS, 3) the minimum number of cue values required to use TTB.

The interpretation of results regarding EAS is less straightforward as there are two criteria concerning the correct amount of information to be considered. First, decision makers can look up all the information. However, it is also possible to use EAS without searching all the available cue values: There exists a point at which, having observed all cue values on one option, the decision maker knows that this option will not be overuled by the second one. For example, having identified that the first object has 8 positive cue values and knowing that the second object scores negatively on the first cue, it is already evident that no combination of cues will overturn the decision. As can be seen in Figure 3.5, both young and older adults' average search falls within the bounds of these criteria, indicating that participants sometimes viewed less than all information available to them but usually not less than the minimum required by EAS. The difference between age groups in what regards this variable is also not significant $t(20) = 1.47$, $p = .158$. In sum, the two age groups were remarkably homogeneous in the amount of information they search, with older adults not differing significantly from younger ones.

Cognitive Capacity Measures

Results of the individual differences measures are displayed in Table 3.4 for the younger and older samples. Compared with younger adults, older adults displayed slower speed of processing ($d = 1.56$) and lower working memory capacity ($d = .51$).

Table 3.4: Individual Difference Measures: Means (and SDs) and Tests of Significant Differences Between Younger and Older Participants

	WM		Speed	
	M	*SD*	*M*	*SD*
Younger Adults (n=12)	36.92	3.97	58.75	9.64
Older Adults (n=10)	34.80	4.83	45.40	8.04
t tests	$t(20) = 1.13$, $p = .272$, $d = .51$		$t(20) = 3.48$, $p = .002$, $d = 1.56$	

Relation between Cognitive Capacity Measures and Strategy Application

One aim of Study 4 was to gather further support for the idea that success in strategy application is related to basic components of cognition. This was investigated by, first, fitting the models' *G* parameter to mimic individual participants' behavior and, second, relating the fitted values to measures of working memory capacity and speed of processing.

Fitting procedure. To be able to assign a value of *G* to each participant, the relation between performance accuracy and *G* was determined by fitting a function linking the average performance of the successfully trained networks when using TTB and EAS and levels of *G*. This resulted in *G* being a nonlinear transformation of average performance accuracy when using the two strategies. The mean *G* values were significantly different between age groups, $t(20)=3.921$, $p = .001$, with older adults showing on average lower *G* values ($M = .78$, $SD = .14$) compared to younger adults ($M = .96$, $SD = .08$). Recall that the original average values of *G* used by Li and Lindenberger (1999) in their modeling efforts were .8 ($Min = .6$; $Max =$

1) and .3 (*Min* = .1; *Max* = .5) for younger and older adults, respectively. Thus, although the fitted *G* values for the younger adults from Study 4 are within the range originally proposed, the value of *G* for older adults is higher than the original. The latter result suggests that the performance of older adults would not have been fit by Li and Lindenberger's original modeling efforts.

Correlation analysis. After assigning a value of *G* to each participant, it was possible to compute the correlation between this parameter and the individual differences measures. Table 3.5 shows the correlation between *G* and cognitive capacity. For comparison purposes, I also report the correlations between these measures and proportions of correct decisions when using TTB and EAS.

Table 3.5: Correlations Between Fitted G and Individual Differences Measures (N = 22)

	WM	Speed
Proportion correct TTB	.14 (.54)	.12 (.58)
Proportion correct EAS	.41 (.06)	-.08 (.73)
Fitted *G*	.32 (.15)	.16 (.48)

Number is brackets represent *p* values

The proportion correct decisions when using TTB was not associated with the cognitive capacity measures. However, when proportion correct EAS was considered, a medium effect emerged between working memory and proportion of correct decisions (albeit this was not the case for the speed measure). These results suggest that while individual differences in working memory capacity were predictive of application accuracy when using EAS this was not the case when TTB was considered.

Regarding the relation between fitted *G* and cognitive capacity measures, there was a medium effect size concerning working memory but only a small one concerning speed, suggesting that the *G* parameter captured some individual differences in working memory but not speed.

As large effects were expected, the sample size used was rather small. However, considering a medium effect size concerning the correlations found here, a post hoc power analysis (Erdfelder, Faul, & Buchner, 1996) indicates that the actual power was low .54 (e.g., $r = .30$; $N = 22$; $\alpha = .05$). To obtain a reasonable power value of .85 the significance level should be reset, $\alpha = .16$. In the future, to properly test these effects a larger sample should be used (a priori power analysis: $r = .40$, power = .85, $\alpha = .05$, $N = 40$).

Discussion

Study 4 set out to test the predictions originating from the neurocomputational approach outlined above. These included 1) differences between young and older adults' mean performance when using decision strategies, with older adults making on average more application errors than younger adults, 2) an age by strategy difficulty effect, with the performance of older adults showing a larger decrement with more cognitively demanding strategies compared to younger adults, and 3) larger intra-individual variability in performance with increased age. The experiment made use of younger and older adults trained to use the frugal TTB and the more information-intensive EAS, and evaluated how accurate the two age groups were in applying the two decision strategies.

The results indicated that older adults made more application errors and that their performance was more variable across similar sets of trials compared to younger adults. However, no age by strategy difficulty effect was found, suggesting that older adults did not have more difficulties in using EAS than TTB compared to younger adults. One possible reason why this was the case is that older adults used EAS as Tallying, thus neglecting weighted cue values. Using weighted cue values implies multiplication and adding up fractions, which older adults may be at a disadvantage compared to younger adults (Siegler & Lemaire, 1997). In contrast, Tallying implies simple addition, an ability which is persevered in old age (Geary & Wiley, 1991). In this study, using weighted or unweighted cue values would lead to the same decision, consequently, it was not possible to investigate whether older adults preferred the less cognitively demanding version. However, the estimates of cognitive demands of strategies suggest that this may have been the case, as older adults on average considered the evidence accumulation strategy as less effortful compared to younger adults. Although this seems a plausible explanation for an otherwise puzzling finding, future studies should include discriminating cases between the two versions of EAS so that a more direct test of this hypothesis is possible.

To my knowledge, there exists only one other study investigating accuracy in the application of decision strategies. In this study, Bettman et al. (1990) reported that 11% of choices made by 7 college students were application errors as measured by choosing the wrong option prescribed by a strategy. Comparatively, the proportion of application errors identified in Study 4 for the younger sample is negligible. However, Bettman et al.

investigated application of strategies in decisions involving larger consideration sets, which may explain the discrepancy between findings. This possibility suggests that, in the future, studies should consider more demanding environments in their designs.

Using more demanding decision environments would also help understanding the relation between individual differences in cognitive capacity and strategy application. In this study, most younger and some older adults performed at ceiling level, making it impossible to reliably estimate the relation between individual differences in cognitive capacity and strategy application. However, heterogeneity in strategy application will most likely be larger in more demanding environments. Consequently, using larger consideration sets should allow a better exam of the role of individual differences, including those due to aging, in the application of simple lexicographic strategies similar to TTB, and other more information-intensive strategies.

General Discussion

Linking brain, mind, and behavior is a major aim of cognitive science (Gardner, 1985) and cognitive neuroscience (Cohen et al., 2002). The work put forward here contributed to this goal by introducing a neurocomputational approach that coupled the adaptive toolbox program (Gigerenzer et al. 1999) with a formal theory of aging (Li et al., 2000) to predict age differences in strategy application.

The remainder of the discussion will focus on the general merits of the approach, its limitations, and potential avenues for future research.

Theoretical and Empirical Contribution

Most efforts in bridging brain, mind, and behavior levels of description have focused on fairly low levels of cognition, such as vision (e.g., Sitton, Mozer, & Farah, 2000), memory (e.g., Li et al., 2000), and attention (e.g., Cohen & Servan-Schreiber, 1992; Braver et al., 2001). The approach presented here, however, joined recent attempts (Anderson, Bothell, Byrne, Douglass, Lebiere, & Qin, 2004; Polk et al., 2002) to consider brain-mind-behavior relations in the domain of higher order cognition, and adopted formal computational methods to understand this connection, a strategy largely underrepresented in cognitive neuroscience research (Chatterjee, 2005). Moreover, it did so in the service of understanding developmental issues, in particular, age-related cognitive decline.

In this respect, the neurocomputational approach provided clear predictions to be tested, which were partly supported in an empirical study. As predicted, older adults made on average more application errors and showed larger intra-individual variability in performance compared to younger adults. However, contrary to expectations, no age by strategy difficulty effect was found: Older adults did not show more difficulties with a more cognitively demanding strategy compared to a simpler one than younger adults, possibly because they used a simplified version of the more cognitively demanding strategy.

This last finding is particularly interesting as it raises that possibility that complexity cost effects are not pervasive in decision making research. For example, Finucane et al. (2005) also failed to find an age by task complexity effect in comprehension of decision problems. This of course contrasts with most aging research showing considerable age by complexity effects (Li et al., 2000). Nevertheless, the results presented here also suggest why this may be the case in decision making situations. In most decision tasks it is possible to adopt strategies that are relatively cognitively undemanding (e.g., TTB, Tallying). Thus, the results contribute to the view of *older adults as adaptive decision makers* which adapt their strategy use to their individual characteristics, namely, their cognitive limitations.

Potential Extensions to the Approach

In the next few paragraphs I consider the implications of relaxing some of the assumptions underlying the approach present above. First, the theory assumed that most of human cognitive aging is related to losses in dopaminergic function and can be modeled as reduced signal-to-noise ratio. As Band, Ridderinkhof and Segalowitz (2002) pointed out "more neurotransmitter changes take place than just the loss of dopamine receptors and more neural changes take place than changes in transmitter systems" (p. 264). Consequently, it may be interesting to explore other aspects of cognitive aging that may impact decision making abilities. Moreover, and although simplicity breeds tractability, the restricted scope of predictions suggests that the approach must be extended, for example, concerning age differences in information search.

Second, the approach addressed the specific issue of aging-related changes, making it a less than general model of the development of human decision making abilities. However, one could in principle extend the present approach to early cognitive development. Working memory and processing speed show an inverted-U shape function over the lifecourse (Kail & Salthouse, 1994; Siegel, 1994) and neuromodulation of prefrontal cortex activity is related to these basic capacities (Arnsten, 1998). Hence, it is conceivable that the methods applied here would be successful in modeling strategy use during early development, thus expanding the approach to the full scope of human ontogeny.

Third, while the approach capitalizes on decline in neural functioning being naturally modeled with subsymbolic systems (e.g., neural networks), it could potentially profit from the potential of symbolic systems to deal with complex problem solving and decision making tasks (see Polk et al., 2002, for a similar argument). One possible future direction would be to model strategy use using a hybrid system, such as ACT-R, which combines a symbolic production system with a set of massively parallel processes (Anderson & Lebiere, 1998). ACT-R is a particularly well-suited architecture for such enterprise for several reasons. First, some aspects of strategy use have already been investigated using this framework (Nellen, 2003; Schooler & Hertwig, 2005). Second, ACT-R's architecture has been shown to model both commission and omission errors (Lebiere, Anderson & Reder, 1994) which are hallmark effects of age-related cognitive decline. Furthermore, ACT-R could help when making predictions about brain functioning underlying strategy use (see Anderson et al., 2004, for applications of ACT-R to understanding brain function).

Finally, one future goal of the theory should be to make direct predictions concerning behavior based on changes in brain functioning. It has been shown that older adults under the effects of a dopamine agonist significantly improve their performance in cognitive tasks (Volkow, Gur, Wang, Fowler, Morberg, Ding, Hitzemann, Smith, & Logan, 1998) while younger adults under the effects of a dopamine antagonist significantly decrease their performance (Luciana, Hanson, & Whitley, 2004). Potentially, the approach could use measures of cognitive capacity to estimate model's parameters, such as *G*, to make a priori predictions concerning decision performance of younger and older adults under the effects of pharmacological drugs. Having a computational model whose parameters relate both to measures of cognitive constructs and decision behavior under pharmacological stimulation could provide new insights concerning the missing link between brain, mind, and behavior.