Personality and Social Psychology Review
2002, Vol. 6, No. 4, 357–369

Virtual Personalities: A Neural Network Model of Personality

Stephen J. Read
Department of Psychology
University of Southern California

Lynn C. Miller
Anneberg School for Communication
University of Southern California

A neural network model of personality is presented. The model has two goal systems: an approach system (BAS) and an avoidance system (BIS), as well as a system that governs the level of disinhibition/constraint (IS) in the two goal systems and the behavior system. Furthermore, within both goal systems, agentic and communal goals are specified. By tweaking the parameters of this system (e.g., chronic activation of goals, sensitivity of systems), and randomly or systematically varying situational arrays, distinct patterns of "behavior" by Virtual Personalities (VPs) across "situations" emerge that fit with classic distinctions (e.g., Big 5, temperaments). Various simulations demonstrate that VPs provide an exciting vehicle for integrating disparate approaches to personality to better understand the dynamics, situational responsiveness, and consistency of persons in situations.

Personality is stable and changing, responsive to situations and "set like plaster" (Costa & McCrae, 1994). The subject matter of personality is "all individual differences in the behavioral realm" (Jensen, 1958, p. 302). Personality is "an individual's…unique pattern of traits" (Guilford, 1959, p. 5). And, it is the "organized whole (system), that is constituted of parts or elements (subsystems), and separated somehow from an environment with which it interacts" (Sanford, 1963, p. 489). It is about that which is idiographic and unique and that which is nomothetic (Allport, 1962). It is all of these things, but it is not at all clear how these truths all come together. In short, the field of personality "itself has struggled with the task of conceptualizing the interplay between the stable and dynamic properties of personality systems" (Pervin, 1990, p. 4). And, "although we are all pretty much interactionists at this point, there remains considerable disagreement about what in the person interacts how with what in the situation"(p. 14).

Personality is complex. Ironically, to study individual differences in personality may require simulating it. Simulations afford a window on the behavioral consequences of underlying complex personality dynamics in the face of changing situations. Simulations could also show links between these dynamics and a structural views of traits, including the Big 5 (Miller & Read, 1987). With virtual personalities we could create and "tweak" the systems from which variability in personality might emerge. In doing so, we provide a framework and a methodology for more holistically grappling with the recurring issues historically central to personality. Thus, the first aim of this article is to present a neural network model of personality informed by what is currently known about the structure of personality and personality processes. Our second goal is to test our model of personality with an initial set of simulations and judge the initial utility of this model for understanding personality. In doing so, we hope to clarify how some of the seemingly disparate facets and definitions of personality might be dynamically linked.

Requests for reprints should be sent to Stephen J. Read, Department of Psychology, University of Southern California, Los Angeles, CA 90089–1061. E-mail: read@usc.edu

Developing Virtual Personalities

Understanding Personality: Integrating Structure and Process

Traits. One focus of the current simulation is to develop a model of the underlying psychological processes from which the structure of personality traits might emerge. As a starting point, we begin with the Big 5 analysis of personality trait structure. Although this structure is arguably not the last word, it provides a useful starting place that seems to capture a large part of the structure of personality. As several authors have noted, the Big 5 is a structural model, and not a model of the underlying psychological dynamics that are responsible for this structure (Pervin, 1990). Although some researchers seem to feel that it is "conceptually problematic" (Cervone, 1999, p. 329) and unnecessary to relate models of personality dynamics to personality structure (e.g., Cervone, 1999; Cervone & Shoda, 1999), others, including Mischel (1999) disagree. We would argue that models of personality dynamics should, in fact, give rise to the patterns of behavior and distinctions that have been encoded in the language of personality (Miller & Read, 1987). Therefore, a viable theory of personality can not leave these connections unaddressed. In this respect, our approach differs from that of Shoda and Mischel (1998) and Shoda, Tiernan and Mischel (this issue) who do not address the relationship between personality structure and personality dynamics in their model.

Traits and goals. The model we propose is heavily goal-based. In line with much of our previous work (e.g., Miller & Read, 1987, 1991; Read & Miller, 1989, 1993, 1998) we argue that most personality traits are fundamentally goal-based, that the goals which people strive to achieve or the outcomes that people strive to avoid provide the fundamental dynamics of personality. In fact, goals account for a tremendous amount of variability in trait judgements (Read, Jones, & Miller, 1990). One of our aims is to develop a model of personality where behavior is conceptualized as the emergent outcome of a set of goals, plans, resources, and beliefs made salient for the person in the situation. Moreover, we argue that behavioral responses to changing situations are the result of a competition among an individual's goals for the control of behavior. Further, putting personality and situational components in a similar language (e.g., of goals, plans, resources, norms), helps us to capture the fluidity of persons-in-situations that the personality versus situation debates had obscured (Miller & Read, 1987; Read & Miller, 1998).

According to this approach, individual differences in personality and behavior can be understood largely in terms of differences in the chronic activation of goals. But further, variability in behavior across situations and time can be understood in terms of how the features of these new situations and events affect the momentary activation of the individual's goals. For example, a helpful person is one who attends to another's needs, wishes to meet them, and has the resources to do so. Individuals for whom reacting to such cues frequently results in the behavior of helping another (compared to other behaviors), is likely to be viewed as a "helpful person:" a trait judgement. In essence, the trait label economically captures the script (who does what to whom, how and why): the recurring patterns of situational cues the person encounters and how he or she responds to them (Miller & Read, 1987). But, this pattern will more frequently impact behavior (e.g., produce more consistent helping behavior across situations) if other goals—with links to competing behaviors—are not as frequently activated for this person.

Agentic and communal goals. Casting traits in goal terms, makes them immediately more dynamic, and starts to shed light on the connection between structure and process. But, there are dozens of goals relevant to personality (Chulef, Read, & Walsh, 2001). How might these be organized to make a simulation more meaningful? As a number of researchers, but most prominently Wiggins and Trapnell (1996), have argued, human traits and motives, as represented in the Big 5 and the Interpersonal Circumplex (Wiggins & Trapnell, 1996), fall into two major classes: agentic and communal. This distinction has gone by several names, but the fundamental difference is between what might be termed individualistic, egoistic or "self-centered" motives versus collectivistic, selfless, "other-centered" or group-centered motives (Bakan, 1966; Spence, 1985; Triandis, 1995; Markus & Kitayama, 1991). Wiggins and Trapnell argue that this distinction can be found throughout the Big 5. They suggest that agentic and communal are 45 degree rotations of the Big 5's extroversion and agreeableness dimensions. Moreover, both agentic and communal components can be found in conscientiousness and neuroticism. Some aspects of conscientiousness are focused on achievement striving, whereas others are focused on duty and responsibility to others. And some aspects of neuroticism, such as fear of rejection, have a communal component, whereas others, such as fear of physical harm, have an agentic component. Other researchers have also provided evidence for the fundamental importance of this distinction in personality structure (e.g., Church, 1994; Digman, 1997; Tellegen & Waller, 1997).

Biology, traits, and systems. Another important aspect of our model is based on recent work on models of temperament. Temperaments are considered to be biologically based traits. Researchers (for a review see Clark & Watson, 1999) have argued for three main dimensions of temperament (extraversion/positive emotionality, neuroticism/negative emotionality, disinhibition /constraint), each with a biological basis in the operation of certain brain structures.

Based on considerable work with animals and humans, researchers have argued that there is one system that governs sensitivity to and approach to rewarding stimuli (Depue, 1996; Gray, 1987; Pickering & Gray, 1999). This has been variously termed a Behavioral Approach System (BAS) or a Behavioral Facilitation System (BFS). Greater sensitivity of the BAS should lead to greater sensitivity to reward cues and to greater striving toward rewarding stimuli. In this model, this forms the basis of extroversion. A second system governs sensitivity to and avoidance of punishing stimuli.

This has been termed the Behavioral Inhibition System (BIS). Greater sensitivity of the BIS should lead to greater sensitivity to punishment and threat cues, and to stronger attempts to avoid punishing stimuli. This is argued to be the basis of neuroticism[1].

Further, temperament researchers have argued that there is a third system, involving biological and heritable differences, that is responsible for disinhibition/constraint of behavior (Eysenck & Eysenck, 1975; Tellegen, 1985; Watson & Clark, 1993). We have labeled this the Inhibitory System (IS). Differences in the disinhibition/constraint system should affect the extent to which there is selectivity and constraint in enacting goal directed behavior. As a number of researchers in cognitive psychology and cognitive neuroscience have demonstrated (see Nigg, 2000, for a review), such inhibitory processes act to enforce selectivity in activation and focus.

Researchers have further argued that the BAS maps onto the Extroversion dimension in the Big 5, that the BIS maps onto Neuroticism, and that the disinhibition / constraint system maps onto Conscientiousness (Clark & Watson, 1999). This provides evidence for a possible biological basis for at least some aspects of 3 dimensions of the Big 5 and may provide some insight into why the underlying dynamics of personality result in the specific structures we find in the language of personality.

Thus, our model is based on the following assumptions: (1) there are two basic goal systems in the brain, one which governs sensitivity to reward (BAS) and one which governs sensitivity to punishment and threat (BIS). (2) There are individual differences in the sensitivity of these systems and these differences in sensitivity are independent of each other. (3) Within each goal system, there are two fundamental types of goals: Agentic goals and communal goals. (4) There are individual differences in the chronic activation or importance of these goals. These differences in importance may lie both at the level of the fundamental types, which could be related to cultural differences in individualistic versus collectivistic orientation, as well as at the level of individual goals. (5) There is a general disinhibition/constraint system, which provides an overall inhibitory "field" or constraint on the level of activation of the goals and the behaviors. With high inhibition, only the most highly activated constructs can become active, whereas with low inhibition a much wider range of constructs can become active. We argue that there are individual differences in such inhibition, which lead to individual differences in focus and behavioral constraint (Nigg, 2000). (6) This inhibitory system influences both the activation of goals and the activation of behaviors, in response to goal activation. These assumptions provide the basis for our model of the underlying dynamics of 3 of the Big 5 dimensions. In other work, we are examining how the remaining Big 5 dimensions, Agreeableness and Openness to Experience, can be incorporated.

Historically, psychologists have often been constrained in their ability to incorporate diverse theoretical distinctions into a single model. When our tools (e.g., 2 × 2 ANOVAs) limited us to manipulating a few variables, simple theoretical models became a virtue born of methodological necessity (Read, Vanman, & Miller, 1997). Nevertheless, the cost was steep: A more unified psychology seemed far beyond our grasp, or even our imagination. But, dynamic approaches—including connectionist models—can accommodate a multitude of factors and free us to explore the richness of their resultant interactions.

Overview of Neural Network Models

Before we present the simulations we first briefly review the general characteristics of neural network models. Neural network models have three important components: the network architecture, the learning rule, and the activation updating function (Bechtel & Abrahamsen, 1991; Ellis & Humphreys, 1999; McClelland & Rumelhart, 1986). The *architecture* refers to the units or nodes in the network and the patterns of connections among them. The nodes are neuron-like units which sum the activations from a number of inputs and then output their own activations as a function of the summed inputs.

The patterns of connections differ for different networks. Some networks are fully connected, whereas in other networks, each node is only connected to a subset of the other nodes. Further, the strengths of the connections can vary. Finally, some connections are unidirectional, whereas other connections are bidirectional. Bidirectional connections result in feedback relationships. Networks with bidirectional connections are referred to as recurrent network.

Processing in a neural network proceeds by presenting an input pattern to the network and then allowing activation to flow among the nodes until the activations of the nodes stop changing. The way in which activation flows in the network is a function of the patterns of connections among the nodes.

The *learning rule* dictates how connections among pairs of nodes are changed during learning in the network. Typically, the learning rule allows the network to learn the patterns of associations among the activations of nodes in the network.

[1]The distinction between an approach and an avoid system also maps onto Higgins recent distinction between promotion focus and prevention focus in goal striving. The basic notion is that frequently what looks like the same goal can be framed either in terms of promotion or prevention. Thus, performance on a test can be framed as either achieving success or avoiding failure. Higgins has shown that these two foci lead to systematically different psychological outcomes (Higgins, 2001).

Finally, the *activation function* determines how activation propagates through the network. This function takes the summed inputs of all the nodes coming into the node and then uses that to compute the output activation of the node. In some cases, the output activation is simply a linear transformation of the summed inputs. However, more frequently the output activation is a nonlinear transformation, typically a sigmoid or S-shaped function. Such non-linear, S-shaped functions have a number of desirable characteristics for processing in networks, as well as being more plausible neurologically then a linear transformation which would lead to unrealistically high output activations.

Network Description

The network architecture is depicted in Figure 1. The model is designed in the PDP++ program (O'Reilly, Dawson, & McClelland, 2001), which is available for several versions of Unix, Linux, Windows, and MacOSX (http://www.cnbc.cmu.edu-/PDP++/PDP++.html). Our model uses O'Reilly's (O'Reilly & Munakata, 2000) Leabra architecture, which is one of several architectures available in PDP++. Leabra is designed to be much more neurally based than most neural network architectures. An extensive introduction to it, with numerous examples can be found in O'Reilly and Munakata.

Layers. The network has 4 layers, an input or situation layer, two goal layers and the behavior layer. There are two separate goal layers, one to model the approach system and the other the avoidance system. In this version there are four goals in each of the two systems, with two goals being communal goals and two being agentic. Obviously a complete model would represent many more goals, but this was sufficient to capture the basic behavior of the model.

The input layer has 16 nodes, 8 designated as situational features that activate a corresponding goal and 8 that designate a corresponding resource for the achievement of that goal. This was done so we could compare the impact of the joint activation of both a goal and resource with the activation of the goal alone. However, we have not yet examined the role of the resource layer.

Connections. For simplicity sake, in the current network all connections are either .5 or 0. In future versions of the model, which will incorporate learning, the connection strengths will vary through the entire range from 0 to 1. The 8 goal feature nodes in the input layer have bidirectional connections (indicated by double headed arrows) only with their corresponding goal in the goal layer. Not only does the increased activation of an input feature increase the activation sent to the corresponding goal, but the increased activation of a goal can increase the activation of its corresponding input feature. Thus, a highly activated goal can bias the likelihood that a corresponding input feature is "perceived." These bidirectional connections make the network recurrent and enable constraint satisfaction.

The 8 resource nodes have unidirectional connections (indicated by single headed arrows) directly to their corresponding behavior node, but have no connections to the goal layer. This was done to capture the notion that the availability of resources should not typically affect the activation of a goal, but should affect the likelihood of enacting a behavior, given that the resource was available.

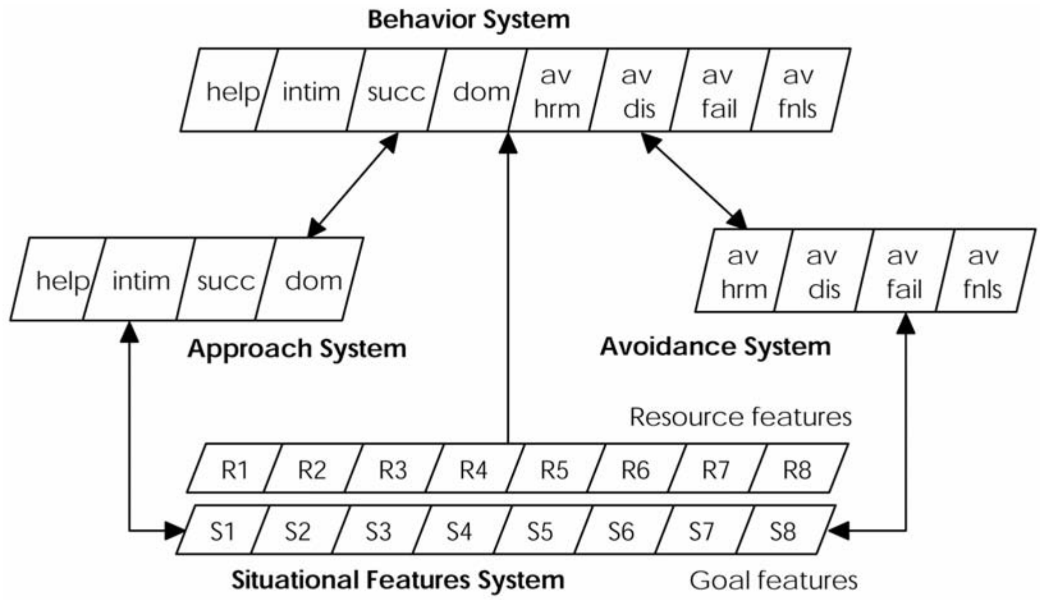

Figure 1. Virtual Personality: A recurrent neural network model of personality. Labels: help = help, intim = intimacy, succ = success, dom = dominance, avhrm = avoid harm, avdis = avoid disappointment, avfail = avoid failure, avfnlss = avoid financial loss.

The goal layers have bi-directional connections to the behavior nodes (again indicated by bidirectional arrows), so that the two layers can mutually influence one another. This also allows us to capture the idea that being able to enact a behavior may increase the activation of its corresponding goal. The bidirectional connections between goals and behaviors add to the recurrent nature of the network. Such models are consistent with feedback mechanisms in self-regulatory models (e.g., Carver & Scheier, 1998).

Biases. Differences in the chronic or resting level of activation of a goal are captured by connecting the goal node to a bias node that is activated at 1.0 (bias nodes are not represented in Figure 1). Differences in the weight from the bias to the goal node represent the degree or level of chronic activation. Thus, a bias weight of .3 will create a higher resting or chronic level of activation on a goal, than a bias weight of .05. It is standard to use bias nodes in neural network models to represent differences in the chronic level of activation of nodes generally. Here, however, we use these bias nodes only to represent chronic individual differences in the activation level of the goals in the goal layers.

Activation equation. Although the activation dynamics for a Leabra neuron have a number of specific features designed to capture many of the characteristics of the behavior of real neurons, many of those characteristics can be ignored here. For our purposes, two unique aspects of the neuron model are important.

First, as in real neurons, there is a threshold for firing. Only if the activation of a neuron exceeds that threshold does the neuron start to "fire" or send activation to other neurons. Thus, neurons that do not exceed the threshold have no influence on other neurons. In the formula below, V_m represents input to the goal node, which is a function of the bias on goals (from the bias node) and other inputs to the goals. It must exceed the threshold of the neuron (Θ) to fire.

Second, each neuron has a gain or sensitivity setting impacting the output activation of the neuron. Higher gain means that the activation growth curve is much sharper or more steeply accelerated, since output is a multiplicative function of the gain and the amount by which input exceeds the neuron's firing threshold. At very high gains, the node behaves in an essentially binary fashion, being close to either 0 or 1.

The activation function can be written as:

$$y_j = \gamma[V_m - \Theta]_+ / (\gamma[V_m - \Theta]_+ + 1)$$

where y_j is the resulting activation, V_m is the activation level of the neuron, Θ is the threshold value for the neuron, γ is the gain setting for the neuron, and the + subscript means that y is set to 0 if this component is negative.

The resulting activation from this equation has a roughly sigmoidal or S-shaped form, as is true of other activation functions commonly used in neural networks. This nonlinear form plays a critical role in the ability of the network to capture a variety of phenomena (Read et al. 1997).

The gain parameter can be varied to capture differences in the sensitivity of a single neuron or a set of neurons. Here we only use the gain parameter to model differences in the sensitivity of the approach and avoid systems, that is, differences in the sensitivity of the system to rewards (approach) and punishment (avoid).

Inhibitory processes. A major component of this model is an inhibitory system (IS) that allows only the most highly activated nodes in a layer to fire, but dampens the rest. The basic function of inhibitory processes in many instances is to enforce focus or selectivity in firing (i.e., of goals and behaviors). A strong inhibitory "field" means that only the most strongly activated nodes can fire, whereas more weakly activated nodes will be held below threshold.

Such inhibitory processes play a major role in a range of cognitive processes (Martindale, 1991; Nigg, 2000), allowing for selectivity and focus in a variety of tasks. Further, as we noted previously, a number of temperament researchers (see Watson & Clark, 1999 for a summary) have argued that a dimension of disinhibition/constraint is a fundamental dimension of temperament.

Although this inhibitory effect could be produced by having a large number of inhibitory inter neurons in a layer connected to all nodes, O'Reilly and Munakata (2000) note that doing so incurs a very large computational cost in large models. To reduce this cost they have developed the kwta algorithm (kwta stands for "k winners take all") which captures the essential behavior of this inhibitory system without the high computational costs. This kwta algorithm results in a strong separation between the activation levels of strongly and weakly activated nodes. In the absence of such inhibition, it is much easier for all activated nodes, even weakly activated ones, to fire. As a result, there is less discrimination or separation between them.

The kwta algorithm captures this inhibitory effect, by requiring that no more than k nodes in a specified group can be activated at a time, and, in fact, fewer than k may be activated if their activation is insufficiently strong. Using this algorithm, the number of nodes that can be activated in a specified group can be set to a specific number. This algorithm enforces the same kind of selectivity and focus that an inhibitory "field" would. However, the kwta algorithm enforces a fairly rigid selectivity. Thus, O'Reilly and Munakata (2000) have also developed an average kwta algorithm, which enforces a less stringent selectivity, by specifying that "on average" only k or fewer nodes will fire. This algorithm allows for more variability in the number of nodes that can be active, depending on how many are strongly activated.

We use the kwta algorithm to capture individual differences in the degree of disinhibition versus constraint. We do this by manipulating (1) the number of nodes that can be active in both the approach and the avoidance system, as well as in the behavior system, and (2) whether we use the more stringent kwta algorithm or the somewhat less stringent, kwta average algorithm.

Summary. The various aspects of the personality system are captured as follows. We attempt to capture three major dimensions of temperament, which have also been mapped onto 3 of the Big 5 dimensions. First, a BAS is modeled, which governs sensitivity to reward (central to extroversion). Second, a BIS is modeled, which governs sensitivity to punishment (central to neuroticism). And, third, a disinhibition/constraint system (IS) is modeled, which governs the selectivity of activation of the other two systems (impacting conscientiousness).

The BAS and the BIS are modeled by two independent layers, each composed of qualitatively different goals. Differential sensitivity of these systems to reward and punishment, is modeled by manipulating the gain parameter for the nodes in the corresponding layer.

Differences in the IS are modeled by differences in the general amount of inhibition in the network. This is implemented by O'Reilly's (O'Reilly & Munakata, 2000) kwta algorithm; reducing the number of possible active nodes corresponds to increasing constraint. We use both the more stringent kwta algorithm and the less stringent average kwta algorithms in order to capture the impact of differences in the degree of inhibition.

In addition to representing three major dimensions of personality, we also represent two other major dimensions of personality: agentic and communal. Wiggins and Trapnell (1996), in particular, have argued that this distinction crosscuts all of the major dimensions of the Big 5.

The agentic–communal distinction is captured by representing two separate sets of goals within both the approach and the avoid system. Individual differences in these sets of goals are represented by differences in chronic activation, which is captured by differences in the bias weights to the goal nodes.

Simulations: Testing the Utility of the Virtual Personalities (VP) Model

Personality psychologists regard two features as critical for ascribing a personality characteristic to people. First, people should differ on that characteristic, showing different behaviors in response to the same situation. Second, people should evidence cross-situational consistency in behavior, showing similar behavior in similar situations. Some personality psychologists have identified an additional aspect of personality: individual differences in the extent of cross-situational consistency. This can be thought of as a meta-level aspect of personality. In the following we demonstrate how our model can capture each of these aspects of personality.

Our first two simulations demonstrate how different virtual personalities (VPs) can respond differently to the same situation. Different people often respond differently to the same situation, in part, because situations contain an array of information and cues, and different people respond to different cues. What are the mechanisms that might lead different VPs to respond differently to a series of situational arrays? Two possibilities are that: (1) different personalities differ in the sensitivity of their BIS and BAS, and (2) they differ in the chronic activation of individual goals. In our first simulation, we compare the behavioral outputs of individuals who differ in the chronic levels of sensitivity of their personality systems to punishment (BIS) and reward (BAS). Our second simulation is similar, but instead of manipulating the BIS/BAS of our VPs, we manipulate the chronic activation of specific goals, specifically Help and Avoid harm. In these simulations, we hope to show that the relative sensitivity of the two systems and the chronic level of activation of goals will influence which cues are responded to. In both simulations, we manipulate inputs relevant to the target goals, Help and Avoid physical harm, and set the rest of the situational inputs to 0.

Our third simulation investigates whether differences in the sensitivity of the different systems and differences in the chronic activation of goals can result in consistent trends in behavior across a wide array of situations (cross-situational consistency). In everyday life, each situation contains an array of information, often relevant to a variety of goals. But how best to capture the seemingly random combinations of salient situational features which occur such that two "situations" are rarely exactly the same? To do so, in the third simulation, we generated a large number of situations, which randomly vary in the strength of all the situational inputs, and then examine whether there are consistent trends across those varying situations as a function of the VP's BIS/BAS sensitivity and chronic goal activation.

In the fourth simulation, we manipulate the degree of inhibition in the system (IS). It is expected that this will influence the extent to which the system's behavior is focused and responsive to the situation and thus, the extent to which the individual exhibits cross-situational consistency in personality. We expect that the greater the inhibition, the less sensitive the system will be to distracting cues. We model this by developing vectors with a central, highly activated situational feature and then adding random noise to all elements in the vector. We expected to find that VPs with higher IS responded more consistently to the same situational inputs and were less influenced by random distractions from situational features that receive some, but not as much, activation (random noise).

Simulation 1: System Sensitivity Simulation

The purpose of this simulation was to examine the impact of individual differences in sensitivity between the approach and the avoid system. To examine this, we had one VP in which the gains for the BAS and BIS systems are both set at 100, simulating an individual who is equally sensitive to reward and punishment cues, and a second VP in which the gain for the BAS was 100 and the gain for the BIS was 200. This second VP should be more sensitive to punishment than to reward cues and thus, can be thought of as somewhat higher on Neuroticism than the first VP. Inhibition for the different layers, which is manipulated through the kwta algorithm, is given in Table 1. We list whether the kwta average algorithm is used or the more stringent kwta algorithm and we list the number of nodes that can be activated. The table also lists the default gain levels of the four layers.

As part of this comparison, we hoped to show that higher gain for the avoidance system of the second VP could actually override the impact of higher levels of input to the approach system. That is, a more sensitive avoidance system might respond more strongly than the approach system, even though the approach system had stronger inputs from the situation. To put this more concretely, an individual who is more fearful (more sensitive avoidance system) might continue to see fear cues and engage in fearful behaviors at a point where a less fearful individual would not.

To demonstrate this, we created a series of 8 vectors, in which all the input nodes were set at 0, except for the help and avoid harm features. The input activation for the avoid harm feature, which is in the more sensitive avoidance system, starts out equal to the activation of the help feature (see Table 2). But, it then slowly decreases across the 8 patterns, while the activation of the help feature stays constant. We used this input pattern for each of the two VPs.

Table 1. *Parameters for Simulation 1: System Sensitivity Simulation*

Parameters	Situation	Approach	Avoid	Behavior
	Parameter Settings for Different Layers			
Inhibition				
Number of nodes active	K = 1	K = 2	K = 2	K = 1
Algorithm used	Kwta avg inhib	Kwta avg inhib	Kwta avg inhib	Kwta inhib
Gain	300	100	100 or 200	400

Note: Kwta = K winners take all.

The results for the two VPs can be seen in Table 2. Let us first look at what happens when the two systems have equal gains. When the input activation to avoid harm is equal to that for help, or even when it drops slightly to .66 compared to .71, then the outputs for the behavior nodes are low and essentially equal. However, once the avoid harm activation drops to .61 there is a sudden shift in activations and the behavior output for help jumps to .99 and the output activation for avoid harm is 00. When the activations to the two goals are close, the system does not really distinguish between them, but once the activations differ sufficiently the competitive dynamics of the system result in high level of activation for the most strongly activated goal and its associated behavior. Such nonlinear behavior is a result of the competition among nodes. Sudden shifts and non-linearity in behavior are important features of the model, and dynamic systems theories (e.g., Catastrophe Theory) more generally (Poston & Stewart, 1978; Thom, 1975; Zeeman, 1976).

But now lets look at the second VP, where the gain for avoidance is higher than the gain for approach. Initially, for the first 3 input vectors the avoid harm behavior is reasonably highly activated and more highly activated than the help behavior, even when the avoid harm goal is less highly activated, .61 versus .71. Thus, this VP, because of the higher gain of its avoidance system, responds strongly to avoid harm cues, whereas the first VP does not. Note also that for the more "neurotic" VP, for input 3 (.71 vs. .61) the avoid harm behavior is highly activated, whereas the help behavior is not. In contrast, for the first VP, for input 3 the output activations are the reverse, with the help behavior being highly activated. Thus, the higher sensitivity of the avoidance system in VP 2 overrides the weaker activation of one of its goals. One way to think about this, is that an individual who is

Table 2. *Simulation 1: Results for system sensitivity simulation*

			Output (Behavior)			
	Input		Virtual Personality 1 BAS Gain 100, BIS Gain 100		Virtual Personality 2 BAS Gain 100, BIS Gain 200	
Pattern Number	Help	Avoid Harm	Help	Avoid Harm	Help	Avoid Harm
1	0.71	0.71	0.21	0.21	0.11	0.59
2	0.71	0.66	0.22	0.21	0.11	0.59
3	0.71	0.61	0.99	0.00	0.12	0.59
4	0.71	0.56	0.97	0.00	0.97	0.00
5	0.71	0.51	0.97	0.00	0.97	0.00
6	0.71	0.45	0.97	0.00	0.97	0.00
7	0.71	0.40	0.97	0.00	0.97	0.00
8	0.71	0.35	0.97	0.00	0.97	0.00

more sensitive to threat (higher gain on the avoidance system), will see a threat where another individual would not and will act to avoid harm. This can happen even when a less "neurotic" individual would respond to the same cues by helping.

However, once the difference between the two inputs is sufficiently large, .71 versus .56, the stronger input to help overcomes the stronger gain of the avoidance system. And at this point, there is a sudden shift, with the helping behavior becoming highly activated (.97), and the avoiding harm behavior reaching 0.

Simulation 2: Chronic Goal Activation Simulation

In the first simulation we examined individual differences by manipulating the activation of whole systems (BIS or BAS via manipulation of gain). Here we wished to show that individual differences can also be captured by differences in the chronic activation of a goal, such that a goal that is chronically activated would be more sensitive to situational inputs than one that is less chronically activated. We compared two VPs, which differed in the chronic activation of one of their goals. We implemented higher chronic activation for the second VP by setting the bias weight for the avoid harm goal node at .3 and setting the bias weights for the other goals at .05. The input vectors were created as in the preceding simulation. The gains for the various layers are as in Table 1, with the avoidance system set at the same gain as the approach system, 100.

As can been in Table 3, differences in the chronic activation of the avoid harm goal affected the point at which the help behavior became highly activated. When the avoid harm goal had a higher bias of .3, it took a larger difference between the inputs before the help behavior became highly activated. One way to think about this is that when the avoid harm goal is more chronically activated, the VP is more sensitive to threat cues, which tends to block the impact of cues to helping. Thus, the second VP will be less likely to help, because it is more likely to be trying to avoid danger.

Table 3. *Simulation 2: Results for chronic goal activation simulation*

	Input		Output (Behavior)			
			Virtual Personality 1 (All Biases = .05)		Virtual Personality 2 (Avoid Harm Bias = .3, All Others = .05)	
Pattern Number	Help	Avoid Harm	Help	Avoid Harm	Help	Avoid Harm
1	0.71	0.71	0.21	0.21	0.20	0.27
2	0.71	0.66	0.22	0.21	0.20	0.27
3	0.71	0.61	0.99	0.00	0.20	0.27
4	0.71	0.56	0.97	0.00	0.97	0.00
5	0.71	0.51	0.97	0.00	0.98	0.00
6	0.71	0.45	0.97	0.00	0.98	0.00
7	0.71	0.40	0.97	0.00	0.98	0.00
8	0.71	0.35	0.97	0.00	0.98	0.00

Also notice that as in the preceding simulation, as the difference between the two inputs increases, there is a sudden, nonlinear shift in the output activations. As in the preceding system sensitivity simulation, this highly nonlinear response is due to the competition among nodes, particularly in the behavior layer.

Simulation 3: Cross-situational Consistency in Behavior

Consistency in behavior across a wide range of situations is one defining characteristic of personality. In this simulation, we examined whether differences in the sensitivity of a system, as well as differences in the chronic activation of goals, would lead to such consistent trends. To examine the impact of differences in sensitivity of systems to different situations, we first set up a network in which the BIS was more sensitive (higher gain of 200) than the BAS (gain of 100). To examine the impact of differences in chronic goal activation on responses to a wide range of situations, a single goal—Avoid harm—had a higher bias weight (.4 versus .05) than the other goals. Finally, to simulate the "highly varied" patterns of multiple-goal-relevant features active in situational arrays, we had the network respond to 400 input vectors that were randomly generated with uniform noise, with a mean of 0 and a variance of 1. The rationale was that this would be akin to a random sampling of situations and that the more sensitive system and the more chronically activated goal would be more likely to respond to the "random sample of situations" with a consistent trend across situations.

Because output activations tended to be either near 1 or 0, to make the results easier to look at, we simply counted the number of times that each of the eight behavior nodes had an activation equal to or greater than .5. Over the 400 random vectors, we found that the four approach behaviors were activated 38, 31, 37, and 31 times each, so that 34% of the time (137/400 = 34%), one of the approach goals fired. In contrast, the four avoid behaviors were responded to 89, 58, 55, and 58 times each, so that 260/400 = 65% of the time, one of the avoid goals fired. If we exclude the avoid harm behavior, to get rid of the effect of its higher bias, the other three avoid behaviors are still activated more frequently than the approach behaviors, 57% of the time. So when the BIS was more sensitive to cues than the BAS, its associated behaviors were much more likely to fire.

We also found that the chronic level of activation of a goal did indeed influence its likelihood of activation. The avoid harm goal, which had a higher bias of .4, was activated the most times of any goal, 89, versus 58 for the next closest goal. Note also that the effects of gain and bias were at least roughly additive.

Simulation 4: Impact of Inhibition on Behavioral Consistency Across Situations

Higher levels of inhibition should lead to greater selectivity and focus in behavior and should be associated with better correspondence between stimulus input levels for a target stimulus and target behaviors. Random noise and distractions should be less likely to activate corresponding non-target behaviors.

To test the impact of inhibition, we generated 800 input vectors of the following form. First, we created 8 input vectors, each with an activation of 1 for a feature corresponding to a different one of the 8 goals. We then created 100 copies of each vector, for a total of 800 vectors, and then added uniform random noise, with a mean of 0 and a variance of .6, to each vector. Thus, for each vector the feature that corresponded to the target goal was highly active (plus or minus the random noise) and other elements in the vector were also randomly activated, although typically not as strongly.

The gain for each of the different layers is the same as in the chronic goal simulation earlier (see Table 4 for all parameters for this simulation). O'Reilly's (O'Reilly & Munakata, 2000) kwta algorithm was used to capture the effect of differing levels of inhibition. In the high inhibition condition, we set the number of nodes in a layer that could be active to 1 for the situation layer, the two goal layers, and the behavior layer. Further, we used kwta average inhibition for the situation layer and the two goal layers. This version of the algorithm is somewhat less stringent than the kwta inhibition version, making it easier for more than k nodes to be activated if they have sufficient input. However, the behavior layer used the more stringent kwta inhibition.

There were two lower inhibition conditions (intermediate and low). In both, for the situation layer and the goal layers, we used the kwta average inhibition algorithm, just as we did in the high inhibition condition. However, for these simulations we set the number of nodes that could be active in a layer to 2 for the Situation layer and the two goal layers (rather than 1, as in the high inhibition condition), while keeping the Behavior layer at 1, as in the high inhibition condition.

Within the lower inhibition conditions, we also decided to see what would happen if we varied the inhibition in the behavior layer. To do this, we used the more stringent kwta algorithm in one version of the simulation (intermediate inhibition) and the less stringent kwta average version in the other (low inhibition). So we had two low inhibition conditions, which varied in the inhibition on the behavior level.

For each goal node, there were 100 vectors designed to primarily activate that node, with random activations for the other elements in the vector. To examine the results, we counted the number of times that a target behavior was activated by the 100 corresponding vectors and the number of times that one of the other behavior nodes was activated. We then summed this across the 8 sets of vectors to get the total number of times a target behavior was activated and the total number of times a non-target behavior was activated.

As can be seen in Table 5, the level of inhibition did make a large difference in the selectivity of the network. Interestingly, within the two lower inhibition simulations (intermediate and low), the amount of inhibition in the behavior layer also made a difference.

Table 4. *Simulation 4: Parameter setting for simulation of impact of inhibition on behavioral consistency across situations*

		Parameter Settings for Different Layers			
Simulation	**Parameters**	**Situation**	**Approach**	**Avoid**	**Behavior**
High Inhibition; high inhibition on all layers	Kwta	K = 1 Kwta avg inhibition	K = 1 Kwta avg inhibition	K = 1 Kwta avg inhibition	K = 1 Kwta inhibition
	Gain	300	100	100	400
Intermediate Inhibition; only behavior layer is high inhibition	Kwta	K = 2 Kwta avg inhibition	K = 2 Kwta avg inhibition	K = 2 Kwta avg inhibition	K = 1 Kwta inhibition
	Gain	300	100	100	400
Lowest Inhibition; low inhibition at all layers	Kwta	K = 2 Kwta avg inhibition	K = 2 Kwta avg inhibition	K = 2 Kwta avg inhibition	K = 1 Kwta avg inhibition
	Gain	300	100	100	400

Note: Kwta = K winners take all.

The high inhibition VP is highly responsive to the situation, without being distracted by "random noise". Although the high inhibition VP and the lowest inhibition VP (with low inhibition on all layers) do not differ very much on the frequency with which target behaviors were activated (89.75% versus 94%), they differ dramatically on the frequency with which non-target behaviors were activated, 5% versus 29.38% respectively. Essentially, the lowest inhibition condition involves much less selectivity and greater firing in response to both target and non-target activations. In this case, the lowest inhibition system looks like an individual who is unable to focus, but seems to bounce around in response to small changes in the situation. Disorganized or impulsive individuals might fit this pattern.

The intermediate inhibition system, with higher inhibition on the behavior layer, has some interesting differences from both of the other simulations. First, although it is similar to the high inhibition system in terms of the number of non-target behaviors activated, approximately 5%, it is less likely to activate target behaviors, 80% versus 89.75%. This system is less likely than the lowest inhibition system to perform non-target behaviors, but it is also less likely to perform target behaviors than the high inhibition system. This looks like an individual who is less likely to enact situationally responsive behaviors. Less efficient and less productive individuals might fit such a pattern.

As this simulation makes clear, inhibition processes in this model can operate separately at multiple layers of the system. The three layers of the simulation we present roughly map onto a sensory layer, a goal layer, and a motor/behavior layer. Prior conceptions of disinhibition/constraint in the temperament literature have focused on qualitative differences in general inhibition in the system (high or low) but not on different inhibitory processes at different points in the system. However, whether biologically or neurochemically such inhibitory layers are differentially set (e.g., via different neurotransmitters) remains an open question. Therefore, simulations that produce behavioral effects that mimic real life human variability might provide clues as to where different inhibitory mechanisms impacting behavior might be operating—potentially supplying hypotheses for neurobiological/chemical exploration. These simulations are also consistent with other work on automatic (versus controlled processes) that suggests, in line with considerable work in cognitive psychology (Martindale, 1991; Nigg, 2000), that many, if not most, inhibitory processes impacting behavior operate non-consciously.

Although it is unclear whether inhibitory processes are different at different levels, there is intriguing evidence that *excitation* is enhanced in different portions of the system via different neurotransmitters. For example, sensory excitation is enhanced with higher levels of acetylcholine and motor system excitation is increased by dopamine (Panksepp, 1998).

Table 5. *Simulation 4: Percentage of activation of behaviors for target and non-target behaviors under high and low levels of inhibition*

	Percentage of Activation of Behaviors	
	Target Behaviors	**Non-Target Behaviors**
High Inhibition	89.75%	5%
Intermediate Inhibition	80%	5.5%
Lowest Inhibition	94%	29.38%

Discussion

With the current, relatively simple model of personality, a surprising amount of the dynamics of personality can be simulated and seems to "fall out" of the model. First, we discuss the implications of the model for how "tweakable" VPs are apt to vary in personality consistency and situational responsiveness. Second, we discuss how, by "tweaking" the underlying personality system (BIS/BAS, IS or specific goals), patterns of behavioral output resembling trait-like behavior (e.g., Big 5, temperament, and other traits) emerge. Finally, we discuss how the VP Model is a "work in progress" and what future modifications of the program are planned.

Personality Bias, Cross-situational Consistency, and Situational Responsiveness

To the extent that there are biases in the underlying personality system, individual VPs—and we would argue, individuals in everyday life—will tend to be more consistent across situations in their behavioral responses, as we saw in simulations 1, 2, and 3. Of course, the power of situational cues often overwhelms those biases—as we would expect from a non-rigid, normal individual. Moreover, to the extent that there are lower biases in the system, the person would be expected to respond based more on the activation level from the situational array (e.g., go with the most heavily activated situational features).

How consistently the individual responds to the situation, as we saw in simulation 4, depends upon the IS system. An individual with a high IS would respond more consistently to important situational cues and be less distracted. In contrast, an individual with a low IS would be more responsive to random noise, and thus, be inconsistent across situations. Such a VP might capture some aspects of a hyperactive individual.

Producing the Big 5 and Other Traits From Personality Dynamics

A variety of well-known traits could be represented in the model. For example, Neuroticism was captured by a higher gain applied to the BIS (see simulations 1 and 3). Extroversion could similarly be modeled via higher gain for the BAS. Conscientiousness is modeled by setting the IS system at a higher strength for the situational cue layer, the two goal layers, and the behavior layer. Such VPs enact higher levels of behavioral output that is responsive to situational input. They are also more selective in situational input, that is, less likely to be "distracted" by situational noise.

Two types of low inhibition individuals were also generated in the current simulation by varying the inhibition separately at the behavioral layer, and at the situation and goal layers. The lower inhibition on all layers resulted in a VP who might have seemed disorganized and hyperactive (generating a high level of simultaneously competing behaviors and behaviors not relevant to target behaviors). The intermediate VP, with lower inhibition on the situation and goal layers, but with a more stringent level of behavioral inhibition (using the more stringent kwta algorithm), produced a VP whose behavior was more consistent with situational input. However, because this VP produced less target behavior than either the lowest inhibition VP or the high inhibition VP, it might be the pattern for a less productive or efficient individual.

The current model captures the behavioral inhibition that several researchers have argued is central to individual differences in temperament (Pickering & Gray, 1999). Because the behaviors activated by the goals in the approach and avoid systems compete with one another in the behavioral layer, the BIS system can act to inhibit the behavior of the individual. That is, if an avoidance goal is more highly activated than any approach goal, avoidance or withdrawal behaviors will typically "win" the competition with approach behaviors. Thus, the individual will be inhibited from approaching potentially rewarding situations.

One benefit of the VP simulation is that it allows us to tweak the system to find a pattern of behavior consistent with a given trait, and then ask, do the parameter system settings fit with our knowledge—or new empirical data? What new features of the model might be necessary?

Because the virtual personality approach is a useful way of testing out these dynamics and resultant behaviors and then using these results to examine the phenomena within real persons, it offers a personality approach or "laboratory" that we haven't had previously to study very complex dynamics and emergent outcomes.

The current model attempts to capture only 3 of the Big 5 dimensions. Another obvious way to expand the model is to examine whether and/or how we can capture the two other dimensions, Openness to Experience and Agreeableness. We are currently considering several possibilities (Read & Miller, 2002). For Agreeableness, one possibility is that agreeable individuals are simply higher on communal goals in both the BIS and the BAS system. Another focus is on the possibility that there are several specific motivational systems that might be responsible for these individual differences. For example, Panksepp (1998) has argued that humans and other mammals have a PLAY system that generates exploration and curiousity, and might provide part of the basis for Openness to Experience.

Shoda et al. (this issue) have also presented a neural network model of personality. Although both models are recurrent models, there are some very important differences between them. Both sets of researchers have focused heavily on understanding personality in terms of underlying goal based structures and have made that a central part of their model. Further, in both models, behavior is not the simple result of activation by situational cues, but is also the result of the interactions among a number of different mediating psychological components.

However, there are also some important differences. We believe that it is important that any model of personality be able to capture many of the important individual differences that have been encoded in our language and identified in the personality literature. Thus, the development of our model was constrained both by what is known about the nature of individual differences in behavior and by what is being discovered about the neurobiological bases of temperament and personality. In contrast, Shoda and his colleagues are not currently focused on trying to capture the structure of personality, but instead are examining other characteristics of personality dynamics, such as the general attractor dynamics of personality and the results of interactions between two individuals.

It is worth noting, that because both models are recurrent models, our model could be rather straightforwardly used to capture the same kinds of phenomena on which Shoda and colleagues focus. In contrast, because Shoda et al. (this issue) do not make any specific commitments to the internal structure of their model, they would be unable to capture the phenomena we have examined in our model.

A work in progress

VP is a "work in progress." Our goal was to start with a core set of dynamics, and to push the simple model as far as possible in terms of explaining personality outcomes. An iterative process of exploring the implications of a series of additional features is planned: What might each change "buy us" in understanding and predicting behavior? Some of our favorite next add-ons include learning, feedback control loops, emotional reactions, and better specification of chronic versus situational resources and their current state of availability or depletion. The current simulation does not model, for example, how the relative parameters at the sensory, goal, and motor levels initially emerge (e.g., the emergence of macro-level properties such as goals and norms). Adding learning to the model might inform our understanding of processes in personality development and change.

Connectionist and other dynamic approaches offer tremendous potential to build a much more holistic understanding of human psychology (Read, Vanman, & Miller, 1997). Ironically, simulating personality may better tell us what personality—in all its richness and complexity—really is.

References

Allport, G. W. (1962). The general and the unique in psychological science. *Journal of Personality*, *30*, 405–422.

Bakan, D. (1966). *The duality of human existence*. Chicago: Rand McNally.

Carver, C. S. (in press). Affect and the functional bases of behavior: On the dimensional structure of affective experience. *Personality and Social Psychology Review.*

Carver, C. S., & Scheier, M. F. (1998). *On the self-regulation of behavior.* Cambridge, England: Cambridge University Press.

Cervone, D. (1999). Bottom-up explanation in personality psychology: The case of cross-situational coherence. In D. Cervone & Y. Shoda (Eds.), *The coherence of personality: Social-cognitive bases of consistency, variability, and organization*. (pp. 303–341). New York: Guilford.

Cervone, D., & Shoda, Y. (Eds.). (1999). *The coherence of personality: Social-cognitive bases of consistency, variability, and organization*. New York: Guilford.

Chulef, A., Read, S. J., & Walsh, D. A. (2001). A hierarchical taxonomy of human goals. *Motivation and Emotion*, *25*, 191–232.

Church, A. T. (1994). Relating the Tellegen and the five-factor models of personality structure. *Journal of Personality and Social Psychology*, *67*, 898–909.

Clark, L. A., & Watson, D. (1999). Temperament: A new paradigm for trait psychology. In L. A. Pervin & O. P. John (Eds.). *Handbook of personality: Theory and research* (2nd ed., pp. 399–423). New York: Guilford.

Costa, P. T., Jr., & McCrae, R. R. (1994). Set like plaster? Evidence for the stability of adult personality. In T. F. Heatherton & J. L. Weinberger (Eds.), *Can Personality Change?* (pp. 21–40). Washington, DC: American Psychological Association.

Depue, R. A. (1996). A neurobiological framework for the structure of personality and emotion: Implications for personality disorders. In J. Clarkin & M. Lenzenweger (Eds.), *Major theories of personality disorders* (pp. 347–390). New York: Guilford.

Digman, J. M. (1997). Higher-order factors of the Big Five. *Journal of Personality and Social Psychology*, *73*, 1246–1256.

Ellis, R., & Humphreys, G. (Eds.) (1999). *Connectionist psychology: A text with readings*. Hove, England: Psychology Press.

Gray, J. A. (1987). *The psychology of fear and stress*. (2nd ed). New York: Cambridge University Press.

Guilford, J. P. (1959). *Personality.* New York: McGraw-Hill.

Higgins, E. T. (2001). Promotion and prevention experiences: Relating emotions to nonemotional motivational states. In J. P. Forgas (Ed.), *Handbook of affect and social cognition*. (pp. 186–211). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Jensen, A. R. (1958). Personality. *Annual Review of Psychology*, *9*, 295–322.

Marcus, H. R, & Kitayama, S. (1991). Culture and the self: Implications for cognition, emotion, and motivation. *Psychological Review*, *98*, 224–253.

Martindale, C. (1991). *Cognitive psychology: A neural-network approach*. Pacific Grove, CA: Brooks/Cole.

Miller, L. C., & Read, S. J. (1987). Why am I telling you this? Self-disclosure in a goal-based model of personality. In V. J. Derlega & J. Berg (Eds.), *Self-disclosure: Theory, research, and therapy.* New York: Plenum.

Miller, L. C., & Read, S. J. (1991).On the coherence of mental models of persons and relationships: A knowledge structure approach. In G. J. O. Fletcher & F. Fincham (Eds.), *Cognition in close relationships*. (pp. 69–99). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Mischel, W. (1999). Personality coherence and dispositions in a cognitive-affective personality system (CAPS) approach. In D. Cervone and Y. Shoda (Eds.), *The coherence of personality: Social–cognitive bases of consistency, variability, and organization* (pp. 37–60). New York: Guilford.

Nigg, J. T. (2000). On inhibition/disinhibition in developmental psychopathology: Views from cognitive and personality psychology and a working inhibition taxonomy. *Psychological Bulletin*, *126*, 220–246.

O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, MA: MIT Press.

Panksepp, J. (1998). *Affective neuroscience: The foundations of human and animal emotions*. New York: Oxford University Press.

Pervin, L. A. (1990). A brief history of modern personality theory. In L. A. Pervin (Ed.), *Handbook of personality theory and research* (pp. 3–18). New York: Guilford.

Pickering, A. D., & Gray, J. A. (1999). The neuroscience of personality. In L. A. Pervin & O. P. John (Eds.), *Handbook of personality: Theory and research* (2nd ed., pp. 277–299). New York: Guilford.

Read, S. J., Jones, D. K., & Miller, L. C. (1990). Traits as goal-based categories: The role of goals in the coherence of dispositional categories. *Journal of Personality and Social Psychology*, *58*, 1048–1061.

Read, S. J., & Miller, L. C. (1989). Inter-personalism: Toward a goal-based theory of persons in relationships. In L. Pervin (Ed.), *Goal concepts in personality and social psychology* (pp. 413–472). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Read, S. J., & Miller, L. C. (1993). Rapist or "regular guy": Explanatory coherence in the construction of mental models of others. *Personality and Social Psychology Bulletin*, *19*, 526–540.

Read, S. J., & Miller, L. C. (1998). On the dynamic construction of meaning: An interactive activation and competition model of social perception. In S. J. Read & L. C. Miller (Eds.), *Connectionist models of social reasoning and behavior.* Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Read, S. J., & Miller, L. C. (2002). *Neural network models of personality and social behavior.* Grant submitted to NIMH.

Read, S. J., Vanman, E. J., & Miller, L. C. (1997). Connectionism, parallel constraint satisfaction processes and Gestalt principles: (Re)Introducing cognitive dynamics to social psychology. *Personality and Social Psychology Review*, *1*, 26–53.

Sanford, N. (1963). Personality: Its place in psychology. In S. Koch (Ed.), *Psychology: A study of science* (Vol. 5, pp. 488–592). New York: McGraw-Hill.

Shoda, Y., LeeTiernan, S., & Mischel, W. (2002). Personality as a dynamical system: Emergence of stability and disinctiveness from intra- and interpersonal interactions. *Personality and Social Psychology Review, 6*, 316–325.

Shoda, Y., & Mischel, W. (1998). Personality as a stable cognitive–affective activation network: Characteristic patterns of behavior variation emerge from a stable personality structure. In S. J. Read & L. C. Miller (Eds.), *Connectionist models of social reasoning and social behavior* (pp. 175–208). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Spence, J. T. (1985). Achievement American style: The rewards and costs of individualism. *American Psychologist*, *40*, 1285–1295.

Tellegen, A. (1985). Structures of mood and personality and their relevance to assessing anxiety, with an emphasis on self-report. In A. H. Tuma & J. D. Maser (Eds.), *Anxiety and the anxiety disorders* (pp. 681–706). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Tellegen, A., & Waller, N. G. (1997). Exploring personality through test construction: Development of the multidimensional personality questionnaire. In S. R. Briggs, J. M. Cheek, & E. M. Donohue (Eds.), *Handbook of adult personality inventories*. New York: Plenum Press.

Triandis, H. C. (1995). *Individualism and collectivism*. Boulder, CO: Westview Press.

Watson, D., & Clark, L. A. (1993). Behavioral disinhibition versus constraint: A dispositional perspective. In D. M. Wegner & J. W. Pennebaker (Eds.), *Handbook of mental control* (pp. 506–527). Upper Saddle River, NJ: Prentice Hall.

Wiggins, J. S., & Trapnell, P. D. (1996). A dyadic-interactional perspective on the five-factor model. In J. S. Wiggins (Ed.), *The five-factor model of personality: Theoretical perspectives* (pp. 88–162). New York: Guilford.