An updated version of this manuscript has been published in D. Wood, S. J. Read, P. D. Harms, & A. Slaughter (Eds.), *Measuring and Modeling Persons and Situations* (pp. 617–649). Academic Press.
<https://doi.org/10.1016/B978-0-12-819200-9.00010-7>

Integrating Cybernetic Big Five Theory with the Free Energy Principle:

A new strategy for modeling personalities as complex systems

Adam Safron*

Indiana University

Colin G. DeYoung**

University of Minnesota

To appear in *Measuring and modeling persons and situations*, edited by Dustin Wood,
Stephen J. Read, Peter D. Harms, and Andrew J. Slaughter.

*The Kinsey Institute & Cognitive Science Program, 428 Lindley Hall, Indiana University,
Bloomington, IN 47405. Email: asafron@gmail.com

**Corresponding author: Department of Psychology, 75 East River Rd., Minneapolis, MN
55455. Email: cdeyoung@umn.edu

Abstract

Cybernetics is the study of goal-directed systems that self-regulate via feedback, a category that includes human beings. Cybernetic Big Five Theory (CB5T) attempts to explain personality in cybernetic terms, conceptualizing personality traits as manifestations of variation in parameters of the neural mechanisms that evolved to facilitate cybernetic control. The Free Energy Principle and Active Inference framework (FEP-AI) is an overarching approach for understanding how it is that complex systems manage to persist in a world governed by the second law of thermodynamics—the inevitable tendency toward entropy. Although these two cybernetic theories were developed independently, they overlap in their theoretical foundations and implications and are complementary in their approaches to understanding persons. FEP-AI contributes a potentially valuable formal modeling framework for CB5T, while CB5T provides detail about the science and structure of personality. In this chapter we explore how CB5T and FEP-AI may begin to be integrated into a unified approach to modeling persons.

Keywords

CB5T, cybernetics, Free Energy Principle, Active Inference, generative models

A Cybernetic Approach to Personality

“SOCRATES: Or again, in a ship, if a man having the power to do what he likes, has no intelligence or skill in navigation [*aretēs kybernetikēs*], do you see what will happen to him and to his fellow-sailors?”

—Plato, *Alcibiades I*

Cybernetics, from the Greek *kybernetikēs*, meaning “steersman,” is the study of principles governing goal-directed systems that self-regulate via feedback. Since its invention by Wiener (1948), cybernetics has revolutionized multiple disciplines, forming a foundation for many aspects of cognitive science, computer science, and robotics and leading to a variety of additional neologisms, like “cyberspace,” “cyborg,” and cybersecurity.” In its modern forms, it is often known as “control theory.” One of Wiener’s key insights was that similar principles must be involved in the regulation of artificial control systems and of organisms, given that the latter need to pursue various goals in order to survive and reproduce. The relevance of cybernetics for the study of human psychology was quickly recognized and it has been highly influential (Austin & Vancouver, 1996; Carver & Scheier, 1998; Miller, Galanter, & Pribram, 1960; Powers, 1973). More recently, cybernetic principles have also begun to be applied in the study of personality—that is, individual differences in psychological variables (Carver, Sutton, & Scheier, 2000; DeYoung & Weisberg, 2019; Van Egeren, 2009).

When cybernetic principles are applied to personality, the focus is typically on evolutionarily conserved control parameters, which are persistent, but can be tuned in various ways, such that they can account for stability in personality over time, as well as accounting for individual differences. Crucially, however, those persistent parameters are conceived as part of a dynamic system that is constantly changing and reacting to different situations; cybernetics thus

helps to bridge the gap between the study of the dynamics of an individual person's personality from moment to moment and the consistencies in behavior that differentiate people from each other (DeYoung & Weisberg, 2019).

One of the most thorough attempts to use cybernetics to guide the construction of personality theory is *Cybernetic Big Five Theory* (CB5T; DeYoung, 2015). CB5T describes all persistent psychological individual differences as personality and asserts that all elements of personality can be categorized as either personality traits or characteristic adaptations.

Personality traits are defined as probabilistic descriptions of relatively stable patterns of behavior, motivation, emotion, and cognition, in response to classes of stimuli that have been present in human environments over evolutionary time. These patterns are dispositional, in that they reflect the ways that people tend to respond to the relevant stimuli, and they are what is assessed by standard questionnaire measures of personality traits. Traits, as dimensions of variation between people, reflect cross-person variation in the parameters of evolved cybernetic mechanisms present in every intact human brain. Inasmuch as these parameters have some stability within the individual, this tends to produce trait-like dispositions.

Characteristic adaptations are relatively stable goals, interpretations, and strategies, specified in relation to an individual's particular life circumstances. They are the brain's persistent but updatable memory contents, learned in response to life experiences, also describable as habits. Notably from a cybernetic perspective, goals, interpretations, and strategies correspond to the three necessary elements of any cybernetic system: a representation of a goal (or goals) physically instantiated within the system, a representation of the current state that can be compared to the goal state via feedback, and a set of operators that can be engaged to attempt to transform the current state into the goal state (DeYoung & Weisberg, 2019).

Figure 1 depicts the general causal model of CB5T. Both genetic and environmental forces influence the relatively stable parameters of cybernetic mechanisms that produce the dispositional patterns described as personality traits. In turn, personality traits shape people's characteristic adaptations, because persistent patterns of behavior influence what goals, strategies, and interpretations are discovered, adopted, and remembered. The traits reflect variations in the underlying mechanisms that allow people's characteristic adaptations to be learned and enacted, inasmuch as personality traits stem from variation in the biological systems that evolved to allow people to learn about the world and to enact their goals. We discuss examples in relation to many specific traits in our penultimate section.

Figure 1. Causal processes in the functioning of personality within the individual. Both genes and the environment directly influence the cybernetic mechanisms underlying personality traits, which are patterns of behavior, motivation, emotion, and cognition. All genetic influences on characteristic adaptations are funneled through traits, but the environment can influence characteristic adaptations independently of the influence of traits. Circular arrows indicate that cybernetic parameters can influence each other, as can characteristic adaptations and other life outcomes, such as physical health. (Adapted with permission from DeYoung, 2015.)

Unlike personality theories that describe characteristic adaptations merely as manifestations of traits, however (e.g., McCrae & Costa, 2008), CB5T recognizes traits and characteristic adaptations as fundamentally distinct entities, such that one may have characteristic adaptations that are incongruent with one's traits because they represent adaptations to some particularly important aspect of one's situation (indicated by the causal effect of Environment directly on Characteristic Adaptations in Figure 1), where trait-typical behavior would be inappropriate. For example, an introvert who works in sales may habitually adopt an outgoing, talkative manner, but only in that context. CB5T also recognizes that characteristic adaptations may influence traits such that they may feed back to re-tune basic cybernetic parameters. For example, if an introvert became good at acting extraverted for a job in sales, this might feed back to re-tune the cybernetic parameters (relating to reward sensitivity) that underlie Extraversion, leading to a general increase in Extraversion across many other situations (DeYoung, 2015).

Because the Big Five personality traits are major dimensions of covariation among many more specific traits (regardless of whether these are assessed via adjectives from the lexicon or by phrase-based items designed to measure other traits; John, Naumman, & Soto, 2008; Markon, Krueger, & Watson, 2005; Waller, DeYoung, & Bouchard, 2016), CB5T focuses on the Big Five as broad traits that any comprehensive theory of personality will need to explain. However, CB5T is not merely a theory of the Big Five, but also of the causal processes contributing to personality as a whole, as well as the ways in which dysfunction in these processes can potentially contribute to psychopathology, conceived as cybernetic dysfunction (DeYoung & Krueger, 2018a, 2018b).

CB5T defines psychopathology as “persistent failure to move toward one’s psychological goals due to failure to generate effective new goals, interpretations, or strategies when existing ones prove unsuccessful.” Merely being temporarily blocked from movement toward one’s goals does not constitute cybernetic dysfunction because cybernetic systems are self-correcting by nature. Only when the system is both thrown off course and then subsequently unable to engage operations to bring itself back on course can it be reasonably described as dysfunctional. From CB5T’s perspective, psychopathology necessarily involves personality because it involves a failure both of current characteristic adaptations and then subsequently of the processes that should generate effective new characteristic adaptations. Although extreme personality traits are not necessary for psychopathology in CB5T, they are typically risk factors for psychopathology because unusually extreme values of functional parameters for components of a cybernetic system render that system more likely to malfunction.

CB5T additionally attempts to identify brain systems and neurobiological parameters that are likely bases of dimensions along which personalities vary (Allen & DeYoung, 2017). CB5T’s mapping of traits to brain function has been essentially piecemeal, considering the psychological functions associated with each trait, and then attempting to develop hypotheses regarding their neural instantiation by identifying brain systems that have been shown to carry out these functions. Thus far, CB5T has made little contact with attempts outside of personality and social psychology to use concepts related to control theory and cybernetics to understand the mind and brain. In this chapter, we attempt to begin a synthesis of CB5T with cybernetic perspectives from cognitive science and neuroscience, in the hope that integrating CB5T with other cybernetic approaches will facilitate progress toward formally modeling personality dynamics.

The Free Energy Principle and CB5T

A paradigm shift is under way in cognitive science toward a predictive processing approach. Predictive processing is inherently cybernetic (Clark, 2013; Seth, 2014), as it assumes that the basic functional architecture of the brain, throughout its whole hierarchy of organization, is to compare observations against predictions, and to attempt to operate in such a manner that minimizes prediction errors. An especially thorough and broad approach to the predictive coding paradigm has been developed by Friston (2010, 2019), embodied in the *Free Energy Principle (FEP)*, a comprehensive formal framework for understanding mind and life as dynamical systems. The scope of the FEP is extremely broad, ranging from individual cells to societies (Friston, 2019; Ramstead et al., 2018; Veissière et al., 2019). The present chapter provides a summary of the basic principles of the FEP and discusses how they can be integrated with CB5T, thereby offering a more formal approach to modeling the cybernetic mechanisms underlying personality.

All living systems are unusual in their ability to persist in a world governed by the second law of thermodynamics, which asserts that any system not already in thermodynamic equilibrium will tend to exhibit increasing degrees of disorder (i.e., entropy). That is, without some kind of continual supply of energy to steer the system in a more organized direction, probability dictates that the elements of any system will become more randomly ordered, such that its particular composition eventually becomes maximally diffuse and disorganized. In this way, the second law does not describe what must be observed in all instances without exception, but rather is a statistical necessity—a trend that holds with such reliability that countervailing processes must be identifiable for any systems that manage to persist (Brillouin, 1951; Tooby et al., 2003).

The FEP describes, at an abstract level, the processes by which adaptive systems are able to steer away from the maximally probable outcome of maximal disorder (Brillouin, 1951; Friston, 2013; Kirchhoff et al., 2018; Schrodinger, 1944). In cybernetics, the requirements for such a governing process are expressed as the *good regulator theorem* and the *law of requisite variety*: the former states that any effective controller must contain a model of the variables it is controlling, and the latter states that these regulating models require sufficient complexity to represent the variety of states of the controlled variables that are likely to be encountered (Boyd et al., 2017; Conant & Ashby, 1970; Heylighen, 1992). As highly complex cybernetic systems with many goals, persons must possess sufficiently accurate and detailed models of the world and themselves, such that they can navigate toward areas of state space in which they are able to continue to exist as particular persons. According to the FEP, persisting complex adaptive systems necessarily minimize prediction-errors with respect to these models by which they navigate the world.

CB5T begins with the realization that minds and the brains that instantiate them are control hierarchies for goal-directed, self-regulating systems. These controllers evolve and develop for the overarching purpose of minimizing uncertainty regarding goal-attainment—in other words, minimizing cybernetic entropy (DeYoung, 2013; DeYoung & Krueger, 2018b; Hirsh, Mar, & Peterson, 2012). This perspective is consistent with the foundations of the FEP, wherein persisting dynamical systems are governed by a singular objective of prediction-error minimization. The FEP provides a first-principles derivation of the normative preconditions for successful goal-seeking, and CB5T provides detailed models of persons, grounded in the same cybernetic foundation.

The FEP and CB5T are founded on the same fundamental principles for what systems must do if they are to achieve their goals, and both frameworks constitute dynamical perspectives, with CB5T describing traits as “equivalent to persistent attractor states of the cybernetic system; they indicate states toward which the person will tend to gravitate but do not preclude that person from being in other states” (DeYoung, 2015). This formulation is consistent with other descriptions of traits as “density distributions of states,” reflecting regularities in the diverse states of the person we might expect to find if we were to measure persons at different time points (e.g., via experience sampling) (Fleeson, 2001; Fleeson & Gallagher, 2009; Fleeson & Law, 2015).

This characterization is also consistent with the FEP, which models systems as probabilistic densities constituted by trajectories through state space. To the degree that systems persist, they possess attracting sets that define them as particular phase space densities with varying probabilities of occurrence, wherein attractor dynamics produce the very generative processes out of which they are constituted. Although this idea may seem complicated, the essence of it is almost (but not quite) tautologically self-evident: systems are likely to exist to the degree they can perpetuate their own forms, including with respect to generating the network of relations that define them as particular systems.

The FEP describes these attractor configurations and ensuing trajectories as *generative models*, where that which is generated is the probabilistic densities of states that define the existence of particular systems. This is a place where the purview of the FEP is considerably broader than CB5T. CB5T focuses on complex adaptive systems (specifically humans) that are governed by psychological models—organized systems of characteristic adaptations that comprise representations of the current state of the world, the desired state, and strategies to

move from the former to the latter. The FEP, in contrast, attempts to model *all* systems in terms of the processes by which they persist (Friston, 2019). In this way, where CB5T is focused on the human cybernetic system, the FEP attempts to characterize all dynamical systems.

The processes enacted by generative models in human beings can take many forms, ranging from unconscious habits to emotionally charged reactive dispositions, to declarative knowledge, and even self-organization via autobiographical narratives (Damasio, 2012; Hirsh et al., 2013). To the extent that persons have identifiable traits and characteristic adaptations (i.e., personalities), these would represent relatively stable parameter values for generative models governing dynamics. This parameter stability may be genetically specified, epigenetically canalized (Waddington, 1942), or could emerge from relatively stable self-reinforcing attractor states whose causal structure involves both innate and learned factors.

Within the FEP, different personality configurations correspond to processes by which persons attempt to achieve their goals via predictive modeling. The FEP quantifies the ability of systems to minimize prediction error according to an information theoretic quantity of *variational (or approximate) free energy*. Derived from statistical physics, this objective function is optimized by minimizing discrepancies between probabilistic beliefs and observations (i.e., *prediction-errors*, or “*surprisal*”), penalized by model complexity (i.e., Occam’s Razor) (Dayan et al., 1995). This optimization occurs via stochastic gradient descent, in which models are iteratively updated according to whichever combinations of parameter values result in the maximal reduction of free energy (i.e., cumulative prediction error).

This way of understanding systems in terms of probabilistic generative modeling has close correspondences with “Bayesian brain” approaches in which all neural functioning is understood as probabilistic inference (Friston, 2010; Seth, 2014). The FEP is specifically

Bayesian in that the models by which complex adaptive systems are governed are updated by combining all available data, whose contributions to updating are weighted by the estimated reliability of each source of information. This mode of inference allows for the ongoing refinement of beliefs through continual observations and iterative updating of hypotheses. It also represents the cutting edge of artificial intelligence in facilitating “self-supervised” learning and data augmentation, wherein systems train themselves by generating predictions, which they can then compare against empirical observations (Hohwy, 2020; Tschantz et al., 2020).

Predictive Processing and Personality Neuroscience

“Every prediction is an operation on the past.”

–Norbert Wiener

“The whole function of the brain is summed up in: error-correction.”

–W. Ross Ashby

Whereas the previous discussion has focused on the functional or computational level of analysis, the FEP has also generated fruitful theorizing and experiments with respect to biological mechanisms. Neural networks can be understood as generative models in a relatively straightforward sense: the directed structure of neurons and their organization into networks generate patterns of effective connectivity (Friston, 1994; Park et al., 2018; Park & Friston, 2013), where directed weighted connections and flows of mechanistic influence are physical instantiations of probabilistic relations. From this perspective, nervous systems can be viewed as entailing *models* to the extent that internal patterns of effective connectivity express information about system and world.

With relevance to the FEP, it has been suggested that all cortical processing may potentially be explainable in terms of a common cortical algorithm of *hierarchical predictive*

coding (HPC)—whose importance for neuroscience may be difficult to overstate (Hawkins & Blakeslee, 2004; Mumford, 1991; Rao & Ballard, 1999). In HPC, neuronal processes constitute hierarchically organized generative models that attempt to infer the most likely (hierarchically organized) world states that could have caused actual sensory observations (Clark, 2013; Friston & Kiebel, 2009). Bottom-up sensory information is simultaneously predicted across all levels of cortical hierarchies by sending predictions downwards in anticipation of sensory observations. These “predictions” correspond to Bayesian statistical expectations (or prior probabilities) in HPC, rather than to folk psychological beliefs. In this scheme for efficiently coding information, only discrepancies (i.e., prediction-errors) are passed upwards toward higher levels, causing modifications that update beliefs into posterior expectations, which then become new (empirical) priors (or predictions) to be passed downwards.

HPC provides an algorithmic and implementational process for Bayesian inference by iteratively updating beliefs based on a weighted combination of priors and observations. In computational terms, the relative weighing of expectations and evidence is determined according to a parameter known as “precision,” which corresponds to the inverse variance of probability distributions; the greater the precision, the less variable the estimate. This weighted combination of probabilities renders HPC Bayesian by combining sources of information according to their estimated relative reliabilities. Within HPC, precision-weighting determines the extent to which prediction-errors are likely to percolate to higher-levels of the cortical hierarchy and thereby drive updating and model selection with respect to consciously-accessible high-level beliefs (Safron, 2020). This precision-weighting has been associated with attention, or salience, and is one of the most important factors responsible for shaping inference and learning in Bayesian brains (Parr & Friston, 2017b). To summarize, in HPC, each level of neural organization models

the level below it, extending down to sensory and effector systems, with all of these models being integrated when they are bound into larger generative models (e.g., brains and organisms).

Evidence continues to accumulate for predictive coding as a fundamental principle of cortical functioning (Bastos et al., 2012; Bastos et al., 2015; Bergmann et al., 2019; Chao et al., 2018; Sedley et al., 2016). If, as HPC suggests, it is indeed the case that all cortical processing is governed by a common algorithm, this will facilitate interpreting specific findings and generating new hypotheses in personality neuroscience, much as in cognitive neuroscience. Further, various psychological individual differences could potentially be explainable in terms of the relative gain on predictions or prediction-errors in hierarchical message passing.

Along these lines, some personality characteristics associated with the autism spectrum may be explainable in terms of excessive prediction-errors (i.e., too much gain on sensory observations) (Lawson et al., 2014; Markram & Markram, 2010), and key features of psychosis may derive from anomalous predictions (Friston et al., 2016), with hallucinations and delusions both potentially understood in terms of overly strong predictions. In the Big Five, risk for psychosis is linked to Openness to Experience (DeYoung & Krueger, 2018a), and personality traits related to Openness to Experience could potentially be explainable in terms of the ability to relax prior perceptual expectations, with potential sources of variation centering on deeper portions of the brain's generative models such as posterior cingulate cortices (Carhart-Harris & Friston, 2019; Carhart-Harris, 2018; Carhart-Harris et al., 2014; Erritzoe et al., 2019; Preller et al., 2019; Smigielski et al., 2019). Brain regions such as the cingulate and association cortices can be considered “deep” with respect to generative modeling in that they appear to host high-level priors for brains as belief networks. Beliefs from these hierarchically higher (or deeper) areas would potentially have the maximal impact on the overall functioning of the network, and

may have strong correspondences with broad personality traits (Hassabis et al., 2014). One of the core goals of CB5T is to develop a theory of personality that is compatible with and informed by neuroscientific knowledge (Allen & DeYoung, 2017), and using the FEP to develop personality models will help them to be consistent with the current state of knowledge regarding HPC.

Active Inference (AI)

“Objects are always imagined as being present in the field of vision as would have to be there in order to produce the same impression on the nervous mechanism.”

—Hermann Ludwig Ferdinand von Helmholtz (1910)

“Each movement we make by which we alter the appearance of objects should be thought of as an experiment designed to test whether we have understood correctly the invariant relations of the phenomena before us, that is, their existence in definite spatial relations.”

—Hermann Ludwig Ferdinand von Helmholtz (1878)

Work within the FEP paradigm has yielded a normative model of behavior in *Active Inference* (Friston et al., 2017). The Active Inference framework describes processes by which free energy is minimized, and provides close ties to other frameworks such as expected utility theory and risk-sensitive control. Further, advances in deep reinforcement learning appear to be converging on the kinds of solutions that are predicted to be necessary for (bounded) optimal decision-making and learning in *the FEP and Active Inference (FEP-AI) framework* (Ha & Schmidhuber, 2018; Hafner et al., 2020; Tschantz et al., 2019, 2019; Wang et al., 2018).

The notion of active inference rests on the insight that perception takes place within the context of adaptively shaping actions, which alter patterns of likely perceptions. Rather than being a simple result of passive sensations, perception is an active process of foraging for information and resolving model uncertainty (Parr & Friston, 2017a; Pezzulo & Friston, 2019;

Wade & Kidd, 2019), often driven by discrete actions as a kind of hypothesis testing (Gopnik et al., 2017; Parr & Friston, 2018). Both perception and action are understood as kinds of inferences in the FEP-AI framework, in that they both represent means by which systems can minimize prediction-error with respect to the models by which they navigate the world. Importantly, goals are just as much components of these models as are representations of the current state of the world, and minimizing prediction error with respect to goals means moving toward them in state space.

One way a system can reduce prediction-error is by updating internal models, thus changing predictions; this is perceptual inference. Another way a system can reduce prediction-error is by updating the world through action, thus making its predictions more accurate by changing likely perceptions; this is active inference. The degree to which system dynamics are governed by these two strategies—of perceptual or active inference—is determined according to whichever combination is expected to minimize overall free energy (i.e., cumulative precision-weighted prediction error). The development of the Active Inference framework has made connections between FEP and cybernetics more obvious, as it highlights that organisms are motivated not only to predict their perceptions accurately, but also to influence states of the world in order to ensure that perceptions correspond to goal attainment (Powers, 1973).

Within the FEP-AI framework, all cybernetic systems necessarily minimize free energy for their generative models. However, in order to effectively achieve this objective, adaptive goal-seeking systems such as organisms select actions anticipated to result in free energy-minimizing consequences in the future. Under this regime of *expected free energy*, model accuracy with respect to the realization of preferred outcomes is equivalent to expected utility, or opportunities for realizing the so-called *extrinsic value* of preference satisfaction (i.e., achieving

goals). In contrast, model complexity contributes to the ambiguity and risk associated with pursuing courses of action under varying states of uncertainty, which is higher given more complex modeling due to a broader range of states being possible. All sources of uncertainty are opportunities for realizing the *intrinsic value* of refining models via foraging for information and learning more about what was unknown.

Optimizing for extrinsic value involves minimizing discrepancies between preferred system-world configurations and observations, which entails pragmatically exploiting particular *policies*. Policies are sets of state-action mappings for goal realization, broadly construed to include the covert behavior of cognition. In CB5T's terms, policies are equivalent to *strategies*. Optimizing for intrinsic value involves model refinement through seeking out sources of uncertainty as opportunities for maximizing information gain, thereby allowing for epistemic exploration of potentially adaptive actions. These two sources of value relate to tradeoffs between exploitation (pursuing extrinsic value) and exploration (pursuing intrinsic value), which are navigated by selecting policies based on whatever combination of actions is estimated to most effectively minimize overall expected free energy (Kaplan & Friston, 2018).

If such actions occur in the context of a novel task environment about which little is known, then policy selection in FEP-AI will tend initially to involve more exploratory behavior in which information gain is optimized, followed by a shift to more exploitative behavior as the task structure becomes sufficiently clear to afford informed actions. However, if actions fail to be as successful as anticipated, then this will tend to result in shifting back to exploratory behavior until a better grasp of the situation can be acquired (Bruineberg & Rietveld, 2014; Kiverstein et al., 2019). In this way, given well-calibrated prior expectations, agents governed by FEP-AI will

tend to exhibit flexible levels of curiosity, thus balancing extrinsic and intrinsic value as they engage in goal-oriented behavior.

The existence of exploration and exploitation as dual strategies for minimizing expected free energy means that cybernetic systems governed by FEP-AI may decide to forego pursuing a goal in favor of model refinement, depending on whatever is anticipated to lead to larger reductions in prediction error (Parr & Friston, 2017b). This may be a surprising implication of FEP-AI—as evolutionary fitness depends more directly on achieving goals than on enhancing the accuracy of perceptual maps—but refining internal models can be preferred over achieving valued goals only because this tradeoff has served the more fundamental goal of promoting continued existence and effective goal pursuit over evolutionary time. In other words, having more accurate models of the world is important from an evolutionary perspective only because it increases the success of achieving goals that increase fitness. FEP-AI tries to address this prioritization by specifying “evolutionary priors” (Zador, 2019), implicit posterior beliefs based on fitness consequences for previous generations. This continuity of prediction error minimization for individual agents and evolutionary processes—where natural selection may be understood as a kind of learning, and vice versa (Campbell, 2016; Ramstead et al., 2018)—helps to ensure that all predictions are ultimately chained to stable goal pursuit, leading evolved systems to minimize free energy with respect to maintaining the preconditions for their existence.

The notion of action as a kind of inferential dynamic is crucial for understanding how cybernetic systems are capable of minimizing free energy. It also suggests a potentially major axis of personality variation across individuals. In any given situation, prediction error will be minimized via some combination of updating internal versus external states of the world (Friston

et al., 2014). Whereas the precise degree to which either action or perception drives model-evolution is determined by whichever combination is expected to result in the best fit for the models by which behavior is governed, these expectations will differ across different kinds of systems, including persons. The implications of this source of variation for personality are potentially extensive. Theoretically, individual differences in tuning the differential gain on action or perception could help to explain certain differences in personality, such as differential risk for externalizing vs. internalizing psychopathology (DeYoung & Krueger, 2018a). An agent that tends to minimize prediction-error by updating world states (rather than its beliefs) may be more likely to “act out” (corresponding to externalizing problems), relative to an agent that tends to optimize via model refinement (potentially corresponding to the behavioral inhibition and rumination associated with internalizing problems).

Interpretations of biological processes in terms of parameter settings for FEP-AI may help provide convergent support for CB5T and link its descriptions to computational models. For example, differing levels of dopaminergic function appear to have major impacts on personality with respect to Extraversion and probably also Openness/Intellect (DeYoung, 2013; Wacker & Smillie, 2015). FEP-AI associates tonic dopaminergic function with precision (as an inverse temperature parameter) over policies, indicating certainty (subjectively, confidence) and more deterministic action selection. Phasic dopamine, however, indicates changing estimates with respect to expected free energy and updating of likelihoods for selecting different policies for enaction, as in reward prediction errors (Adams et al., 2020; FitzGerald et al., 2015; Friston et al., 2012, 2014), including both overt behavior and cognitive operations (Parr et al., 2019). That is, nervous systems tuned to exhibit stronger dopaminergic signaling may more readily deploy mental acts such as discrete attentional shifts and simulated plans, potentially contributing to

more exploratory and flexible cognitive styles. Although the conceptual framework is different, FEP-AI's account of dopaminergic function has strong correspondences with CB5T's interpretation of dopamine as a “neuromodulator of exploration” and contributor to the broad personality trait Plasticity—the shared variance of Extraversion and Openness/Intellect—which is posited to reflect the exploratory tendency to generate new goals, interpretations, and strategies (DeYoung, 2013).

As illustrated with this analysis of dopamine signaling, CB5T and FEP-AI converge on highly similar interpretations of biological phenomena and their psychological implications, albeit from quite different perspectives. With the addition of Active Inference as a process theory for the FEP, this framework can be seen to align even more clearly with CB5T's perspective on personality as reflecting cybernetic control parameters and their functional consequences. In what follows, we discuss how the principles of FEP-AI can be used to create formal models of personality that are consistent with CB5T.

Generative Models and Personality Modeling

A primary research strategy employed within FEP-AI is as follows: a) starting from formal foundations, specify a generative model for a system, b) simulate that system in a task environment, and c) compare *in vitro* and *in vivo* data for the modeled system. If the system has agentic properties, then these simulations typically take the form of Markov decision processes (MDPs), wherein agents (as generative models) make sequential decisions (i.e., policy selections), often under conditions of incomplete information (so constituting partially observable MDPs, or POMDPs). In this section we will review some of the major components of generative modeling of agents within an MDP framework. For readers who would like a more detailed handling of generative models of agents, some excellent tutorials have recently become

available online for public viewing (*"Deeply felt affect: Understanding emotions through deep active inference"* — Casper Hesp, 2020; *Maxwell Ramstead — A Tutorial on Active Inference*, 2020; *"Modeling Task Behavior with Active Inference"* — Ryan Smith, 2020).

Generative modeling in FEP-AI can be complex (Hesp, Smith, et al., 2019; Hesp, Steenbeek, et al., 2019), but the basic setup involves specifying agents (and their situations) as systems of interacting variables and parameters. These variables take the form of mutually referencing matrices, whose entries specify probabilistic mappings between different portions of the generative model. As previously described, a single objective function of expected free energy is used to determine which policies are selected, according to whichever actions are anticipated to minimize overall prediction-error. An example FEP-AI generative model (Figure 2) may include the following matrices, where each mapping is conditional upon the probabilistic mappings from preceding matrices:

1. A-matrix: Likelihood mapping from latent world states to sensory observations, or how sensations are interpreted via perception and imagination. From the agent's perspective, this would correspond to answering the question, "What is likely to have caused my sensations?"
2. B-matrix: Transition probabilities, or how dynamics are expected to evolve from one moment to the next. From the agent's perspective, this would correspond to answering the question, "How do I believe things are likely to change through time?"
3. C-matrix: Prior beliefs over policies and outcomes, or which states an agent predicts itself occupying. Via mechanisms of prediction-error minimization, these prior expectations constitute the attracting states of a system, and function as preferences (Samuelson, 1948). From the agent's perspective, this would correspond to answering the

question, “What do I want to happen?” (and thereby doing/enacting it through active inference in the ways described above). From a CB5T perspective, these preferences are an agent’s high-level goals.

4. D-matrix: An agent’s belief regarding its initial position within the decision environment. From the agent’s perspective, this would correspond to answering the question, “Where am I likely to be in this kind of situation?”
5. E-matrix: Baseline priors over policies, governing policy selection in a habitual (or heuristic) fashion. From the agent’s perspective, this would correspond to answering the question, “What am I likely to be doing right now?” (and thereby doing/enacting through active inference, but with actions deployed in a fashion that ignores uncertainty with respect to the likelihood of goal attainment in a particular situation). From a CB5T perspective, these habit-priors would constitute a major class of characteristic adaptations.

Figure 2. Example of a generative model in an FEP-AI Markov decision process (MDP).

The direction of arrows indicates the data generation process via active inference, which in this case takes the form of selecting policies (π) as sequences of actions which result in influencing various world states. Selected policies specifically take the form of specifying sequences of state transitions, chosen according to whichever actions are expected to minimize overall prediction-error, or free energy (G). These states are hidden, or latent, and must be probabilistically inferred on the basis of sensory observations, where this perceptual inference flows in an opposite direction to that indicated by the arrows corresponding to the processes generating action sequences. The action selection resulting from this generative process constitutes a kind of inference, in that it reflects expectations as to what is likely to minimize the divergence between predictions and observations. The precision with which agents probabilistically select among actions is reflected by “inverse temperature” parameters (γ and β) to specify thresholds for policy deployment, so influencing exploration/exploitation tradeoffs.

The generative model defined by these joint matrices can then be inverted using linear algebraic operations (implemented by scripts pre-packaged with SPM12) (Friston et al., 1994; Penny et al., 2011), in order to generate various patterns of behavior. These patterns can range from discrete choices in either simple or complex tasks (Constant et al., 2019) to patterns of eye movements during informational foraging (Friston, Lin, et al., 2017) to likely patterns of neuronal responses expected to accompany psychological processes ranging from decision making to emotional categorization (Smith et al., 2019). This simulated in vitro data can then be compared against empirical data in vivo, with different hypothetical generative models chosen

and refined via Bayesian model selection and reduction (Friston et al., 2019). Generative models can be elaborated as much as desired by adding additional factors to further parameterize A, B, C, D or E matrices in any combination. Graphically, this would correspond to specific directed connections to generative model components, entailing specific conditional probabilistic relationships.

The basic setup just described as A-E matrices forms a generic model of an agent. Additional variables and parameters could be included to model systematic variation between agents corresponding to personality traits, understood as high-level expectations regarding the attracting states of the system. Here we will refer to these elaborations of generative models as P-structures, whose elements would include specifications of which parameters of A, B, C, D, and E matrices should be varied, and in which ways, to produce agents varying in a particular trait. A P-structure could specifically take the form of a graphical model, entailing a system of structured conditional probability distributions (Koller & Friedman, 2009). For example, variation in Openness to Experience could correspond to a P-structure parameterizing an A matrix in terms of its perceptual components (e.g. sensitivity to novel information), B and D matrices in terms of beliefs about the world (e.g., likelihood of detecting meaningful patterns), and C and E matrices in terms of the kinds of behaviors that tend to be enacted by more open people (e.g. selecting policies that realize epistemic value).

To provide another example, degrees of Neuroticism could be specified as a separate P-structure, which could parameterize A (e.g. sensitivity to interoceptive prediction-errors), B (e.g. expecting volatile state transitions), C (e.g. preferring risk-minimizing policies), D (e.g. believing that one tends to be in high-risk circumstances), and E matrices (e.g. habitual tendencies for large responses to error signals). Different P-structures will produce different

patterns of overt and covert behavior when applied as parameterizations of generative models, and so can be evaluated based on their ability to fit empirical data from actual people or hypothesized data based on assumptions about what variation in a given trait should look like in an artificial agent.

Personality modelers interested in using these methods can find a diverse assortment of demonstrations (and commented scripts) by installing SPM12 and running the Matlab command, DEM. These scripts can be downloaded from supplementary materials accompanying most FEP-AI publications, which cover a wide variety of topics of interest for understanding different aspects of persons. For illustrative purposes, Figure 3 reprints example figures generated by one of these pre-packaged demos (DEM_demo_MDP_rule), which explores curiosity in terms of patterns of maximizing expected value and foraging for information during a structure-learning task (for more background on this simulation and associated hypotheses, see online lecture by Friston (2016), “Active inference and artificial curiosity”). This kind of generative modeling can simulate a wide variety of phenomena, ranging from choice behavior to predicted neural responses to parameter estimates for agents (as generative models). Further, this rich diversity of simulated data, via Bayesian model inversion, allows for FEP-AI based generative modeling to make highly specific predictions along multiple dimensions, thereby affording multiple opportunities for hypothesis testing and potential falsification of particular models.

This demo and preexisting work in FEP-AI do not include P-structures that would allow one to systematically vary the parameters of the matrices to specifically produce agents (and populations of agents) varying in particular personality traits. With generative modeling in SPM, Bayesian model inversion scripts allow multiple simulations to be run with stochastically varying parameter combinations, so allowing for probabilistic estimation of variables via distributions of

simulated behavior and outcomes, including with respect to personality characteristics such as curiosity. However, this is somewhat different than what we have in mind with our P-structure proposal, which would require identifying the parameters of the model that one believed were responsible for variation in curiosity, and then introducing functions specifying these associations as empirical priors. By observing the behavior of the agents as a function of these variations in parameters, one could test whether the parameters identified were plausible contributors to variations in curiosity. This provides a general way forward for investigating how personality differences might be instantiated in a system operating under FEP-AI.

Figure 3. Illustrative figure accompanying a generative modeling demo pre-packaged with SPM12. This particular simulation involved modeling curiosity with respect to a task with unknown rules, where agents attempt to maximize valued outcomes while also needing to discover the latent task structure. Generated patterns of data include inferred rules (top left), patterns of task choices (middle left), and neuronal activity anticipated to occur based on mechanistic process theories associated with FEP-AI (bottom left, all right panels).

Personality Traits from the Perspectives of CB5T and FEP-AI

The normative principles of FEP-AI provide a basis for the formal modeling of persons (and situations) and the regularities in behavior that constitute personality. They also provide a set of fundamental axes along which individuals can be differentiated as cybernetic systems. Within CB5T, as noted above, personality traits are understood as parameter settings for evolved cybernetic control mechanisms, whereas characteristic adaptations are learned goals, interpretations, and strategies defined in relation to particular life circumstances. Within FEP-AI, personality traits would correspond descriptively to variations in the states in which the system is likely to be found, caused by variation in the persistent parameters of the matrices and functions of the generative model, and characteristic adaptations would correspond to tendencies for selecting from hierarchically organized policies for particular actions in the system's current repertoire (Table 1). Importantly, characteristic adaptations would correspond not only to the policies themselves, but to all updatable memory contents capable of contributing to policy selection, including typical interpretations (i.e., likelihood mappings), strategies (i.e., policies), and goals (i.e., consistently predicted states of the situation, understood as entailing preferences).

The goals in question would be characteristic adaptations as long as they are learned from the experience of the agent, rather than being built into the system from the beginning (in which case they would constitute personality traits, according to CB5T).

Figure 4. The personality trait hierarchy. Note that the structure indicated here is the between-person covariance structure of traits in the population, not a within-person causal structure. Nonetheless, CB5T assumes that traits aggregated at a higher level are likely to share some of their mechanisms. Top level = metatraits, second level = Big Five, third level = aspects, fourth level = facets. Minus sign indicates that Neuroticism is negatively associated with Stability. The diagram is a simplification, in that some associations that exist among traits at the aspect level (e.g., a negative correlation between Politeness and Assertiveness), are not deducible from the paths depicted.

Numerous points of intersection can be identified between the FEP-AI framework and the account of traits provided by CB5T. Most of the personality traits explicitly treated by CB5T are

depicted in Figure 4. Below we will briefly describe how a number of these traits might be accounted for in terms of these two complementary perspectives, with reference to the matrices in Figure 2 that might be used to model the dynamics involved. Note that this summary is intended to be illustrative, rather than exhaustive; we are not attempting to identify every aspect of the generative model that might be varied to simulate differences in any given personality trait. In this section, we focus on functional descriptions, but we also reference neurobiological details occasionally in order to elucidate parallels between CB5T and FEP-AI because both frameworks use such details as a source of evidence about function.

The Metatraits

In CB5T, the major factors of personality have functional significance in terms of control parameters for the basic cybernetic process of identifying and prioritizing goals, identifying appropriate operators (policies/strategies), and comparing the current state to goal states (DeYoung, 2015; Leibo, Zambaldi, Lanctot, Marecki, & Graepel, 2017). CB5T provides an evolutionary functional account of the mechanisms underlying the traits in Figure 4.

At a level above the Big Five, the metatraits of *Stability* (shared variance of Conscientiousness, Agreeableness, and inverse Neuroticism) and *Plasticity* (shared variance of Extraversion and Openness/Intellect) are particularly notable as the highest-level descriptions of the parameter settings for cybernetic control systems. Stability reflects optimization for resisting disruption of ongoing goal pursuit, an instrumentally useful capacity for a cybernetic system (DeYoung, 2015; DeYoung & Krueger, 2018a). Interpreted in this light, Neuroticism is the opposite pole of Emotional Stability, which discourages potential threats from disrupting goal pursuit due to anxiety or other negative emotions; Conscientiousness reflects motivational stability, such that distractions and disorganization tend not to disrupt goal pursuit; and

Agreeableness reflects social stability, maintaining social harmony and avoiding conflicts with others. Plasticity, in contrast, reflects optimization for a capacity to generate new characteristic adaptations, including the exploration of novel circumstances (DeYoung, 2013, 2015).

Interpreted in this light, Openness/Intellect involves a cognitive/perceptual tendency toward exploration, whereas Extraversion involves a behavioral tendency toward exploration.

Although Stability and Plasticity may seem conceptually opposed, in reality they have the potential to be mutually supporting factors. Indeed, a capacity for plasticity must evolve to enable stability (via enhanced adaptability) in changing and unpredictable environments (Grossberg, 2013). Stability provides a secure foundation for plastic exploration, and plasticity is necessary to update characteristic adaptations so that the person can continue stable goal pursuit in changing circumstances. Nonetheless, the metatraits are in dynamic tension with each other because high levels of plasticity may pose a challenge for stability and vice versa. Certainly not every person has an optimal balance of the two. As described above, this sort of dual-optimization, with both tradeoffs and synergy, can also be found in FEP-AI; the existential imperative of system preservation necessarily requires balancing the extrinsic value of achieving particular goals with the intrinsic value of learning new information (de Abril & Kanai, 2018; Friston, Lin, et al., 2017; Yufik & Friston, 2016).

Within CB5T, the metatrait of Stability corresponds to the protection of goals, interpretations, and strategies from disruption by impulses, and the metatrait of Plasticity corresponds to exploration and the creation of new goals, interpretations, and strategies. Within FEP-AI (Friston et al., 2017), Stability would have its closest correspondence with optimization for the extrinsic/pragmatic value of minimizing prediction-error with respect to achieving particular goals in the world. The most obvious parameter of the generative model in Figure 2

that should correspond to variation in Stability is precision over the C matrix. The higher the precision, the less likely the systems goals (high-level preferences) are to be disrupted by other factors in the model or its environment. (In this context, it is notable that the general risk factor for all forms of psychopathology, known as the *p*-factor, appears to correspond fairly closely to inverse Stability; DeYoung & Krueger, 2018a, 2018b.)

Plasticity, in contrast, may have its closest correspondence with optimization for the intrinsic/epistemic value of information, achieved either via foraging for information through overt attentional orienting or via the exploration of novel policies for their epistemic affordance (that is, their potential for increasing information gain). As discussed above in the discussion of convergence between CB5T and FEP-AI's perspectives on dopamine, Plasticity has its clearest correspondence in FEP-AI with precision over policies, corresponding to the inverse temperature parameters (γ/β) in Figure 2. The higher the precision over policies, the more the agent will be willing to attempt to engage existing policies even under conditions of uncertainty, leading it to explore more prolifically in novel or unpredictable environments.

The Big Five

Within CB5T, Extraversion is understood as corresponding to tendencies for behavioral exploration and engagement with specific rewards. Whereas its exploratory function has already been discussed in relation to Plasticity, engagement with specific rewards would tend to correspond in FEP-AI to the likelihood of releasing value-realizing policies focused on maximizing extrinsic value, which is another aspect of precision over policies (γ/β in Figure 2). Biological evidence supporting this functional hypothesis includes the fact that Extraversion has been convincingly demonstrated to relate to variation in dopaminergic function (Wacker & Smillie, 2015), and fluctuating dopamine levels in the striatum have been identified as a major

factor in determining the likelihood with which cortical ensembles are likely to be disinhibited (Mannella et al., 2013). Elevated dopamine would be viewed as signaling that conditions are likely to be such that deploying policies will be likely to satisfy prior preferences and, therefore, that deployment thresholds ought to be lowered (Hesp, Smith, et al., 2019). Thus, Extraversion in CB5T corresponds to differences in the motivational force the person is likely to bring to bear on any given goal.

CB5T describes Openness/Intellect as cognitive exploration and engagement with information to generate new interpretations of the world. In FEP-AI terms, Openness/Intellect would correspond to the degree to which an agent predicts that it will minimize overall prediction-error by selecting policies that tend to optimize for epistemic value (Friston, Lin, et al., 2017; Schwartenbeck et al., 2019). With respect to generative models, like the one in Figure 2, this could correspond to having increased precision assigned to prediction errors from perceptual A matrices (and so greater sensitivity to novel information), and as particular configurations of C and E matrices in which novelty and learning are predicted to be particularly likely to realize valued outcomes.

Within CB5T, Neuroticism is viewed as reflecting defensive responses to uncertainty, threat, and punishment, and its underlying mechanisms are described in terms of responses to “psychological entropy” (Hirsh et al., 2012), which corresponds to uncertainty regarding what action or interpretation to adopt at a given time. Detection of prediction errors increases psychological entropy, and Neuroticism should reflect greater sensitivity of error detection and also greater tendency to interrupt ongoing goal-directed action when error is detected, as a defense against threat (Gray & McNaughton, 2003). In cybernetic terms, all threats are increases in entropy because they increase uncertainty about the capacity of the system to reach one or

more of its goals. (This account of Neuroticism is also compatible with notions from the recently introduced “Attitudinal Entropy” framework (Dalege et al., 2018).)

In FEP-AI, degree of sensitivity to goal-discrepant observations likely depend heavily on the precision assigned to prediction errors, and hence their abilities to ascend belief hierarchies and drive updating (Barrett & Simmons, 2015; Hesp, Smith, et al., 2019; Joffily & Coricelli, 2013; Seth & Friston, 2016). A P-structure for Neuroticism, therefore, would be likely to involve variation in elements of the B matrix—confidence in how things will play out—as well as to variation in parameters of the function (G) that computes deviation from expectations as free energy. Minimizing prediction errors is always the fundamental imperative in FEP-AI (Solms & Friston, 2018), but excessive sensitivity to individual prediction errors is unlikely to effectively minimize prediction error in the longer term. This is because (1) over-reactions to small errors is likely to lead to over-fitting and reduction in the ability to generalize to new situations, and (2) interrupting action in response to minor errors is likely to increase prediction errors with respect to the prediction of satisfying prior preferences (i.e., achieving goals). In personality terms, this is to say that high Neuroticism—although potentially making people more likely to avoid unpredictable and therefore risky situations—may also undermine goal pursuit by causing myopic attention to minutiae and abandonment of strategies likely to be successful (DeYoung, 2015; DeYoung & Krueger, 2018a).

Conscientiousness is understood, in CB5T, as a trait reflecting the ability to prioritize non-immediate and/or abstract goals and to avoid disruption in progress toward them. It has been found to relate to patterns of connectivity in a large brain network combining what have traditionally been labeled as the salience and ventral attention networks, but also including a node in dorsolateral prefrontal cortex, which is likely to be central to prioritizing goals

effectively (Rueter et al., 2018). Complex cybernetic systems governed by a control hierarchy must have means of differentially weighing goals across contexts, and a combination of higher-level control operations and salience correspondences is precisely what is needed for adjudicating among competing objectives.

The pattern of findings for Conscientiousness is highly compatible with FEP-AI, where the nature of hierarchical policy selection represents an active area of research (Pezzulo et al., 2018). Successfully minimizing expected free energy requires evaluating the likely consequences of policies deployed in the future and so requires inter-temporal modeling of coherent hierarchical active inference (i.e., multi-step planning). Although some aspects of hierarchical control may be explained relatively straightforwardly in terms of hierarchical predictive coding mechanisms (Adams et al., 2013), more complex models have been developed for understanding more abstract and temporally extended goals. For example, it may be necessary to integrate heterogeneous kinds of information into a common currency of value that can allow comparison across goals of widely differing scope (FitzGerald et al., 2009; Levy & Glimcher, 2012; Mannella et al., 2013; Sescousse et al., 2010).

It is notable that personality factors corresponding to Conscientiousness have tended to be identified outside of humans only in our nearest evolutionary neighbors, chimpanzees and bonobos, suggesting how rare it is for organisms to be able to consider goals that are relatively distant in time or to prioritize relatively abstract, rule-oriented goals (Altschul et al., n.d., 2017; Call & Tomasello, 2008; Freeman & Gosling, 2010; Gosling & John, 1999). Importantly, this is not to deny that many other species vary on behavioral dimensions relating to impulsivity, but these forms of impulsivity are probably related to a Stability factor rather than to Conscientiousness specifically (DeYoung, 2015). Conscientiousness may be one of the more

challenging personality traits to characterize with FEP-AI, but recent work on more elaborate generative models may provide advances with respect to simulating metacognition and higher-level awareness for complex planning (Friston et al., 2020; Smith et al., 2020). It will be difficult for an agent to coordinate its behavior across time and various contexts without having useful models of how it may be likely to behave in different situations. At this point, our best guess is that such models of the agent itself might be incorporated into the A matrix in Figure 2, in terms of likelihood mappings between inferred states and observations of oneself.

To conclude our tour of the Big Five, CB5T describes Agreeableness as the shared variance among all traits reflecting altruism, cooperation, and the coordination of goal-pursuit with other persons. Whereas the other Big Five trait dimensions are more clearly related to the fundamental cybernetic processes underlying the pursuit of any goal, Agreeableness exists only because we are a social species and must take the goals of our conspecifics into account in our own actions. Nonetheless, the mechanisms underlying Agreeableness are fundamental for a highly social species like human beings.

Within FEP-AI and machine learning more generally, degrees of Agreeableness would correspond to tendencies towards mutually beneficial policy selection within multi-agent contexts (Rabinowitz et al., 2018). To the extent that an agent has common ground and shared modes of function with other agents, then there may be substantial overlap in representations, so allowing for useful priors to be shared between self and others, and also between others and oneself. Further, interacting agents share causes from latent world states, which influence (and are influenced by) sensory observations, and so they must be able to share model structure to at least some degree. However, in this process of converging upon shared generative models (Friston & Frith, 2015a, 2015b) agents may differ in the degree to which they update their

predictions (which then drive action) based on the activities of their interaction partners, and such differences should be reflected in variation in Agreeableness.

Within the FEP-AI paradigm, the deeply social nature of humans is increasingly being recognized as a necessary starting place for understanding the ways in which mental processes are bootstrapped over cognitive development (Fotopoulou & Tsakiris, 2017; Friston, 2017; Veissière et al., 2019). Such considerations are beyond the scope of the computational model that we discuss here, however, and we do not have any specific suggestions for how to integrate variation in Agreeableness into such models beyond pointing out that a P-structure of Agreeableness would entail variation in the contents of several matrices, such as A, C, and E, such that those high in Agreeableness will have different goals, policies, and representations of the world than someone low in Agreeableness. Of course, this is generally true for all of the other traits as well, but in relation to Agreeableness these contents would vary specifically in terms of their prosociality.

Conclusion: Levels of Analysis in Personality Modeling

One of us is on record expressing skepticism about the value of formally modeling psychological phenomena as a method for understanding psychological individual differences (DeYoung & Krueger, 2018b, in press). The concern is that the actual human agent is so complex that it is likely to be difficult for any artificial agent to resemble the real thing with sufficient accuracy to afford new insights into human psychology. Rather than learning anything new about the human mind and/or brain, one is often likely to learn only about the limitations of the model. This raises an important question: why should one be interested in the approach advocated here? Why should one have any optimism that an agent modeled using FEP-AI may be useful for understanding the nature and sources of personality traits?

The answer has to do with the frequently drawn distinction between computational, algorithmic, and implementational levels of analysis (Marr, 1982). The modeling approach we describe here operates entirely at the computational level, which involves identifying and explaining functions that are carried out by the human cybernetic system. Our modeling approach need not identify features of the specific representations and algorithms that the human brain uses to carry out the computational function of creating a generative model and engaging in active inference to minimize expected free energy. Nor does it require identifying features of the specific neural systems that implement those representations and algorithms (Rudrauf & Debbané, 2018). The availability of algorithmic and implementational process theories such as hierarchical predictive coding (HPC) allows for hypotheses to be generated (and compared against empirical data) at less abstracted levels of analysis, but this is not strictly necessary for conducting research using FEP-AI.

As should be obvious from the rest of the chapter, we do value scientific inquiry into both the algorithmic and implementational levels as a strategy for gaining insight about the sources of personality. Variations in neural function and in the representations and algorithms that such function encodes form the neurobiological basis of differences in personality, and we can learn much by studying them. Nonetheless, we think it is an advantage of the FEP-AI approach that it can ignore these levels in creating models of agents. (This is not to say that it ignores these levels in terms of its scientific applications—far from it—but rather that its generative models are focused on reproducing systems' computational functions, without requiring the specification of particular details of algorithmic/implementational instantiation.) This apparent paradox can be resolved by observing that focusing on the computational level provides a greater opportunity for

verisimilitude, without needing to worry about the whether the model resembles the brain's specific algorithms and their neural implementations.

The specific algorithm used in Bayesian computational models, like the one we suggest here, may be quite different from the algorithms employed by the human brain, and the computer hardware that typically implements the algorithm is certainly nothing like a brain (Griffiths, Vul, & Sanborn, 2012; Jones & Love, 2011; Marcus, 2009). However, by focusing on the computational level in modeling, one can potentially gain insight into human functioning without needing to worry about resemblance at other levels of analysis. We believe that FEP-AI provides an accurate and mathematically precise computational depiction of basic aspects of cybernetic function and, therefore, can potentially provide real insight into the functional properties of human personality traits. Convergent evidence for this computational verisimilitude derives both from more abstract computational analyses rooted in FEP-AI and cybernetics, and also from more concrete analysis of algorithmic and implementational levels in the form of HPC and Bayesian-brain approaches generally (Griffiths et al., 2012; T. L. Griffiths et al., 2015). Given the extent of evidence for HPC as one of the brain's fundamental algorithms, working with a computational model based in FEP-AI is a good bet.

From the perspective of modeling personality, parameters that vary the behavioral patterns of an FEP-AI agent are reasonable candidates for gleaning insights into the functional significance of different personality traits. We suspect this is most true for very broad personality traits, including the Big Five and especially their metatraits, Stability and Plasticity, because such broad traits are likely to reflect variation in parameters of basic cybernetic functions (DeYoung, 2015). Such insights would not tell us anything directly about the specific algorithms or implementations that underlie human personality, due to the many potential ways that any given

computational function may be realized. However, the computational level of analysis is of interest in its own right, and gaining a handle on the functional significance of a given personality trait can help generate hypotheses about relevant brain systems, to the extent that anything is known about which neural processes carry out the computational function in question.

In this chapter, we have begun to outline how artificial agents based on FEP-AI might be used to test CB5T's theoretical propositions regarding the functional significance of dimensions in the Big Five trait hierarchy. Although a full integration between CB5T and FEP-AI will be a non-trivial undertaking, we believe it would be worthwhile in helping us to understand the psychological dimensions along which persons vary. If deep correspondences can be found between CB5T and FEP-AI—as we suggest is likely to be the case due to their common cybernetic foundations—then this integration may allow personality modelers to obtain new understanding with respect to functions and then to parlay that understanding of function into new hypotheses about mechanisms. Finally, identifying the configurations of personality mechanisms that correspond to effective cybernetic function may yield insights into how personality imbalances can lead to mental disorders, thereby furthering the aims of computational psychiatry and other attempts at developing detailed theories of human mental health and well-being (DeYoung & Krueger, 2018a, 2018b; Friston, Redish, et al., 2017).

Table 1. *Definitions of key terms for CB5T and correspondences with FEP-AI.*

CB5T concepts	Definitions	FEP-AI Correspondences
Cybernetics	The study of principles governing goal-directed systems that self-regulate via feedback.	Model-based control. Predictive modeling. Good regulator theorem. Law of requisite variety. Self-information/entropy.
Psychological goal	A representation (conscious or unconscious) of the desired state of some variable, capable of being pursued in part via output through the voluntary muscular system or through the operation of selective attention and working memory.	Bayesian beliefs (as prior expectations) realized via mechanisms of prediction-error minimization, functionally understood as satisfying preferences via policy selection. May be either implicitly or explicitly represented.
Psychopathology	Persistent failure to move toward one's psychological goals due to failure to generate effective new goals, interpretations, or strategies when existing ones prove unsuccessful.	Persistent inability to minimize prediction-error with respect to prior preferences. Grounding of value in homeostatic imperatives may provide partially objective criteria for mental health. Although these would nonetheless contain irreducible subjective aspects (e.g. one need not value health, even if it is instrumentally valuable in most contexts).
Personality	All reasonably enduring psychological individual differences.	Enduring attractors that define systems as (weakly) ergodic processes.
Personality Traits	Probabilistic descriptions of relatively stable patterns of emotion, motivation, cognition, and behavior, in response to classes of stimuli that have been present in human environments over evolutionary time.	Minimal message-length, maximally informative (i.e., “normal form”) descriptions of agents as generative models. Evolutionary and developmentally specified priors, predictions, and default modes of operation.

Characteristic Adaptations	Relatively stable goals, interpretations, and strategies, specified in relation to an individual's particular life circumstances.	Policies and updatable components of memory (as Bayesian beliefs) that influence which actions are likely to be selected in different situations.
-----------------------------------	---	---

References

- Adams, R. A., Moutoussis, M., Nour, M. M., Dahoun, T., Lewis, D., Illingworth, B., Veronese, M., Mathys, C., de Boer, L., Guitart-Masip, M., Friston, K. J., Howes, O. D., & Roiser, J. P. (n.d.). Variability in Action Selection Relates to Striatal Dopamine 2/3 Receptor Availability in Humans: A PET Neuroimaging Study Using Reinforcement Learning and Active Inference Models. *Cerebral Cortex*. <https://doi.org/10.1093/cercor/bhz327>
- Adams, R., Shipp, S., & Friston, K. J. (2013). Predictions not commands: Active inference in the motor system. *Brain Structure & Function*, 218(3), 611–643.
<https://doi.org/10.1007/s00429-012-0475-5>
- Allen, T. A., & DeYoung, C. G. (2017). Personality neuroscience and the Five Factor Model. In T. A. Widiger (Ed.). *Oxford handbook of the Five Factor Model* (pp. 319–349). New York: Oxford University Press.
- Altschul, D. M., Hopkins, W. D., Herrelko, E. S., Inoue-Murayama, M., Matsuzawa, T., King, J. E., Ross, S. R., & Weiss, A. (n.d.). Personality links with lifespan in chimpanzees. *ELife*, 7. <https://doi.org/10.7554/eLife.33781>
- Altschul, D. M., Wallace, E. K., Sonnweber, R., Tomonaga, M., & Weiss, A. (2017). Chimpanzee intellect: Personality, performance and motivation with touchscreen tasks. *Royal Society Open Science*, 4(5). <https://doi.org/10.1098/rsos.170169>
- Austin, J. T., & Vancouver, J. B. (1996). Goal constructs in psychology: Structure, process, and content. *Psychological Bulletin*, 120(3), 338–375.
- Barrett, L. F., & Simmons, W. K. (2015). Interoceptive predictions in the brain. *Nature Reviews. Neuroscience*, 16(7), 419–429. <https://doi.org/10.1038/nrn3950>

- Bastos, Andre M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron*, 76(4), 695–711. <https://doi.org/10.1016/j.neuron.2012.10.038>
- Bastos, André Moraes, Vezoli, J., Bosman, C. A., Schoffelen, J.-M., Oostenveld, R., Dowdall, J. R., De Weerd, P., Kennedy, H., & Fries, P. (2015). Visual Areas Exert Feedforward and Feedback Influences through Distinct Frequency Channels. *Neuron*, 85(2), 390–401. <https://doi.org/10.1016/j.neuron.2014.12.018>
- Bergmann, J., Morgan, A. T., & Muckli, L. (2019). Two distinct feedback codes in V1 for “real” and “imaginary” internal experiences. *BioRxiv*, 664870. <https://doi.org/10.1101/664870>
- Boyd, A. B., Mandal, D., & Crutchfield, J. P. (2017). Leveraging Environmental Correlations: The Thermodynamics of Requisite Variety. *Journal of Statistical Physics*, 167(6), 1555–1585. <https://doi.org/10.1007/s10955-017-1776-0>
- Brillouin, L. (1951). Maxwell’s Demon Cannot Operate: Information and Entropy. I. *Journal of Applied Physics*, 22(3), 334–337. <https://doi.org/10.1063/1.1699951>
- Bruineberg, J., & Rietveld, E. (2014). Self-organization, free energy minimization, and optimal grip on a field of affordances. *Frontiers in Human Neuroscience*, 8. <https://doi.org/10.3389/fnhum.2014.00599>
- Call, J., & Tomasello, M. (2008). Does the chimpanzee have a theory of mind? 30 years later. *Trends in Cognitive Sciences*, 12(5), 187–192. <https://doi.org/10.1016/j.tics.2008.02.010>
- Campbell, J. O. (2016). Universal Darwinism As a Process of Bayesian Inference. *Frontiers in Systems Neuroscience*, 10, 49. <https://doi.org/10.3389/fnsys.2016.00049>

- Carhart-Harris, R. L., & Friston, K. J. (2019). REBUS and the Anarchic Brain: Toward a Unified Model of the Brain Action of Psychedelics. *Pharmacological Reviews*, 71(3), 316–344.
<https://doi.org/10.1124/pr.118.017160>
- Carhart-Harris, Robin L. (2018). The entropic brain—Revisited. *Neuropharmacology*, 142, 167–178. <https://doi.org/10.1016/j.neuropharm.2018.03.010>
- Carhart-Harris, Robin Lester, Leech, R., Hellyer, P. J., Shanahan, M., Feilding, A., Tagliazucchi, E., Chialvo, D. R., & Nutt, D. (2014). The entropic brain: A theory of conscious states informed by neuroimaging research with psychedelic drugs. *Frontiers in Human Neuroscience*, 8, 20.
- Carver, C. S., & Scheier, M. F. (1998). *On the self-regulation of behavior*. Cambridge, England: Cambridge University Press.
- Carver, C. S., Sutton, S. K., & Scheier, M. F. (2000). Action, emotion, and personality: Emerging conceptual integration. *Personality and Social Psychology Bulletin*, 26(6), 741–775.
- Chao, Z. C., Takaura, K., Wang, L., Fujii, N., & Dehaene, S. (2018). Large-Scale Cortical Networks for Hierarchical Prediction and Prediction Error in the Primate Brain. *Neuron*, 100(5), 1252–1266.e3. <https://doi.org/10.1016/j.neuron.2018.10.004>
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *The Behavioral and Brain Sciences*, 36(3), 181–204.
<https://doi.org/10.1017/S0140525X12000477>
- Conant, R., C., & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2), 89–97.
<https://doi.org/10.1080/00207727008920220>

- Constant, A., Ramstead, M. J. D., Veissière, S. P. L., & Friston, K. J. (2019). Regimes of Expectations: An Active Inference Model of Social Conformity and Human Decision Making. *Frontiers in Psychology, 10*. <https://doi.org/10.3389/fpsyg.2019.00679>
- Dalege, J., Borsboom, D., Harreveld, F. van, & Maas, H. L. J. van der. (2018). The Attitudinal Entropy (AE) Framework as a General Theory of Individual Attitudes. *Psychological Inquiry, 29*(4), 175–193. <https://doi.org/10.1080/1047840X.2018.1537246>
- Damasio, A. (2012). *Self Comes to Mind: Constructing the Conscious Brain* (Reprint edition). Vintage.
- Dayan, P., Hinton, G. E., Neal, R. M., & Zemel, R. S. (1995). The Helmholtz machine. *Neural Computation, 7*(5), 889–904.
- de Abril, I. M., & Kanai, R. (2018). A unified strategy for implementing curiosity and empowerment driven reinforcement learning. *ArXiv:1806.06505 [Cs]*.
<http://arxiv.org/abs/1806.06505>
- DeYoung, C.G., & Krueger, R. F. (In Press). A cybernetic theory of psychopathology. *Psychological Inquiry*.
- DeYoung, C. G. (2013). The neuromodulator of exploration: A unifying theory of the role of dopamine in personality. *Frontiers in Human Neuroscience, 7*.
<https://doi.org/10.3389/fnhum.2013.00762>
- DeYoung, C. G. (2015). Cybernetic Big Five Theory. *Journal of Research in Personality, 56*, 33–58. <https://doi.org/10.1016/j.jrp.2014.07.004>
- DeYoung, C. G., & Krueger, R. F. (2018a). A Cybernetic Theory of Psychopathology. *Psychological Inquiry, 29*(3), 117–138. <https://doi.org/10.1080/1047840X.2018.1513680>

- DeYoung, C. G., & Krueger, R. F. (2018b). Understanding Psychopathology: Cybernetics and Psychology on the Boundary between Order and Chaos. *Psychological Inquiry*.
<https://www.tandfonline.com/doi/pdf/10.1080/1047840X.2018.1513690>
- DeYoung, C. G., & Krueger, R. F. (in press). To wish impossible things: On the ontological status of latent variables and the prospects for theory in psychology. *Psychological Inquiry*.
- DeYoung, C. G., & Weisberg, Y. J. (2019). Cybernetic approaches to personality and social behavior. In M. Snyder & K. Deaux (Eds). *Oxford handbook of personality and social psychology, Second edition* (387–414). New York: Oxford University Press.
- Erritzoe, D., Smith, J., Fisher, P. M., Carhart-Harris, R., Frokjaer, V. G., & Knudsen, G. M. (2019). Recreational use of psychedelics is associated with elevated personality trait openness: Exploration of associations with brain serotonin markers. *Journal of Psychopharmacology (Oxford, England)*, 269881119827891.
<https://doi.org/10.1177/0269881119827891>
- FitzGerald, T. H. B., Dolan, R. J., & Friston, K. J. (2015). Dopamine, reward learning, and active inference. *Frontiers in Computational Neuroscience*, 9.
<https://doi.org/10.3389/fncom.2015.00136>
- FitzGerald, T. H. B., Seymour, B., & Dolan, R. J. (2009). The role of human orbitofrontal cortex in value comparison for incommensurable objects. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 29(26), 8388–8395.
<https://doi.org/10.1523/JNEUROSCI.0717-09.2009>

- Fleeson, W. (2001). Toward a structure- and process-integrated view of personality: Traits as density distributions of states. *Journal of Personality and Social Psychology*, 80(6), 1011–1027. <https://doi.org/10.1037/0022-3514.80.6.1011>
- Fleeson, W., & Gallagher, M. P. (2009). The Implications of Big-Five Standing for the Distribution of Trait Manifestation in Behavior: Fifteen Experience-Sampling Studies and a Meta-Analysis. *Journal of Personality and Social Psychology*, 97(6), 1097–1114. <https://doi.org/10.1037/a0016786>
- Fleeson, W., & Law, M. K. (2015). Trait Enactments as Density Distributions: The Role of Actors, Situations, and Observers in Explaining Stability and Variability. *Journal of Personality and Social Psychology*, 109(6), 1090–1104. <https://doi.org/10.1037/a0039517>
- Fotopoulou, A., & Tsakiris, M. (2017). Mentalizing homeostasis: The social origins of interoceptive inference—replies to Commentaries. *Neuropsychanalysis*, 19(1), 71–76.
- Freeman, H. D., & Gosling, S. D. (2010). Personality in nonhuman primates: A review and evaluation of past research. *American Journal of Primatology*, 72(8), 653–671. <https://doi.org/10.1002/ajp.20833>
- Friston, K. (n.d.) *Karl Friston: Active inference and artificial curiosity*. Retrieved June 18, 2020, from <https://www.youtube.com/watch?v=Y1egnoCWgUg>
- Friston, K. (2013). Life as we know it. *Journal of The Royal Society Interface*, 10(86), 20130475. <https://doi.org/10.1098/rsif.2013.0475>
- Friston, K., Da Costa, L., Hafner, D., Hesp, C., & Parr, T. (2020). *Sophisticated Inference*. <https://arxiv.org/abs/2006.04120v1>

- Friston, K. J. (1994). Functional and effective connectivity in neuroimaging: A synthesis. *Human Brain Mapping*, 2(1–2), 56–78. <https://doi.org/10.1002/hbm.460020107>
- Friston, K. J. (2017). Self-evidencing babies: Commentary on “Mentalizing homeostasis: The social origins of interoceptive inference” by Fotopoulou & Tsakiris. *Neuropsychanalysis*, 19(1), 43–47.
- Friston, K. J. (2019). A free energy principle for a particular physics. *ArXiv:1906.10184 [q-Bio]*. <http://arxiv.org/abs/1906.10184>
- Friston, K. J., Brown, H. R., Siemerikus, J., & Stephan, K. E. (2016). The dysconnection hypothesis (2016). *Schizophrenia Research*, 176(2–3), 83–94. <https://doi.org/10.1016/j.schres.2016.07.014>
- Friston, K. J., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017). Active Inference: A Process Theory. *Neural Computation*, 29(1), 1–49. https://doi.org/10.1162/NECO_a_00912
- Friston, K. J., & Frith, C. (2015a). A Duet for one. *Consciousness and Cognition*, 36, 390–405. <https://doi.org/10.1016/j.concog.2014.12.003>
- Friston, K. J., & Frith, C. D. (2015b). Active inference, communication and hermeneutics. *Cortex; a Journal Devoted to the Study of the Nervous System and Behavior*, 68, 129–143. <https://doi.org/10.1016/j.cortex.2015.03.025>
- Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J.-P., Frith, C. D., & Frackowiak, R. S. J. (1994). Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping*, 2(4), 189–210. <https://doi.org/10.1002/hbm.460020402>

- Friston, K. J., & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 364(1521), 1211–1221. <https://doi.org/10.1098/rstb.2008.0300>
- Friston, K. J., Lin, M., Frith, C. D., Pezzulo, G., Hobson, J. A., & Ondobaka, S. (2017). Active Inference, Curiosity and Insight. *Neural Computation*, 29(10), 2633–2683. https://doi.org/10.1162/neco_a_00999
- Friston, K., Parr, T., & Zeidman, P. (2019). Bayesian model reduction. *ArXiv:1805.07092 [Stat]*. <http://arxiv.org/abs/1805.07092>
- Friston, K. J., Redish, A. D., & Gordon, J. A. (2017). Computational Nosology and Precision Psychiatry. *Computational Psychiatry*, 1, 2–23. https://doi.org/10.1162/CPSY_a_00001
- Friston, K. J., Schwartenbeck, P., FitzGerald, T., Moutoussis, M., Behrens, T., & Dolan, R. J. (2014). The anatomy of choice: Dopamine and decision-making. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1655). <https://doi.org/10.1098/rstb.2013.0481>
- Friston, K. J., Shiner, T., FitzGerald, T., Galea, J. M., Adams, R., Brown, H., Dolan, R. J., Moran, R., Stephan, K. E., & Bestmann, S. (2012). Dopamine, affordance and active inference. *PLoS Computational Biology*, 8(1), e1002327. <https://doi.org/10.1371/journal.pcbi.1002327>
- Gopnik, A., O’Grady, S., Lucas, C. G., Griffiths, T. L., Wente, A., Bridgers, S., Aboody, R., Fung, H., & Dahl, R. E. (2017). Changes in cognitive flexibility and hypothesis search across human life history from childhood to adolescence to adulthood. *Proceedings of the National Academy of Sciences*, 114(30), 7892–7899. <https://doi.org/10.1073/pnas.1700811114>

- Gosling, S. D., & John, O. P. (1999). Personality Dimensions in Nonhuman Animals: A Cross-Species Review. *Current Directions in Psychological Science*.
<https://journals.sagepub.com/doi/10.1111/1467-8721.00017>
- Gray, J. A., & McNaughton, N. (2003). *The Neuropsychology of Anxiety: An enquiry into the function of the septo-hippocampal system*. OUP Oxford.
- Griffiths, T. L., Lieder, F., & Goodman, N. D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science*, 7(2), 217–229. <https://doi.org/10.1111/tops.12142>
- Griffiths, T., Vul, E., & Sanborn, A. (2012). *Bridging Levels of Analysis for Probabilistic Models of Cognition*. <https://doi.org/10.1177/0963721412447619>
- Grossberg, S. (2013). Adaptive Resonance Theory: How a brain learns to consciously attend, learn, and recognize a changing world. *Neural Networks*, 37, 1–47.
- Ha, D., & Schmidhuber, J. (2018). World Models. *ArXiv:1803.10122 [Cs, Stat]*.
<https://doi.org/10.5281/zenodo.1207631>
- Hafner, D., Lillicrap, T., Ba, J., & Norouzi, M. (2020). Dream to Control: Learning Behaviors by Latent Imagination. *ArXiv:1912.01603 [Cs]*. <http://arxiv.org/abs/1912.01603>
- Hassabis, D., Spreng, R. N., Rusu, A. A., Robbins, C. A., Mar, R. A., & Schacter, D. L. (2014). Imagine All the People: How the Brain Creates and Uses Personality Models to Predict Behavior. *Cerebral Cortex*, 24(8), 1979–1987. <https://doi.org/10.1093/cercor/bht042>
- Hawkins, J., & Blakeslee, S. (2004). *On Intelligence* (Adapted). Times Books.
- Helmholtz, H. (1878). The Facts in Perception. In R. Kahl (Ed.), *Selected Writings of Hermann Helmholtz*. Wesleyan University Press.
- Helmholtz, H. von. (1910). *Treatise on Physiological Optics*. Courier Corporation.

Hesp, C. (n. d.) “*Deeply felt affect: Understanding emotions through deep active inference*” —

Retrieved April 23, 2020, from <https://www.youtube.com/watch?v=BLDqaERoZUQ>

Hesp, C., Smith, R., Allen, M., Friston, K., & Ramstead, M. (2019). *Deeply Felt Affect: The Emergence of Valence in Deep Active Inference* [Preprint]. PsyArXiv.

<https://doi.org/10.31234/osf.io/62pfd>

Hesp, C., Steenbeek, H. W., & van Geert, P. L. C. (2019). Socio-Emotional Concern Dynamics in a Model of Real-Time Dyadic Interaction: Parent-Child Play in Autism. *Frontiers in Psychology, 10*. <https://doi.org/10.3389/fpsyg.2019.01635>

Heylighen, F. (1992). Principles of Systems and Cybernetics: An evolutionary perspective. *Cybernetics and Systems, 92*(1992), 3–10.

Hirsh, J. B., Mar, R. A., & Peterson, J. B. (2012). Psychological entropy: A framework for understanding uncertainty-related anxiety. *Psychological Review, 119*(2), 304–320. <https://doi.org/10.1037/a0026767>

Hirsh, J. B., Mar, R. A., & Peterson, J. B. (2013). Personal narratives as the highest level of cognitive integration. *The Behavioral and Brain Sciences, 36*(3), 216–217. <https://doi.org/10.1017/S0140525X12002269>

Hohwy, J. (2020). Self-supervision, normativity and the free energy principle. *Synthese*. <https://doi.org/10.1007/s11229-020-02622-2>

Joffily, M., & Coricelli, G. (2013). Emotional valence and the free-energy principle. *PLoS Computational Biology, 9*(6), e1003094. <https://doi.org/10.1371/journal.pcbi.1003094>

John, O. P., Naumann, L. P., & Soto, C. J. (2008). Paradigm shift to the integrative Big Five trait taxonomy: History: measurement, and conceptual issue. In O. P. John, R. W. Robins, &

- L. A. Pervin (Eds.). *Handbook of personality: Theory and research* (pp. 114–158). New York, NY: Guilford Press.
- Jones, M., & Love, B. C. (2011). Bayesian Fundamentalism or Enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *The Behavioral and Brain Sciences*, 34(4), 169–188; discussion 188-231.
<https://doi.org/10.1017/S0140525X10003134>
- Kaplan, R., & Friston, K. J. (2018). Planning and navigation as active inference. *Biological Cybernetics*, 112(4), 323–343. <https://doi.org/10.1007/s00422-018-0753-2>
- Kirchhoff, M., Parr, T., Palacios, E., Friston, K. J., & Kiverstein, J. (2018). The Markov blankets of life: Autonomy, active inference and the free energy principle. *Journal of the Royal Society Interface*, 15(138). <https://doi.org/10.1098/rsif.2017.0792>
- Kiverstein, J., Miller, M., & Rietveld, E. (2019). The feeling of grip: Novelty, error dynamics, and the predictive brain. *Synthese*, 196(7), 2847–2869. <https://doi.org/10.1007/s11229-017-1583-9>
- Koller, D., & Friedman, N. (2009). *Probabilistic Graphical Models: Principles and Techniques*. MIT Press.
- Lawson, R. P., Rees, G., & Friston, K. J. (2014). An aberrant precision account of autism. *Frontiers in Human Neuroscience*, 8. <https://doi.org/10.3389/fnhum.2014.00302>
- Leibo, J. Z., Zambaldi, V., Lanctot, M., Marecki, J., & Graepel, T. (2017). Multi-agent Reinforcement Learning in Sequential Social Dilemmas. *ArXiv:1702.03037 [Cs]*.
<http://arxiv.org/abs/1702.03037>

- Levy, D. J., & Glimcher, P. W. (2012). The root of all value: A neural common currency for choice. *Current Opinion in Neurobiology*, 22(6), 1027–1038.
<https://doi.org/10.1016/j.conb.2012.06.001>
- Mannella, F., Gurney, K., & Baldassarre, G. (2013). The nucleus accumbens as a nexus between values and goals in goal-directed behavior: A review and a new hypothesis. *Frontiers in Behavioral Neuroscience*, 7, 135. <https://doi.org/10.3389/fnbeh.2013.00135>
- Marcus, G. (2009). *Kluge: The Haphazard Evolution of the Human Mind*. Houghton Mifflin Harcourt.
- Markon, K. E., Krueger, R. F., & Watson, D. (2005). Delineating the structure of normal and abnormal personality: An integrative hierarchical approach. *Journal of Personality and Social Psychology*, 88, 139–157. doi:10.1037/0022-3514.88.1.139
- Markram, K., & Markram, H. (2010). The intense world theory—A unifying theory of the neurobiology of autism. *Frontiers in Human Neuroscience*, 4, 224.
<https://doi.org/10.3389/fnhum.2010.00224>
- Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. Henry Holt and Company.
- Miller, G. A., Galanter, E., & Pribram, K. H. (1960). *Plans and the structure of behavior*. New York: Holt.
- Ramstead, M. (n.d.). *Maxwell Ramstead —A tutorial on active inference*. Retrieved April 23, 2020, from <https://www.youtube.com/watch?v=WzFQzFZiwzk>
- Smith, R. (n.d.). “*Modeling Task Behavior with Active Inference*”— Retrieved April 23, 2020, from <https://www.youtube.com/watch?v=e7nIK8leDGU>

- Mumford, D. (1991). On the computational architecture of the neocortex. *Biological Cybernetics*, 65(2), 135–145. <https://doi.org/10.1007/BF00202389>
- Park, H.-J., & Friston, K. (2013). Structural and functional brain networks: From connections to cognition. *Science (New York, N.Y.)*, 342(6158), 1238411. <https://doi.org/10.1126/science.1238411>
- Park, H.-J., Friston, K. J., Pae, C., Park, B., & Razi, A. (2018). Dynamic effective connectivity in resting state fMRI. *NeuroImage*, 180(Pt B), 594–608. <https://doi.org/10.1016/j.neuroimage.2017.11.033>
- Parr, T., Corcoran, A. W., Friston, K. J., & Hohwy, J. (2019). Perceptual awareness and active inference. *Neuroscience of Consciousness*, 2019(1). <https://doi.org/10.1093/nc/niz012>
- Parr, T., & Friston, K. J. (2017a). Uncertainty, epistemics and active inference. *Journal of the Royal Society, Interface*, 14(136). <https://doi.org/10.1098/rsif.2017.0376>
- Parr, T., & Friston, K. J. (2017b). Working memory, attention, and salience in active inference. *Scientific Reports*, 7(1), 14678. <https://doi.org/10.1038/s41598-017-15249-0>
- Parr, T., & Friston, K. J. (2018). The Discrete and Continuous Brain: From Decisions to Movement-And Back Again. *Neural Computation*, 30(9), 2319–2347. https://doi.org/10.1162/neco_a_01102
- Penny, W. D., Friston, K. J., Ashburner, J. T., Kiebel, S. J., & Nichols, T. E. (2011). *Statistical Parametric Mapping: The Analysis of Functional Brain Images*. Elsevier.
- Pezzulo, G., & Friston, K. J. (2019). The value of uncertainty: An active inference perspective. *Behavioral and Brain Sciences*, 42. <https://doi.org/10.1017/S0140525X18002066>

- Pezzulo, G., Rigoli, F., & Friston, K. J. (2018). Hierarchical Active Inference: A Theory of Motivated Control. *Trends in Cognitive Sciences*, 22(4), 294–306.
<https://doi.org/10.1016/j.tics.2018.01.009>
- Powers, W. T. (1973). *Behaviour: The control of perception*. Chicago, IL: Aldine.
- Preller, K. H., Razi, A., Zeidman, P., Stämpfli, P., Friston, K. J., & Vollenweider, F. X. (2019). Effective connectivity changes in LSD-induced altered states of consciousness in humans. *Proceedings of the National Academy of Sciences of the United States of America*, 116(7), 2743–2748. <https://doi.org/10.1073/pnas.1815129116>
- Rabinowitz, N. C., Perbet, F., Song, H. F., Zhang, C., Eslami, S. M. A., & Botvinick, M. (2018). Machine Theory of Mind. *ArXiv:1802.07740 [Cs]*. <http://arxiv.org/abs/1802.07740>
- Ramstead, M. J. D., Badcock, P. B., & Friston, K. J. (2018). Answering Schrödinger's question: A free-energy formulation. *Physics of Life Reviews*, 24, 1–16.
<https://doi.org/10.1016/j.plrev.2017.09.001>
- Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 79–87. <https://doi.org/10.1038/4580>
- Rueter, A. R., Abram, S. V., MacDonald, A. W., Rustichini, A., & DeYoung, C. G. (2018). The goal priority network as a neural substrate of Conscientiousness. *Human Brain Mapping*, 39(9), 3574–3585. <https://doi.org/10.1002/hbm.24195>
- Rudrauf, D., & Debbané, M. (2018). Building a cybernetic model of psychopathology: Beyond the metaphor. *Psychological Inquiry*, 29(3), 156–164.
- Safron, A. (2020). An Integrated World Modeling Theory (IWMT) of Consciousness: Combining Integrated Information and Global Neuronal Workspace Theories With the

- Free Energy Principle and Active Inference Framework; Toward Solving the Hard Problem and Characterizing Agentic Causation. *Frontiers in Artificial Intelligence*, 3. <https://doi.org/10.3389/frai.2020.00030>
- Samuelson, P. A. (1948). Consumption Theory in Terms of Revealed Preference. *Economica*, 15(60), 243–253. JSTOR. <https://doi.org/10.2307/2549561>
- Sandved Smith, L., Hesp, C., Lutz, A., Mattout, J., Friston, K., & Ramstead, M. (2020). *Towards a formal neurophenomenology of metacognition: Modelling meta-awareness, mental action, and attentional control with deep active inference* [Preprint]. PsyArXiv. <https://doi.org/10.31234/osf.io/5jh3c>
- Schrodinger, E. (1944). *What is Life?: With Mind and Matter and Autobiographical Sketches* (Reprint edition). Cambridge University Press.
- Schwartenbeck, P., Passecker, J., Hauser, T. U., FitzGerald, T. H., Kronbichler, M., & Friston, K. J. (2019). Computational mechanisms of curiosity and goal-directed exploration. *ELife*, 8. <https://doi.org/10.7554/eLife.41703>
- Sedley, W., Gander, P. E., Kumar, S., Kovach, C. K., Oya, H., Kawasaki, H., Howard, M. A., III, & Griffiths, T. D. (2016). Neural signatures of perceptual inference. *ELife*, 5, e11476. <https://doi.org/10.7554/eLife.11476>
- Sescousse, G., Redouté, J., & Dreher, J.-C. (2010). The architecture of reward value coding in the human orbitofrontal cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 30(39), 13095–13104. <https://doi.org/10.1523/JNEUROSCI.3501-10.2010>
- Seth, A. K. (2014). *The Cybernetic Bayesian Brain*. Open MIND. Frankfurt am Main: MIND Group. <https://doi.org/10.15502/9783958570108>

- Seth, A. K., & Friston, K. J. (2016). Active interoceptive inference and the emotional brain. *Phil. Trans. R. Soc. B*, 371(1708), 20160007. <https://doi.org/10.1098/rstb.2016.0007>
- Smigielski, L., Scheidegger, M., Komater, M., & Vollenweider, F. X. (2019). Psilocybin-assisted mindfulness training modulates self-consciousness and brain default mode network connectivity with lasting effects. *NeuroImage*, 196, 207–215. <https://doi.org/10.1016/j.neuroimage.2019.04.009>
- Smith, R., Parr, T., & Friston, K. J. (2019). Simulating Emotions: An Active Inference Model of Emotional State Inference and Emotion Concept Learning. *Frontiers in Psychology*, 10. <https://doi.org/10.3389/fpsyg.2019.02844>
- Tooby, J., Cosmides, L., & Barrett, H. C. (2003). The second law of thermodynamics is the first law of psychology: Evolutionary developmental psychology and the theory of tandem, coordinated inheritances: comment on Lickliter and Honeycutt (2003). *Psychological Bulletin*, 129(6), 858–865. <https://doi.org/10.1037/0033-2909.129.6.858>
- Tschantz, A., Baltieri, M., Seth, A. K., & Buckley, C. L. (2019). Scaling active inference. *ArXiv:1911.10601 [Cs, Eess, Math, Stat]*. <http://arxiv.org/abs/1911.10601>
- Tschantz, A., Millidge, B., Seth, A. K., & Buckley, C. L. (2020). Reinforcement Learning through Active Inference. *ArXiv:2002.12636 [Cs, Eess, Math, Stat]*. <http://arxiv.org/abs/2002.12636>
- Van Egeren, L. F. (2009). A cybernetic model of global personality traits. *Personality and Social Psychology Review*, 13(2), 92–108.
- Veissière, S. P. L., Constant, A., Ramstead, M. J. D., Friston, K. J., & Kirmayer, L. J. (2019). Thinking Through Other Minds: A Variational Approach to Cognition and Culture. *The Behavioral and Brain Sciences*, 1–97. <https://doi.org/10.1017/S0140525X19001213>

- Wacker, J., & Smillie, L. D. (2015). Trait Extraversion and Dopamine Function. *Social and Personality Psychology Compass*, 9(6), 225–238. <https://doi.org/10.1111/spc3.12175>
- Waddington, C. H. (1942). Canalization of development and the inheritance of acquired characteristics. *Nature*, 150(3811), 150563a0. <https://doi.org/10.1038/150563a0>
- Wade, S., & Kidd, C. (2019). The role of prior knowledge and curiosity in learning. *Psychonomic Bulletin & Review*. <https://doi.org/10.3758/s13423-019-01598-6>
- Waller, N. G., DeYoung, C. G., & Bouchard, T. J. (2016). The recaptured scale technique: A method for testing the structural robustness of personality scales. *Multivariate Behavioral Research*, 51, 433–445.
- Wang, J. X., Kurth-Nelson, Z., Kumaran, D., Tirumala, D., Soyer, H., Leibo, J. Z., Hassabis, D., & Botvinick, M. (2018). Prefrontal cortex as a meta-reinforcement learning system. *Nature Neuroscience*, 21(6), 860. <https://doi.org/10.1038/s41593-018-0147-8>
- Wiener, N. (1961). *Cybernetics, or control and communication in the animal and the machine*, 2nd ed. John Wiley & Sons Inc. <https://doi.org/10.1037/13140-000>
- Yufik, Y. M., & Friston, K. J. (2016). Life and Understanding: The Origins of “Understanding” in Self-Organizing Nervous Systems. *Frontiers in Systems Neuroscience*, 10. <https://doi.org/10.3389/fnsys.2016.00098>
- Zador, A. M. (2019). A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature Communications*, 10(1), 1–7. <https://doi.org/10.1038/s41467-019-11786-6>