

Allostasis, interoception, and the free energy principle: Feeling our way forward

Andrew W. Corcoran & Jakob Hohwy

Cognition & Philosophy Laboratory

Monash University

Melbourne, Australia

This material has been accepted for publication in the forthcoming volume entitled *The interoceptive mind: From homeostasis to awareness*, edited by Manos Tsakiris and Helena De Preester. It has been reprinted by permission of Oxford University Press:

[https://global.oup.com/academic/product/the-interoceptive-mind-9780198811930?q=](https://global.oup.com/academic/product/the-interoceptive-mind-9780198811930?q=interoception&lang=en&cc=gb)

[interoception&lang=en&cc=gb](https://global.oup.com/academic/product/the-interoceptive-mind-9780198811930?q=interoception&lang=en&cc=gb). For permission to reuse this material, please visit

<http://www.oup.com/academic/rights/permissions>.

Abstract

Interoceptive processing is commonly understood in terms of the monitoring and representation of the body's current physiological (i.e. homeostatic) status, with aversive sensory experiences encoding some impending threat to tissue viability. However, claims that homeostasis fails to fully account for the sophisticated regulatory dynamics observed in complex organisms have led some theorists to incorporate predictive (i.e. allostatic) regulatory mechanisms within broader accounts of interoceptive processing. Critically, these frameworks invoke diverse – and potentially mutually inconsistent – interpretations of the role allostasis plays in the scheme of biological regulation. This chapter argues in favour of a moderate, reconciliatory position in which homeostasis and allostasis are conceived as equally vital (but functionally distinct) modes of physiological control. It explores the implications of this interpretation for free energy-based accounts of interoceptive inference, advocating a similarly complementary (and hierarchical) view of homeostatic and allostatic processing.

Keywords: allostasis, homeostasis, interoception, free energy principle, prediction error minimisation, biological regulation

1 Introduction

The free energy principle (Friston, 2010) invokes variational Bayesian methods to explain how biological systems maximise evidence for their predictive models via the minimisation of variational free energy, a tractable information-theoretic quantification of prediction error. This account, which was originally proposed to explain sensory learning, has evolved into a much broader scheme encompassing action and motor control, decision-making, attention, communication, and many other aspects of mental function (for overviews see Clark, 2013, 2016; Hohwy, 2013). Under the free energy principle, minimisation of free energy is what any self-organising system is compelled to do in order to resist dissipation and maximise the evidence for its own existence (i.e. self-evidencing through active inference; Hohwy, 2016).

Recent years have witnessed a growing interest in extending the conceptual apparatus of the free energy principle to the interoceptive domain. A number of investigators have sought to explain the influence of interoceptive prediction error minimisation on various physical and cognitive processes (for recent reviews, see Barrett, 2017; Khalsa et al., in press; Seth & Friston, 2016; Smith, Thayer, Khalsa, & Lane, 2017). Central to such *interoceptive inference* perspectives is the notion that interoceptive signals encode representations of the internal (physiological) state of the body, thus providing vital information about how well the organism is managing to preserve the biological viability of its internal environment. Traditionally, the latter has been conceived in terms of *homeostasis*, a concept that usually refers (minimally) to the process of maintaining

44 the internal conditions of complex, thermodynamically open, self-organising biological systems in stable,
45 far-from-equilibrium states (Yates, 1996). From the perspective of the free energy principle, homeostasis
46 translates to the process of restricting the organism to visiting a relatively small number of states that are
47 conducive to its ongoing existence, with interoceptive prediction error playing a particularly important role in
48 signaling deviation from these attractive states (technically, these are known as attracting sets).

49 Notably, the centrality of homeostasis in some free energy-inspired accounts of interoceptive processing has
50 started to give way to the newer concept of *allostasis*. According to proponents of the latter, homeostasis
51 fails to capture the rich variety of self-regulatory processes that complex biological systems engage in in
52 order to conserve their own integrity. Allostasis tries to address this shortcoming through various theoretical
53 innovations, chief amongst which is a core emphasis on predictive or anticipatory modes of regulation. This
54 is to say that, rather than merely responding to physiological perturbations in order to ensure the internal
55 conditions of the body remain within homeostatic bounds, allostasis enables the organism to proactively
56 prepare for such disturbances *before* they occur. While this account carries obvious appeal from the perspective
57 of predictive model-based theories of interoceptive processing, attempts to marry the two have given rise
58 to a number of divergent interpretations of allostatic regulation. As it turns out, the history of allostasis
59 is a history of contested definitions; some 30 years on from its inception, there appears to be no definitive
60 consensus as to its precise meaning. The key aims of this chapter, then, are to establish (a) how allostasis
61 might be best understood as a distinctive concept in the overall scheme of biological regulation, and (b)
62 how this construal might inform (and indeed, be informed by) free energy-inspired theories of interoceptive
63 inference.

64 2 Discovering ‘the wisdom of the body’: Homeostasis

65 A standard account of the history of homeostasis might trace its source to the nineteenth century physiologist
66 Claude Bernard, whose pioneering work on the role of the nervous system in maintaining the relative
67 constancy of internal states (*le milieu intérieur*, i.e. the extracellular fluid environment that envelops the
68 cell) would prove highly influential (Cooper, 2008; Woods & Ramsay, 2007). The key ideas at the core of
69 Bernard’s thinking – notions of harmony, equilibrium, and regulation – are, however, much older, dating
70 as far back as the pre-Socratics (see Adolph, 1961, for a historical review). Bernard refined the ancient
71 insight that organisms maintain a healthy constitution by engaging in certain self-regulatory behaviours (e.g.,
72 consuming nutrients, excreting waste) – and deviate from well-being whenever subject to certain unfavourable
73 physiological imbalances – by drawing attention to the physiological mechanisms that ensure the continuity of
74 a stable internal environment. Such compensatory adjustments act to cancel out internal disturbances that
75 would otherwise be caused by fluctuations in the external environment. This capacity to meet environmental
76 impingements with countervailing responses thus grants the organism an adaptive coupling with – and a
77 special kind of autonomy from – its ecological niche.

78 Benefitting from Bernard’s keen insights and some 50 years of subsequent experimental research, Walter
79 Cannon (1929, 1939) coined the term ‘homeostasis’ to describe the organism’s capacity to maintain a ‘steady
80 state’ or intrinsic uniformity despite ongoing fluctuations in its internal and external processes. Cannon
81 was at pains, however, to stress that his neologism was intended to characterise a complex process in which
82 multiple physiological mechanisms are recruited to ensure the continued stability of the organism’s internal
83 milieu, where stability is construed in terms of a more or less variable range of acceptable (i.e. viable) values.
84 This latter point is crucial for distinguishing Cannon’s conception of homeostasis from Bernard’s emphasis
85 on the fixed, invariant nature of internal conditions, in as much as homeostatic processes admit a space of
86 permissible states. Also important was Cannon’s concern to elucidate the autonomic mechanisms responsible
87 for mediating adaptive physiological responses (e.g., increased respiratory rate) to altered internal conditions
88 (e.g., decreased blood oxygen concentration) (Cooper, 2008), a subtle reorientation that would prove highly
89 influential for later work in cybernetics.

90 The basic concept of homeostasis elaborated by Cannon (and extended by contemporaries such as Curt
91 Richter; see Woods & Ramsay, 2007) would become one of, if not *the* core theoretical principle of modern
92 physiology (Michael & McFarland, 2011; Michael, Modell, McFarland, & Cliff, 2009). One essential element of

93 modern conceptions of homeostasis that was however still missing from the Cannonian picture was a formal
94 account of negative feedback (Modell et al., 2015). From a control-theoretic perspective, Cannon’s careful
95 analysis of particular homeostatic processes can be conceived according to a generic scheme of error detection
96 (i.e. where some regulated variable, for instance blood glucose concentration, is found to deviate from some
97 desirable value or *setpoint*) and correction (i.e. where some effector mechanism is activated in order to restore
98 the regulated variable to the prescribed setpoint). It is important to note here that the notion of a setpoint
99 generally conforms to Cannon’s conception of a (broader or narrower) range of acceptable values, rather than
100 any singular, fixed level (Modell et al., 2015). This set of values can thus be construed as a model against
101 which the actual (sensed) state of the regulated variable is compared. The error signal elicited when the
102 current state of the monitored tissue deviates from its setpoint reference represents a threat to organismic
103 viability, and must therefore be corrected via mobilisation of the appropriate effector system(s). Recasting
104 homeostasis in this light thus furnishes a powerful conceptual framework in which the processes responsible for
105 maintaining internal stability achieve this goal through the communication of information between peripheral
106 tissues and a central controller (such as the central nervous system).

107 3 Allostasis: The future of homeostatic regulation?

108 As mentioned in our introduction, several recent theoretical frameworks of interoceptive inference co-opt notions
109 of allostasis in order to situate the autonomic regulation of the internal milieu within the broader scheme
110 of hierarchical predictive processing. Various theorists have argued that the basic concept of homeostasis
111 is somehow insufficient to account for the rich complexity of self-regulatory behaviour evinced by humans
112 and other animals, advocating allostasis as a necessary theoretical supplement or corrective. To what
113 extent allostasis extends, encompasses, or eliminates homeostasis is, however, unclear, not least because the
114 characteristic features of allostatic regulation have been espoused in ambiguous or inconsistent terms across
115 the literature (Lowe, Almér, & Dodig-Crnkovic, 2017; Power, 2004; Schulkin, 2004). This section thus aims to
116 canvass some of the most influential accounts of allostasis to have emerged over the past three decades.

117 3.1 Achieving stability through change

118 The term ‘allostasis’ was originally introduced by Sterling and Eyer (1988) to describe the integrated,
119 hierarchical mechanisms through which the nervous system maintains organismic integrity. In this scheme,
120 the brain is responsible for orchestrating complex, multi-system responses to physiological perturbations,
121 resulting in a cascade of mutually reinforcing effects that are designed to maintain “stability through change”
122 (Sterling & Eyer, 1988, p. 636). Multi-level allostatic regulation is supposedly accomplished through a
123 fine-grained network of feedforward and feedback mechanisms, thus affording a more flexible and co-ordinated
124 means of physiological control than the rather more primitive negative feedback loops typically attributed to
125 homeostatic regulation. One key advantage of this arrangement is that it enables anticipatory alterations
126 of physiological parameters *prior* to undergoing some perturbation (e.g., increasing blood pressure before
127 standing up from a chair, rather than correcting the hypotension induced by the postural change after the
128 fact). Under this allostatic regime, the body benefits from the brain’s capacity to learn from experience by
129 forecasting the organism’s physiological needs ahead of time. As such, allostasis represents a rather more
130 sophisticated system of internal regulation, one which minimises reliance upon the kind of error signaling
131 required to drive homeostatic correction.

132 Sterling and Eyer argued that the concept of homeostasis is fatally deficient, and ought thus to be “superseded”
133 by their notion of allostasis (1988, p. 646; see also Sterling, 2004, 2012; Sterling & Laughlin, 2015). However,
134 the validity of this assertion has been challenged by critics who argue that it turns on a fundamentally
135 mistaken construal of homeostatic regulation (Carpenter, 2004; Day, 2005). The source of this error is two-fold.
136 First, the careful nuance of Cannon’s (1929, 1939) definition of homeostasis is ignored in this account, giving
137 rise to the overly simplistic (and arguably misleading) impression that homeostasis is supposed to “clamp each
138 internal parameter at a ‘setpoint’” (Sterling, 2004, p. 17), except in response to emergency (i.e. potentially
139 life-threatening) situations. Second, Sterling and Eyer (1988) conflate the physiological variables that are the

140 target of homeostatic regulation with the control mechanisms tasked with the job of maintaining such variables
141 within acceptable bounds. The idea that physiological parameters such as blood pressure should fluctuate
142 significantly throughout the day does not constitute a counterexample to the homeostatic model; rather, these
143 fluctuations are in the service of homeostasis precisely insofar as they ensure that the vital constituents and
144 properties of the fluid matrix (e.g., blood pH, oxygen tension) remain suitable for cell functioning. On this
145 reading then, allostasis appears little more than “an unnecessary re-statement of the concept of homeostasis”
146 (Day, 2005, p. 1196).

147 3.2 Allostatic means for homeostatic ends

148 Since Sterling and Eyer’s (1988) introduction of the concept, less radical versions of allostasis have been
149 developed that seek to complement or extend the scope of homeostatic regulation, rather than reject it
150 wholesale. Early work by McEwen, Schulkin, and colleagues (McEwen & Stellar, 1993; Schulkin, McEwen,
151 & Gold, 1994) embraced allostasis as a promising framework for studying complex relations between stress,
152 behaviour, and chronic disease, and set about developing the concept of *allostatic load* to account for the
153 potentially deleterious consequences of resisting stressful stimuli. (Although an important dimension of the
154 allostatic framework developed by McEwen and others, notions relating to allostatic load/overload will not be
155 considered here — but see Peters, McEwen, & Friston, 2017).

156 As these theories matured, however, a more distinctive articulation of the base concept of allostasis started to
157 emerge. McEwen began to conceive of allostasis as “an essential component of maintaining homeostasis” (1998,
158 p. 37); where the latter is limited to “systems... that are truly essential for life” (2000b, p. 173). According
159 to this view, allostasis describes “the process for actively maintaining homeostasis” (McEwen, 2000b, p. 173);
160 or alternatively, “the means by which the body re-establishes homeostasis in the face of a challenge” (McEwen,
161 2000a, p. 25). In collaboration with Wingfield, McEwen’s notion of allostatic regulation was further expanded
162 to include setpoint adjustments in anticipation of cyclical changes across various temporal scales (McEwen &
163 Wingfield, 2003, 2010). This conceptual development highlighted the circadian modulation of homeostatic
164 parameters implicit in Sterling and Eyer’s (1988) paradigmatic example of allostatic change (i.e. the diurnal
165 variation of blood pressure upon which phasic modulations are superposed), while also extending the scope
166 of allostatic processes to incorporate broader aspects of animal well-being, reproduction, and ontogenetic
167 adaptation (e.g., seasonal variation in physiology and behaviour in preparation for hibernation or migration).

168 McEwen concedes that his construal of allostasis might seem almost identical to broader conceptions of
169 homeostasis, such as the view promulgated by Cannon (McEwen, 2000b, 2004; McEwen & Wingfield, 2003). He
170 insists, however, that the notion of the ‘steady state’ at the core of Cannonian homeostasis is inherently vague,
171 insofar as it fails to delineate vital (homeostatic) systems from those mechanisms which work to maintain
172 their stability. It is not entirely clear though why such a distinction ought to be desired; or indeed, if it is
173 even coherent in the context of McEwen’s broader framework. Dallman (2003) argued that so-called allostatic
174 systems do not manifest qualitatively distinct properties as compared to their homeostatic counterparts, on
175 the basis that such systems are responsible for a great deal of essential physiological and behavioural functions.
176 Indeed, it seems strange to claim that allostatic mechanisms are not equally essential to survival if such
177 adaptive systems play a crucial role in enabling the organism to flee (or better yet, entirely avoid) a deadly
178 predator, for example.

179 Although arguments of this sort might be blunted by a more charitable interpretation of the key idea
180 underlying McEwen’s proposed distinction (namely, that allostatic systems accommodate large fluctuations
181 precisely so that those physiological parameters which cannot tolerate such lability are not pushed beyond
182 their narrow limits; e.g., McEwen, 1998), it seems plausible that significant enough deviations in allostatic
183 systems should likewise prove fatal. Furthermore, cross-species analysis suggests that setpoint flexibility does
184 not constitute a reliable indicator of the relative importance of a given physiological parameter (see Boulos &
185 Rosenwasser, 2004). Nevertheless, McEwen and Wingfield’s (2003) thematisation of the multiple layers of
186 predictive regulation that unfold across the life cycle strikes us a valuable addition to the allostasis framework;
187 one which we take to be a genuine departure from traditional notions of homeostasis.

188 3.3 Two modes of sustained viability

189 Another account of allostatic regulation that seeks to integrate (rather than replace) conventional notions of
190 homeostatic control was put forth by Schulkin and colleagues (Power & Schulkin, 2012; Rosen & Schulkin, 2004;
191 Schulkin, 2003a, 2003b). Schulkin (2003a, 2003b) credits Cannon’s conception of homeostasis with greater
192 scope and sophistication than Sterling and Eyer (1988), while maintaining that some kind of supplementary
193 concept is necessary in order to capture the full gamut of regulatory strategies exhibited by complex organisms
194 (Power & Schulkin, 2012; Schulkin, 2003b). Schulkin expounds a version of allostasis in which brain-driven
195 regulatory mechanisms effect fluctuating physiological and psychological states in the absence of any clear
196 setpoint boundary. In particular, anticipatory (feedforward) hormonal processes are posited to play a crucial
197 role in the emergence of many appetitive, self-protective, and socially-orientated motivational drives (Schulkin,
198 2003b, 2011), as well as explaining the affective valence of emotional experiences that accompany such states
199 (Rosen & Schulkin, 2004). Schulkin and colleagues (Power & Schulkin, 2012; Rosen & Schulkin, 2004; Schulkin,
200 2003b, 2004) thus advocate a broad conception of biological regulation, one in which homeostasis and allostasis
201 constitute equally important (yet functionally opponent) mechanisms for maintaining the biological viability
202 of the internal milieu.

203 In some sense, we might regard Schulkin’s framework as a kind of synthesis of prior allostatic concepts. It
204 clearly inherits from Sterling and Eyer’s (1988) original conception of allostasis, retaining as it does an explicit
205 emphasis on the role of anticipatory physiological changes in efficient adaptation to environmental diversity.
206 It also takes up McEwen and Wingfield’s (2003) temporal expansion of the concept to account for longer-term
207 adaptive changes in response to various ecological and life cycle contexts (Schulkin, 2003b, 2004). However, by
208 balancing the homeostatic imperative to conserve stability with the allostatic impulse towards dynamic state
209 transition, Schulkin and colleagues thematise the deeper continuity uniting these apparently contradictory
210 concepts. At the heart of these regulatory principles is not so much the immediate influence they exert over
211 target physiological parameters (i.e. internal constancy vs. variability), but rather the overarching goal that
212 these mechanisms dually subserve: namely, the ongoing survival and reproductive success (i.e. evolutionary
213 fitness) of the organism (see also Power, 2004; Power & Schulkin, 2012; Schulkin, 2004).

214 This is not to say that the regulatory frameworks described by Sterling, McEwen, and others do not also
215 ground the emergence of allostatic mechanisms in the selective advantages they confer. The point here, rather,
216 is that sustained biological viability (rather than some other criterion such as internal stability) seems to us
217 the most plausible target towards which physiological and behavioural regulatory mechanisms are striving.
218 By these lights, there is no inherent contradiction between homeostatic and allostatic principles; they are
219 merely different routes to the same end.

220 4 Allostasis and interoceptive inference

221 The imperative to maintain biological viability over time is at the very core of the free energy principle
222 (Friston, 2010). Briefly, this principle begins with the observation that living entities must “maintain their
223 sensory states within physiological bounds,” and that they do so by engaging in actions which maintain the
224 integrity of their structural and dynamical organisation (Friston, 2013, pp. 1–2). This restates the cybernetic
225 insight that biological organisms resist the tendency towards disorder wrought by variable external conditions
226 (Ashby, 1947, 1962). The central element of the principle is that such self-preserving adaptation is achieved via
227 environmental exchanges enabled by the minimisation of free energy (or, under simplifying assumptions, the
228 long-term average of prediction error; Friston, 2010). Under most accounts invoking the free energy principle,
229 the process of maintaining the biological agent’s internal milieu within the limited subset of states conducive
230 to its ongoing existence is that of homeostasis (where homeostasis is understood more precisely in terms of
231 minimising the free energy of internal state trajectories in order to avoid surprise, i.e. minimise prediction
232 error; Friston, 2010).

233 The concept of allostasis started to infiltrate this picture in conjunction with remarks on the necessity of
234 maintaining homeostasis for survival (e.g., Friston, 2012; Friston et al., 2014; Moran, Symmonds, Dolan, &
235 Friston, 2014). Such comments typically invoked allostasis in the same breath as homeostasis, without offering

any indication as to how the two terms might refer to differentiated aspects of biological regulation. To our knowledge, the first attempt at characterising a substantive notion of allostasis as an independent mode of physiological regulation within the context of free energy minimisation was made by Gu and FitzGerald (2014). In the short period that has elapsed since, a number of investigators have imported allostasis into their own free energy-inspired accounts of interoceptive inference. Much like the original development of allostasis in the biomedical and ethological literatures however, the precise nature of allostatic control in these schemes has been elaborated in various ways. The time is ripe then to take stock of this nascent body of research; both to establish its continuities with – and departures from – pre-existing notions of allostasis, and to assess which interpretation(s) of the concept seem most promising from the free energy perspective.

For convenience, we divide these recent allostatic treatments of interoceptive inference into three broad classes: *behavioural*, *teleological*, and *diachronic* (see Figure 1). This division is not meant to be taken as absolute; indeed, these accounts share many similarities by dint of their common theoretical origins.

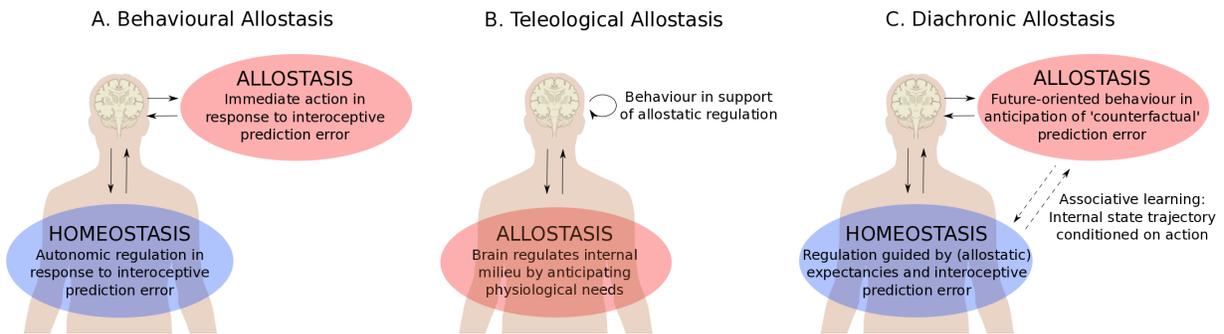


Figure 1: Schematic summary of the key conceptual distinctions between (A) behavioural, (B) teleological, and (C) diachronic accounts of allostatic regulation (see main text for details).

4.1 Behavioural allostasis

In their commentary on Seth’s (2013) theory of interoceptive inference, Gu and FitzGerald argue that the scope of predictive interoceptive processing should be extended beyond “homeostatic control of the internal milieu” to incorporate “allostatic actions on the external world” (2014, p. 269). At first, their position sounds isomorphic to that espoused by McEwen, insofar as allostasis is defined as “the process of achieving homeostasis” (Gu & FitzGerald, 2014, p. 269). It quickly becomes apparent, however, that Gu and FitzGerald (2014) conceive of homeostasis and allostasis in rather different terms. Here, homeostasis consists of autonomic reflexes that resist substantial fluctuations in the physiological conditions of the body (e.g., metabolising stored fat in response to declining blood glucose levels), while allostasis corresponds to the behavioural actions that the agent undertakes in order to ameliorate some internal perturbation (e.g., consuming food in response to glucose decline). Gu and FitzGerald (2014) thus advocate a framework in which homeostatic (brain–internal world) and allostatic (brain–external world) loops offer alternative pathways to the same ultimate goal; namely, that of keeping the organism within the subset of biophysical states most conducive to its survival (in other words, minimising the surprise or free energy indexed by interoceptive prediction error).

Gu and FitzGerald’s (2014) behaviour-orientated characterisation of allostasis is adopted and further elaborated by Seth (2015). Thematising the continuity between Ross Ashby’s pioneering work in cybernetics (Ashby, 1956, 1960) and the free energy principle, Seth (2015) seeks to map homeostasis and allostasis onto the ‘ultrastable’ scheme exemplified by Ashby’s (1960) *homeostat*. Briefly, this device consists of four modular subsystems which dynamically interact to influence one another’s essential variables. If these interactions fail to preserve essential variables within an acceptable range, a regulatory switch intervenes to randomly reconfigure the system’s behaviour. If the homeostat’s new organisation fails to stabilise essential variables within range, it will continue to transition through its repertoire of possible configurations until stability is

270 restored, or until the system disintegrates (see Cariani, 2009, for a more detailed explication of the homeostat's
271 functional architecture). Seth compares homeostasis to the first-order feedback loop constituted by the
272 dynamic interplay of each module's inputs and outputs, and allostasis to the second-order reorganisation of
273 these interactions (although allostatic behaviour constitutes a purposeful, rather than random, attempt to
274 transform system dynamics). On this account, then, allostatic behaviour functions to alter the organism's
275 relation to its environment when homeostatic compensation fails to maintain physiological parameters within
276 viable bounds.

277 Implementing these ideas within the context of free energy minimisation, Seth (2015) argues that interoceptive
278 prediction error can be minimised in one of three ways: (a) by adjusting model predictions in order to
279 better approximate the incoming sensory signal (equivalently, updating one's emotional state; i.e. perceptual
280 inference); (b) by enlisting autonomic reflexes to alter internal conditions such that they correspond with
281 the predicted internal state (i.e. active inference or first-order (homeostatic) control); or (c) by engaging in
282 goal-directed behaviour to act on the environment in such a way that brings about the predicted internal state
283 (i.e. second-order (allostatic) control). Here, then, allostasis is not only distinguished from the physiological
284 mechanisms responsible for regulating the internal milieu, but also construed as an alternative mode of
285 achieving organismic viability.

286 An interesting aspect of Seth's (2015) analysis is the claim that perception simply 'falls out' of the fundamental
287 necessity to achieve homeostatic control. It is not entirely clear whether Seth subscribes to a kind of
288 anti-realism which denies the veridicality of perceptual experience, or whether he wants to say that our
289 rich perceptual experiences of the world are merely an accidental consequence of (or a useful tool for) the
290 homeostatic imperative. In any case, Seth (2015) interprets the free energy principle in a way that assigns
291 primacy to interoception (over exteroceptive perception), insofar as interoceptive inference is regarded as
292 playing an instrumental role in steering the agent towards its homeostatic states. We shall encounter a similar
293 view in the following section, hence we postpone further consideration of its implications until later. Let us
294 first review the allostatic picture presented here.

295 Perhaps the most striking feature of these initial attempts to assimilate allostatic principles within a broader
296 predictive processing framework is the surprisingly *reactive* way in which allostasis is depicted. Rather than
297 presenting a paradigmatic example of *anticipatory* behaviour in the service of some homeostatic goal (e.g.,
298 consuming food *prior* to the decline of blood glucose concentration), Gu and FitzGerald (2014) portray
299 allostatic actions as a kind of external-world equivalent to the corrective autonomic responses orchestrated
300 by homeostatic control mechanisms. Seth (2015) likewise articulates what seems to be a distinctly reactive
301 form of allostatic regulation. Indeed, the ultrastable system to which Seth draws conceptual allusion is
302 entirely dependent on negative feedback responses to the perturbation of essential variables. Although the
303 second-order feedback loop is functionally analogous to McEwen's conception of allostasis as the means by
304 which homeostatic variables are stabilised, this arrangement lacks the capacity to anticipate and offset such
305 deviations before they occur (a vital feature of all allostatic frameworks reviewed in Section 3). Consequently,
306 the notion of allostasis invoked by these 'behavioural' accounts does not obviously pick out any process that
307 is distinctively predictive in nature.

308 Arguably, these fundamentally reactive models of allostasis derive from a partitioning of biological regulation
309 along the lines of internal/autonomic (i.e. homeostatic) and external/goal-directed (i.e. allostatic) responses.
310 Such a distinction is to our knowledge unprecedented in the allostasis literature, in as much as allostatic
311 mechanisms have always been conceived as a suite of actions traversing the physiological-behavioural continuum.
312 Here, notably, allostasis seems instead to refer exclusively to the behavioural strategies an agent can engage in
313 response to mounting interoceptive prediction error, rather than a process that participates in the proactive
314 avoidance of such surprising states. It is however unclear to us what substantive insights can be gleaned
315 from this sort of picture. Indeed, it is so obvious that organisms must interact with their environment in
316 order to satisfy their basic homeostatic needs (e.g., seeking out and drinking fluids to quench thirst) that such
317 behavioural repertoires are a well-established feature of homeostatic theory (see for e.g., Richter, 1942–1943).
318 Simply reassigning such activities under the rubric of allostasis is thus likely to revive the kind of criticism
319 engendered by earlier renditions of the theory (e.g., that allostasis is essentially redundant insofar as it
320 "represent[s] nothing that has not always been part of the ordinary conceptual basis of homeostatic control,"
321 Carpenter, 2004, p. 180).

322 On balance then, these interpretations risk diluting the concept of allostasis to the point where it constitutes
323 little more than a particular mode of homeostasis, a behavioural rearguard for occasions when autonomic
324 mechanisms prove insufficient. As such, these inherently reactive accounts do not seem to carry us far beyond
325 the insights availed by traditional homeostatic principles.

326 4.2 Teleological allostasis

327 The Embodied Predictive Interoception Coding model (EPIC; Barrett & Simmons, 2015) offers another
328 free-energy inspired account grounding interoceptive experience in the physiological status of homeostatic
329 variables. Initially, the authors of this model also defined allostasis in instrumental terms, describing it as the
330 “process of activating physiological systems (such as hormonal, autonomic, or immune systems) with the aim
331 of returning the body to homeostasis” (Barrett & Simmons, 2015, p. 422; Chanes & Barrett, 2016, p. 97).
332 However, allostasis assumes a more pivotal role in subsequent work by Barrett and colleagues (Barrett, 2017;
333 Barrett, Quigley, & Hamilton, 2016; Kleckner et al., 2017); the focus shifting from a reactive-mechanistic
334 interpretation (i.e. where allostatic processes are recruited in response to homeostatic perturbation, similar
335 to McEwen’s (1998, 2004) definition), to a broader perspective emphasizing its fundamentally predictive
336 nature (i.e. where bodily conditions are efficiently regulated through the co-ordinated allocation of energy
337 resources in anticipation of upcoming demands, similar to Sterling’s (2004, 2012; Sterling & Eyer, 1988)
338 position). In this view, allostasis (and its interoceptive consequents) is assigned primary importance in the
339 brain’s computational economy, such that the predictive models posited to underpin cognitive representation
340 are entirely subservient to the efficient satisfaction of the body’s physiological requirements (Barrett, 2017;
341 Barrett et al., 2016).

342 Barrett and colleagues’ more recent characterisations of allostasis as the primary design feature driving brain
343 evolution involve a number of important theoretical commitments. First, this expanded version of allostasis
344 apparently subsumes the homeostatic functions that allostatic processes had previously been supposed to
345 support. In eliminating all talk of homeostasis in favour of a more comprehensively encompassing model of
346 predictive regulation, Barrett and colleagues (Barrett, 2017; Barrett et al., 2016; Kleckner et al., 2017) align
347 themselves with Sterling’s (2004, 2012; Sterling & Laughlin, 2015) more radical allostatic agenda. It is not
348 immediately clear that this sort of move is necessary for Barrett and colleagues’ more recent formulations to
349 cohere, especially since their explicit concern with metabolic exchange and energy regulation would seem to
350 sit just as comfortably within McEwen and Wingfield’s (2003) framework.

351 The second notable claim deriving from this framework is that the brain’s computational architecture has
352 evolved in order to optimise allostatic regulation, rather than for purposes such as veridical perception or
353 reasoned action (Barrett, 2017; Barrett et al., 2016; Kleckner et al., 2017). This is to say that the brain’s
354 internal model (or ‘embodied simulation’) of the body and the ecological niche it inhabits is fundamentally
355 attuned to its physiological needs, such that only those features (i.e. statistical regularities) of the body–niche
356 dyad relevant to allostatic regulation are represented (Barrett, 2017). Furthermore, Barrett and colleagues
357 (Barrett, 2017; Barrett et al., 2016) propose that interoceptive representations emerge as a consequence of
358 allostatic processing, and that such affective sensations form a fundamental and pervasive feature of conscious
359 awareness. By implication, other sensory domains (and presumably, volitional motor activity) figure as
360 secondary or derivative phenomena, the metabolic costs of which are tolerated only insofar as they furnish
361 additional support to the brain’s primary allostatic–interoceptive axis (Barrett, 2017).

362 This picture is reminiscent of Seth’s (2015) argument for the primacy of interoceptive inference and physiological
363 regulation. It is not entirely clear whether Barrett and colleagues consider higher-level cognitive functions
364 to be useful adjuncts for maintaining allostasis, or whether they simply emerge as a byproduct of the
365 brain’s allostatic machinery. It is clear, however, that Barrett (2017) considers perceptual experience to be
366 fundamentally driven by allostatic and interoceptive processing, such that one’s subjective grasp of reality is
367 modeled according to one’s physiological needs. The upshot of this hypothesis is a constructivist account
368 in which allostasis functions as the author and arbiter of phenomenological experience, both insofar as the
369 imperative to optimise allostasis has carved out an evolutionary trajectory that has endowed the creature with
370 a particular cognitive architecture and set of sensory capacities, and insofar as the experiential possibilities

371 afforded by these devices are constrained and modulated in ways designed to realise this imperative in a given
372 context.

373 The brain’s evolution into a highly efficient allostatic machine, rather than (say) a rational decision-maker
374 or accurate perceiver of the world, does not necessarily preclude the possibility that it should realise these
375 additional properties also. Indeed, Seth, Barrett, and their colleagues may well agree that providing a creature
376 with the capacity to accurately model the hidden causes of its external perturbations would, over the long-run,
377 improve its capacity to maintain the viability of its internal milieu, as well as engage in other intrinsically
378 rewarding (and evolutionarily relevant) projects such as reproductive activity. As Barrett (2017) points out,
379 however, creatures need only be informed about hidden causes that are (potentially) relevant to their ongoing
380 allostatic needs and priorities (for instance, evolution has endowed humans with a sensorium that is indifferent
381 to infrared light stimulation). In this sense, then, these authors are correct to say that human perception
382 does not afford a ‘true’ picture of the world, at least insofar as the latter is construed as some complete
383 account of the totality of measurable phenomena. Indeed, it is hard to imagine how the kind of experience
384 that would obtain in the event that we really could perceive ‘everything’ could be of much use, as dense
385 with (predominantly irrelevant) information as it would be. There seems to be good *prima facie* reason then
386 to think that (exteroceptive) sensation has evolved precisely to the extent that it is *useful*, and adaptive
387 self-regulatory activity (maximising the likelihood of well-being and successful reproduction) would seem a
388 reasonable object *for which* it ought to be useful.

389 These considerations notwithstanding, we note a general doubt about the plausibility of any thoroughgoing
390 distinction between interoception and exteroception (independent of the specific role accorded to allostasis).
391 Although it is true that the free energy principle allows for the possibility of inherited model parameters, and
392 hence the newborn may come into the world equipped with certain expectations about the kinds of states
393 its various sensory receptors ought to entertain, it is unclear why information conveyed via interoceptive
394 afferents should be recognised by the brain as somehow different in kind to that received via exteroceptive (or
395 proprioceptive) channels. From a brain-centric perspective, the external world to be modeled is that which
396 lies beyond its neural projections, irrespective of whether this environment happens to be within or without
397 the boundary formed by the body (Friston, 2010). In this respect, then, there is no meaningful distinction
398 (for the brain) between the internal and external milieu; rather, there is only a Markov blanket (see Hohwy,
399 2017) separating a nervous system on the one side, and a hidden world of glucose molecules, blood vessels,
400 muscles, fires, kittens, and so on, on the other. Collapsing this distinction leaves no principled rationale for
401 privileging interoception over alternative forms sensory input; all channels furnish the brain with equally vital
402 information about the state of play beyond the Markov blanket, from which its models profit.

403 A further, rather abstract concern about the teleological perspective presented here relates subtly to the
404 conceptual role of the free energy principle. A key justification for the subordination of perceptual experience to
405 homeostatic or allostatic regulation is made by way of appeal to the free energy principle’s central concern with
406 the persistent integrity of self-organised systems in the face of uncertain environmental conditions. Although
407 we opened this section with a somewhat similar comment on the vital import of sustained biological viability
408 in Friston’s (2010) account, we urge caution in equating this with any so-called “fundamental imperative
409 towards homeostasis” (Seth, 2015, p. 3). Rather, it would be more precise to say that the free energy principle
410 captures something essential about the sorts of properties a biological system must possess in order to live
411 (e.g., Friston & Stephan, 2007). It might be better then to say something like the following: any biological
412 entity that consists of some form of sensorimotor interface through which it can enter into a dynamic exchange
413 of energy and information with its environment, and which comprises an internal organisation that enables it
414 to minimise the free energy that bounds the surprise on its sensory states, is likely to endure; and in so doing,
415 any such entity will thus *appear* to conform to the assumed imperative for the conservation of its biophysical
416 integrity via self-regulatory processes. In other words, if a free energy-minimising system exists, then it must
417 indeed do so in virtue of possessing the right kind of internal configuration, and having entered into the right
418 kind of circular-causal relationship with its environment, to be able to model the causes of its sensory states
419 and engage in (what will look like) adaptive, self-regulatory activity (cf. Allen & Friston, 2016). As such, the
420 apparent imperative towards self-regulatory behaviour (be it homeostatic, allostatic, or whatever) seems to
421 fall out of the ongoing minimisation of free energy, in much the same way as the apparent teleological force
422 driving evolutionary ‘design’ emerges as a consequence of the intricate, non-teleological dynamics driving
423 natural selection.

4.3 Diachronic allostasis

We turn finally to two remaining inferential formulations of allostasis, which we refer to as ‘diachronic’ on account of the important implications they have for regulatory activity over various timescales.

Pezzulo, Rigoli, and Friston (2015) set out to explain how prospective and goal-directed (i.e. allostatic) forms of control might have evolved from more primitive mechanisms subserving homeostatic regulation. Here, homeostasis is construed along control-theoretic/cybernetic lines of negative feedback and setpoint control, where autonomic and behavioural reflexes are enlisted to correct deviations in physiological variables (see also Pezzulo, 2013; Seth, 2013). By contrast, allostasis refers to the flexible, context-specific engagement of complex, adaptive behavioural repertoires for the purposes of achieving some future outcome. Like the behavioural accounts surveyed above, then, homeostasis and allostasis are equated with ‘direct’ and ‘indirect’ modes of eliminating interoceptive prediction error, respectively. Note however that the distinction here is more nuanced, insofar as homeostatic responses extend to the innate behavioural sets (e.g., approach/avoidance behaviour) that equip animals to survive in the absence of associative learning.

If complex behavioural policies are to offer an effective means of controlling the physiological conditions of the body, it is essential that they deliver the right kinds of state transitions at the right time. This requirement is inherently challenging, however, since the consequences of a particular policy are necessarily realised some time after those conditions that triggered its initiation. Such delays are nontrivial in the context of homeostatic control, where a process causing physiological conditions to deteriorate may precipitate catastrophic damage if not promptly addressed. Pezzulo and colleagues’ (2015) solution to this problem leverages the free energy minimising agent’s ability to acquire sophisticated internal models of the hidden environmental causes of its sensory states. Specifically, they argue that such generative interoceptive models enable such agents to predict the temporal evolution of interoceptive state trajectories (i.e. how interoceptive signals are likely to change over time), and encode how these trajectories correlate with sensorimotor events in the external world (cf. Friston, Rosch, Parr, Price, & Bowman, 2017). In virtue of the higher-level integration of sensory information converging from interoceptive, exteroceptive, and proprioceptive streams, the agent is thus able to acquire a rich understanding of how behavioural activities come to influence interoceptive states across various contexts. By linking interoceptive prediction errors and their suppression through active inference (i.e. engagement of allostatic behaviour) via such associative learning processes, Pezzulo and colleagues (2015) provide a compelling explanation of (a) how the allostatic anticipation of future homeostatic needs might systematically arise, and (b) why allostatic behavioural policies should be endorsed despite potentially lengthy delays in their homeostatic payoff.

On this construal, allostatic processing turns out to be fundamentally *counterfactual* in nature. Higher (or deeper) hierarchical representations map the relation between increasingly distal outcome states and the behavioural policies that would lead towards their accomplishment. This account thus renders a smooth continuum of adaptive action selection, ranging from the primitive drives that work, for instance, to sate appetite via exploitation of the immediate environment, to the complex deliberative activities serving various motivations extending well beyond the basic requirements of the internal milieu (see also Pezzulo, 2017). Indeed, Pezzulo and colleagues (2015) observe that the capacity to learn the counterfactual relations that enable the agent to engage in prospective planning, and to choose amongst various available policies, confers an unparalleled degree of autonomy from the exigencies of the homeostatic imperative. Thus, in much the same way as Bernard and Cannon recognised how the capacity to maintain the stability of the internal milieu granted complex biological systems a remarkable degree of autonomy from the caprices of their external environments, allostasis under this scheme extends such freedom even further. Capable of holding the immediate demands of homeostasis in abeyance to some supraordinate desired (i.e. unsurprising and attracting) state, the autonomous horizon of the allostatic organism expands beyond the conditions of the present into a predictable (albeit uncertain) future.

Finally, Stephan and colleagues (2016) propose a formalised Bayesian implementation of hierarchical allostatic control that likewise operates across various temporal grains. Allostasis is defined here as the mode of active inference which performs “anticipatory homeostatic control” (Stephan et al., 2016, p. 5). This is achieved via the modulation of prior beliefs concerning the expected state trajectory of a given homeostatic setpoint. Expectancies about setpoint values are construed in terms of a probability distribution, such that beliefs

475 propagated from higher-level circuits influence both the mean value of the controlled variable, and its associated
476 variability (or precision). In other words, the traditional notion of a homeostatic negative feedback loop is
477 situated at the lowest level of the processing hierarchy, with its target setpoint (i.e. the expected physiological
478 state) conditioned by information from higher (allostatic) circuits. These higher (or deeper) hierarchical levels
479 are posited to model increasingly broader, domain-general representations of the present state of the body
480 and its environment, as well as predictions about changes in those states. Consequently, this account of
481 allostatic regulation incorporates an important temporal dimension, where higher-level generative models are
482 able to inform and update lower-level homeostatic control mechanisms in accordance with predictions about
483 upcoming state transitions.

484 Stephan and colleagues (2016) set out their model of allostatically regulated homeostatic reflexes in accordance
485 with the basic computational architecture assumed by the free energy principle. Homeostatic control thus
486 depends on both the perception of salient features within the internal and external milieu (comprising
487 both physical and social dynamics), and selection of appropriate actions designed to prevent dangerous
488 (i.e. surprising) deviations of physiological parameters. Inference is divided into interoceptive and exteroceptive
489 sensory processing. Prediction concerns how internal and external states will evolve over time, as well as
490 the degree to which possible actions will maintain internal states within the bounds of a given homeostatic
491 setpoint over time. In other words, allostatic prior beliefs set expectations about the space of bodily states
492 that the organism ought to inhabit (i.e. that delimited set of attracting states which engender low entropy),
493 which homeostatic systems subsequently attempt to realise. Importantly, this generic active inference scheme
494 is extended beyond the context of low-level homeostatic reflexes to encompass the higher-level implementation
495 of flexible behavioural policies designed to avoid homeostatic surprise (in a similar vein to Pezzulo et al.,
496 2015).

497 Stephan and colleagues (2016) present the first mathematically concrete account of allostatic control within
498 the context of free energy minimisation. Although more work needs to be done to flesh out this formal scheme
499 with respect to the complex dynamics involved in the integrated regulation of complex physiological systems, it
500 provides a plausible theoretical framework for explaining a number of core allostatic phenomena. The notion of
501 a Bayesian reflex arc whose setpoint is adaptively defined and constrained by higher-order (allostatic) dynamics
502 provides an elegant explanation of setpoint variability; one that seems equally capable of incorporating other
503 (i.e. non-allostatic) accounts of flexible setpoint control (e.g., Cabanac, 2006). Embedding this arc within
504 a hierarchical architecture also provides a principled mechanistic explanation of how certain higher-order
505 parameters might be prioritised at the expense of less-urgent homeostatic needs, and how maladaptive
506 psychological states might be entrained by persistent interoceptive prediction error. This perspective thus
507 offers a deeply unifying picture of homeostatic and allostatic control as a dynamic coupling or closed loop, with
508 lower-level homeostatic inferences and higher-level allostatic predictions reciprocally informing and modulating
509 one another as the joint conditions of the agent–niche dyad evolve.

510 Aside from some minor technicalities concerning the precise definitional boundaries of homeostatic and
511 allostatic control, we consider the two diachronic theories outlined above to be broadly compatible and
512 complementary. We prefer Stephan and colleagues’ (2016) Bayesian reflex formulation insofar as it expands
513 the scope of allostatic control to the modulation of internal conditions (rather than limiting it to the domain
514 of external, goal-directed behaviour). This perspective is more consistent with the historical development
515 of the allostatic framework (as examined in section 3), all prominent versions of which assume allostasis
516 to consist of a repertoire of mechanisms that include the capacity to influence internal conditions directly
517 by harnessing physiological effectors. Happily, the Bayesian reflex account invokes a principled distinction
518 between homeostatic and allostatic control which succeeds in preserving the key functional characteristics
519 of both modes of regulation (i.e. it neither collapses one concept into the other, nor relies on arbitrary or
520 vague criteria for distinguishing their respective remits), while still allowing for the kind of higher-level,
521 temporally-extended allostatic behaviour articulated by Pezzulo and colleagues (2015). Furthermore, we find
522 Stephan and colleagues’ (2016) framework a potentially more useful starting point for future inquiry into the
523 general nature of biological regulation, insofar as it affords the basic computational elements for scaffolding
524 the emergence of less flexible, non-counterfactual forms of allostatic regulation (e.g., circadian, circannual,
525 and ontogenetic). By integrating the complementary perspectives provided by both diachronic theories, we
526 arrive at a nuanced and fecund account of self-regulation that accommodates multiple scales of biological and
527 cognitive complexity.

5 The future of the history of allostasis

Our review of the origins of allostasis, and analysis of its recent uptake in theories of interoceptive inference, might give the impression that the concept is as protean as the phenomena which inspired its coinage. This may be a consequence of zealous category splitting on our part, motivated by our intent to differentiate meaningful distinctions amongst a cluster of intersecting (and not entirely consistent) theoretical perspectives. However, the various interpretations and treatments allostasis has received over the years have tended to congeal around a more or less stable core of organising principles (e.g., Schulkin, 2004). Mature versions of Sterling’s (2004, 2012) and McEwen’s (e.g., 2004, 2007) frameworks have understandably evolved into more expansive and nuanced iterations of their progenitors, benefitting from empirical advances and critical discussion. These influential accounts have thus reached a point of quasi-consensus, in as much as they lack the diversity of a genuine pluralism, but fail to converge fully on a coherent, unified account of what allostasis is or does. This leaves us in the somewhat precarious position of possessing a theoretical construct that appears well established and valid, but comprises a heterogeneous and not entirely coherent set of commitments. Part of the motivation of this chapter was therefore to highlight this situation, given that free energy theorists have started helping themselves to aspects of the allostasis construct without necessarily being explicit about which particular interpretation(s) of it they wish to endorse.

A useful illustration might be drawn from our distinction between what we dubbed the teleological and the diachronic interpretations of allostasis. Indeed, those familiar with the former might protest that it too invokes a hierarchical architecture which, much like the diachronic accounts, also admits of higher generative models encoding predictions extending across increasingly extended temporal windows. As such, it might seem somewhat disingenuous to exclude this model from our favoured diachronic category. Our point, however, is that these frameworks are founded on rather different understandings of allostasis, giving rise to subtle (but deep) conceptual disagreements. The teleological perspective considers exteroception as secondary to interoception, which in turn emerges as a consequence of allostasis. The diachronic perspectives, on the other hand, hold each domain of sensory information processing in equal standing; interoception, exteroception, and proprioception are blended together at a suitably high level of hierarchical modeling and without any indication that any one stream is more fundamental than the others.

We urge care about which aspects of allostatic theory are imported into predictive model-based accounts of interoception. Indeed, it is notable that none of the interoceptive inference theories examined above acknowledge the accusations of redundancy, inconsistency, and ambiguity that have been levelled against the allostasis literature, even after some of these authors had substantially revised their own application of the concept. Ignoring such issues not only belies the contested nature of allostatic control, it has the potential to propagate further confusion as disparate elements of the construct are selectively sampled and fused together.

If the future of allostasis is to disclose meaningful theoretical insights concerning the predictive processes that support biological regulation and interoceptive inference, then the next phase of its conceptual development requires us to work out a clear and precise understanding of its core principles and entailments. We have tried to clarify some of the confusion that has plagued the allostasis literature since its inception, and argued in favour of an inclusive view that reconciles homeostasis and allostasis as complementary strategies for sustaining biological viability. We have also attempted to shed light on some of the idiosyncratic ways in which allostasis has been deployed in recent characterisations of interoceptive inference, and suggest that future progress in this line of research will be hindered if these conceptual inconsistencies are not subject to critical scrutiny.

Acknowledgements

We thank the anonymous reviewer for their suggested improvements to an earlier version of this chapter. AWC is supported by an Australian Government Research Training Program (RTP) scholarship. JH is supported by The Australian Research Council DP160102770 and by the Research School Bochum and the Center for Mind, Brain and Cognitive Evolution, Ruhr-University Bochum.

References

- 575
576 Adolph, E. F. (1961). Early concepts of physiological regulations. *Physiological Reviews*, *41*(4), 737–770.
- 577 Allen, M., & Friston, K. J. (2016). From cognitivism to autopoiesis: Towards a computational framework for
578 the embodied mind. *Synthese*.
- 579 Ashby, W. R. (1947). The nervous system as physical machine: With special reference to the origin of adaptive
580 behavior. *Mind*, *56*(221), 44–59.
- 581 Ashby, W. R. (1956). *An introduction to cybernetics*. London: Chapman & Hall Ltd.
- 582 Ashby, W. R. (1960). *Design for a brain: The origin of adaptive behaviour* (2nd ed.). London: Chapman &
583 Hall Ltd.
- 584 Ashby, W. R. (1962). Principles of the self-organizing system. In H. Von Foerster & G. W. Zopf Jr. (Eds.),
585 *Principles of self-organization: Transactions of the university of illinois symposium* (pp. 255–278). London:
586 Permagon Press.
- 587 Barrett, L. F. (2017). The theory of constructed emotion: An active inference account of interoception and
588 categorization. *Social Cognitive & Affective Neuroscience*, *12*(1), 1–23.
- 589 Barrett, L. F., & Simmons, W. K. (2015). Interoceptive predictions in the brain. *Nature Reviews Neuroscience*,
590 *16*(7), 419–429.
- 591 Barrett, L. F., Quigley, K. S., & Hamilton, P. (2016). An active inference theory of allostasis and interoception
592 in depression. *Philosophical Transactions of the Royal Society B*, *371*(20160011), 1–17.
- 593 Boulos, Z., & Rosenwasser, A. M. (2004). A chronobiological perspective on allostasis and its application
594 to shift work. In J. Schulkin (Ed.), *Allostasis, homeostasis, and the costs of physiological adaptation* (pp.
595 228–301). Cambridge: Cambridge University Press.
- 596 Cabanac, M. (2006). Adjustable set point: To honor harold t. hammel. *Journal of Applied Physiology*, *100*(4),
597 1338–1346.
- 598 Cannon, W. B. (1929). Organization for physiological homeostasis. *Physiological Reviews*, *9*(3), 399–431.
- 599 Cannon, W. B. (1939). *The wisdom of the body: Revised and enlarged edition*. New York: W. W. Norton &
600 Company, Inc.
- 601 Cariani, P. A. (2009). The homeostat as embodiment of adaptive control. *International Journal of General*
602 *Systems*, *38*(2), 139–154.
- 603 Carpenter, R. H. S. (2004). Homeostasis: A plea for a unified approach. *Advances in Physiology Education*,
604 *28*, 180–187.
- 605 Chanes, L., & Barrett, L. F. (2016). Redefining the role of limbic areas in cortical processing. *Trends in*
606 *Cognitive Sciences*, *20*(2), 96–106.
- 607 Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science.
608 *Behavioral Brain Sciences*, *36*(3), 181–253.
- 609 Clark, A. (2016). *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford: Oxford University
610 Press.
- 611 Cooper, S. J. (2008). From claude bernard to walter cannon. emergence of the concept of homeostasis.
612 *Appetite*, *51*(3), 419–427.
- 613 Dallman, M. F. (2003). Stress by any other name.? *Hormones & Behavior*, *43*(1), 18–20.
- 614 Day, T. A. (2005). Defining stress as a prelude to mapping its neurocircuitry: No help from allostasis. *Progress*
615 *in Neuro-Psychopharmacology & Biological Psychiatry*, *29*(8), 1195–1200.
- 616 Friston, K. J. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, *11*(2),

- 617 127–138.
- 618 Friston, K. J. (2012). Embodied inference and spatial cognition. *Cognitive Processing*, 13(S1), 171–177.
- 619 Friston, K. J. (2013). Life as we know it. *Journal of the Royal Society Interface*, 10(86), 20130475.
- 620 Friston, K. J., & Stephan, K. E. (2007). Free-energy and the brain. *Synthese*, 159(3), 417–458.
- 621 Friston, K. J., Rosch, R., Parr, T., Price, C., & Bowman, H. (2017). Deep temporal models and active
622 inference. *Neuroscience & Biobehavioral Reviews*, 77, 388–402.
- 623 Friston, K. J., Schwartenbeck, P., FitzGerald, T., Moutoussis, M., Behrens, T., & Dolan, R. J. (2014).
624 The anatomy of choice: Dopamine and decision-making. *Philosophical Transactions of the Royal Society B*,
625 369(20130481), 1–12.
- 626 Gu, X., & FitzGerald, T. H. B. (2014). Interoceptive inference: Homeostasis and decision-making. *Trends in*
627 *Cognitive Sciences*, 18(6), 269–270.
- 628 Hohwy, J. (2013). *The predictive mind*. Oxford: Oxford University Press.
- 629 Hohwy, J. (2016). The self-evidencing brain. *Noûs*, 50(2), 259–285.
- 630 Hohwy, J. (2017). How to entrain your evil demon. In T. Metzinger & W. Wiese (Eds.), *Philosophy and*
631 *predictive processing* (pp. 1–15). Frankfurt am Main: MIND Group.
- 632 Khalsa, S. S., Adolphs, R., Cameron, O. G., Critchley, H. D., Davenport, J. S., Feinstein, J. S., . . . Paulus,
633 M. P. (in press). Interoception and mental health: A roadmap. *Biological Psychiatry: Cognitive Neuroscience*
634 *& Neuroimaging*.
- 635 Kleckner, I. R., Zhang, J., Touroutoglou, A., Chanes, L., Xia, C., Simmons, W. K., . . . Barrett, L. F. (2017).
636 Evidence for a large-scale brain system supporting allostasis and interoception in humans. *Nature Human*
637 *Behaviour*, 1(0069), 1–14.
- 638 Lowe, R., Almér, A., & Dodig-Crnkovic, G. (2017). Predictive regulation in affective and adaptive behaviour:
639 An allostatic-cybernetics perspective. In J. Vallverdú, M. Mazzara, M. Talanov, S. Distefano, & R. Lowe
640 (Eds.), *Advanced research on biologically inspired cognitive architectures*. Hershey, PA: IGI Global.
- 641 McEwen, B. S. (1998). Stress, adaptation, and disease: Allostasis and allostatic load. *Annals of the New York*
642 *Academy of Sciences*, 840, 33–44.
- 643 McEwen, B. S. (2000a). Protective and damaging effects of stress mediators: Central role of the brain. In
644 E. A. Mayer & C. B. Saper (Eds.), *Progress in brain research* (Vol. 122, pp. 25–34). Amsterdam: Elsevier
645 Science BV.
- 646 McEwen, B. S. (2000b). The neurobiology of stress: From serendipity to clinical relevance. *Brain Research*,
647 886(1-2), 172–189.
- 648 McEwen, B. S. (2004). Protective and damaging effects of mediators of stress: Allostasis and allostatic load.
649 In J. Schulkin (Ed.), *Allostasis, homeostasis, and the costs of physiological adaptation* (pp. 65–98). Cambridge,
650 MA: MIT Press.
- 651 McEwen, B. S. (2007). Physiology and neurobiology of stress and adaptation: Central role of the brain.
652 *Physiological Reviews*, 87(3), 873–904.
- 653 McEwen, B. S., & Stellar, E. (1993). Stress and the individual: Mechanisms leading to disease. *Archives of*
654 *Internal Medicine*, 153(18), 2093–2101.
- 655 McEwen, B. S., & Wingfield, J. C. (2003). The concept of allostasis in biology and biomedicine. *Hormones &*
656 *Behavior*, 43(1), 2–15.
- 657 McEwen, B. S., & Wingfield, J. C. (2010). What is in a name? Integrating homeostasis, allostasis and stress.
658 *Hormones & Behavior*, 57(2), 105–111.
- 659 Michael, J., & McFarland, J. (2011). The core principles (“big ideas”) of physiology: Results of faculty

660 surveys. *Advances in Physiology Education*, 35(4), 336–341.

661 Michael, J., Modell, H., McFarland, J., & Cliff, W. (2009). The “core principles” of physiology: What should
662 students understand? *Advances in Physiology Education*, 33(1), 10–16.

663 Modell, H., Cliff, W., Michael, J., McFarland, J., Wenderoth, M. P., & Wright, A. (2015). A physiologist’s
664 view of homeostasis. *Advances in Physiology Education*, 39(4), 259–266.

665 Moran, R. J., Symmonds, M., Dolan, R. J., & Friston, K. J. (2014). The brain ages optimally to model its
666 environment: Evidence from sensory learning over the adult lifespan. *PLoS Computational Biology*, 10(1),
667 e1003422.

668 Peters, A., McEwen, B. S., & Friston, K. J. (2017). Uncertainty and stress: Why it causes diseases and how it
669 is mastered by the brain. *Progress in Neurobiology*, 156, 164–188.

670 Pezzulo, G. (2013). Why do you fear the bogeyman? An embodied predictive coding model of perceptual
671 inference. *Cognitive, Affective, & Behavioral Neuroscience*, 14(3), 902–911.

672 Pezzulo, G. (2017). Tracing the roots of cognition in predictive processing. In T. Metzinger & W. Wiese
673 (Eds.), *Philosophy and predictive processing* (pp. 1–20). Frankfurt am Main: MIND Group.

674 Pezzulo, G., Rigoli, F., & Friston, K. J. (2015). Active inference, homeostatic regulation and adaptive
675 behavioural control. *Progress in Neurobiology*, 134, 17–35.

676 Power, M. L. (2004). Commentary: Viability as opposed to stability: An evolutionary perspective on
677 physiological regulation. In J. Schulkin (Ed.), *Allostasis, homeostasis, and the costs of physiological adaptation*
678 (pp. 343–364). Cambridge: Cambridge University Press.

679 Power, M. L., & Schulkin, J. (2012). Maternal obesity, metabolic disease, and allostatic load. *Physiology &*
680 *Behavior*, 106(1), 22–28.

681 Richter, C. P. (1942–1943). Total self regulatory functions in animals and human beings. *Harvey Lecture*
682 *Series*, 38, 63–103.

683 Rosen, J. B., & Schulkin, J. (2004). Adaptive fear, allostasis, and the pathology of anxiety and depression. In
684 J. Schulkin (Ed.), *Allostasis, homeostasis, and the costs of physiological adaptation* (pp. 164–227). Cambridge:
685 Cambridge University Press.

686 Schulkin, J. (2003a). Allostasis: A neural behavioral perspective. *Hormones & Behavior*, 43(1), 21–27.

687 Schulkin, J. (2003b). *Rethinking homeostasis: Allostatic regulation in physiology and pathophysiology*. Cam-
688 bridge, MA: MIT Press.

689 Schulkin, J. (2004). Introduction. In J. Schulkin (Ed.), *Allostasis, homeostasis, and the costs of physiological*
690 *adaptation* (pp. 1–16). Cambridge: Cambridge University Press.

691 Schulkin, J. (2011). Social allostasis: Anticipatory regulation of the internal milieu. *Frontiers in Evolutionary*
692 *Neuroscience*, 2(111), 1–15.

693 Schulkin, J., McEwen, B. S., & Gold, P. W. (1994). Allostasis, amygdala, and anticipatory angst. *Neuroscience*
694 *& Biobehavioral Reviews*, 18(3), 385–396.

695 Seth, A. K. (2013). Interoceptive inference, emotion, and the embodied self. *Trends in Cognitive Sciences*,
696 17(11), 565–573.

697 Seth, A. K. (2015). The cybernetic bayesian brain: From interoceptive inference to sensorimotor contingencies.
698 In T. Metzinger & J. M. Windt (Eds.), *Open mind* (pp. 1–24). Frankfurt am Main: MIND Group.

699 Seth, A. K., & Friston, K. J. (2016). Active interoceptive inference and the emotional brain. *Philosophical*
700 *Transactions of the Royal Society B*, 371(1708), 1–10.

701 Smith, R., Thayer, J. F., Khalsa, S. S., & Lane, R. D. (2017). The hierarchical basis of neurovisceral
702 integration. *Neuroscience & Biobehavioral Reviews*, 75, 274–296.

703 Stephan, K. E., Manjaly, Z. M., Mathys, C. D., Weber, L. A. E., Paliwal, S., Gard, T., . . . Petzschner, F.

- 704 H. (2016). Allostatic self-efficacy: A metacognitive theory of dyshomeostasis-induced fatigue and depression.
705 *Frontiers in Human Neuroscience*, *10*(550), 1–27.
- 706 Sterling, P. (2004). Principles of allostasis: Optimal design, predictive regulation, pathophysiology and
707 rational therapeutics. In J. Schulkin (Ed.), *Allostasis, homeostasis, and the costs of physiological adaptation*
708 (pp. 17–64). Cambridge: Cambridge University Press.
- 709 Sterling, P. (2012). Allostasis: A model of predictive regulation. *Physiology & Behavior*, *106*(1), 5–15.
- 710 Sterling, P., & Eyer, J. (1988). Allostasis: A new paradigm to explain arousal pathology. In S. Fisher & J.
711 Reason (Eds.), *Handbook of life stress, cognition and health* (pp. 629–649). John Wiley & Sons Ltd.
- 712 Sterling, P., & Laughlin, S. (2015). *Principles of neural design*. Cambridge, MA: MIT Press.
- 713 Woods, S. C., & Ramsay, D. S. (2007). Homeostasis: Beyond curt richter. *Appetite*, *49*(2), 388–398.
- 714 Yates, F. E. (1996). Homeostasis. In J. E. Birren (Ed.), *Encyclopedia of gerontology: Age, aging, and the*
715 *aged* (Vol. 1, pp. 679–686). San Diego, CA: Academic Press.