Interoception as modeling, allostasis as control

Eli Sennesh, Jordan Theriault, Dana Brooks, Jan-Willem van de Meent,
Lisa Feldman Barrett, Karen S. Quigley

Northeastern University, Boston, Massachusetts, United States

Abstract

The brain regulates the body by anticipating its needs and attempting to meet them before they arise – a process called allostasis. Allostasis requires a model of the changing sensory conditions within the body, a process called interoception. In this paper, we examine how interoception may provide performance feedback for allostasis. We suggest studying allostasis in terms of control theory, reviewing control theory’s applications to related issues in physiology, motor control, and decision making. We synthesize these by relating them to the important properties of allostatic regulation as a control problem. We then sketch a novel formalism for how the brain might perform allostatic control of the viscera by analogy to skeletomotor control, including a mathematical view on how interoception acts as performance feedback for allostasis. Finally, we suggest ways to test implications of our hypotheses.

Keywords: Interoception, Allostasis, Predictive processing

1 Contents

2	1 Introduction: the functions of the brain in the body	3
3	2 Interoception: modeling the body, estimating its efficiency	6
4	2.1 Regulated resources and controlled processes in physiology . .	6
5	2.2 Predicting and modeling the ranges of regulated and controlled	
6	processes	8
7	2.3 Modeling the viable ranges of multiple controlled processes to	
8	support multi-system regulation and coordinated action	14
9	2.4 Viable ranges and capacities could obviate a modular “reward	
10	system”	17

11	2.5 Summary	19
12	3 Control theory: A unifying lens for physiology, motor control, and decision making	20
13		
14	3.1 Control theory for physiology: A reliable body built from un-	
15	reliable parts	20
16	3.2 Moving the body: The referent control hypothesis	24
17	3.3 Making decisions: constructing future reference trajectories . .	27
18	3.4 Allostatic control: Motivating movements with an interocep-	
19	tive model	30
20	3.5 Summary	33
21	4 Allostasis as trajectory-tracking stochastic optimal control	35
22	4.1 Transforming capacity curves into objective functions	35
23	4.2 Optimal foraging theory suggests a functional form for allo-	
24	static control	40
25	4.3 Feedforward control with generative action concepts	41
26	4.4 Feedback control with generative action concepts	45
27	4.5 Summary	50
28	5 Discussion: Interoception stabilizes action and constructs al-	
29	lostatic references	51
30	5.1 Viewing the brain as an allostatic optimal controller	51
31	5.2 The body and brain through the lens of control theory	54
32	5.3 Interoception and capacity curves	55
33	5.4 Conclusions	58
34	A Glossary	83
35	B How ergodicity interacts with control	86
36	B.1 Ergodicity: whether uncertainty and noise make a difference	
37	over time	86
38	B.2 Formulating perception and control ergodically leads to the	
39	Free Energy Principle (FEP)	88
40	C APIC model derivation details	92
41	C.1 A hierarchical generative model with reference distributions	92
42	C.2 Full derivation of optimal value function and transition	
43	dynamics	94

44 1. Introduction: the functions of the brain in the body

45 Imagine that you are learning to play dodgeball as a beginner. You
46 stand with the other players, divided into two teams, and when the game
47 begins you need to pick up a large inflated ball from a pile in the middle
48 and hit a member of the other team with it. As you run, throw, dodge,
49 catch, and reach, your muscle cells require metabolic fuel in the form of
50 molecules such as oxygen and glucose, which must be conveyed to those
51 muscle cells via the blood. Your vascular system must deliver and distribute
52 blood with speed, bringing nutrients and removing metabolites. Despite
53 rapid muscle movements generating waste heat, your body temperature must
54 remain within a narrow, viable range. As blood circulates more quickly
55 throughout your body, your lungs must also increase the rate with which
56 they breathe oxygen in and carbon dioxide out.

57 Playing a simple game of dodgeball, then, requires your brain to contin-
58 ually coordinate the systems of your body. At the same time, your body
59 sends sensory information about internal events up the spinal cord and va-
60 gus nerve to the brain. It is standard practice in neuroscience to distinguish
61 the brain’s “physiological sense of the condition of the body” ([interoception](#)
62 [\[1, 2, 3\]](#)) from the collection of sensory modalities that inform the brain about
63 the world outside the body ([exteroception](#)).

64 Interoception includes, but is not limited to, the brain’s modeling of the
65 sensory signals from innervated visceral organs. Nociception, temperature,
66 and C-tactile afferent-mediated (affective) touch on the skin are also consid-
67 ered interoceptive modalities, by virtue of their conveyance of sensory inputs
68 to the brain via unmyelinated or lightly myelinated ascending fibers in the
69 lamina 1 spinothalamic tract [\[1, 4, 2\]](#). A broad view of interoception also
70 includes modeling chemosensation from within the body’s interior, such as
71 changes in the endocrine system [\[5\]](#), changes in the immune system [\[6\]](#), and
72 changes in the digestive system and gut [\[7, 8\]](#). For simplicity’s sake, however,
73 this paper will treat all these systems as “visceral”.

74 [Viscerosensory signaling](#) (i.e., the ascending signals from the sensory sur-
75 faces inside the body and the skin) informs the brain of the state of the
76 body in an ever-changing and only partly predictable world. Since sensory
77 signals themselves are ambiguous and noisy, this poses an inverse problem
78 for the brain, one of inferring causes (the state of the body) from effects (the
79 ascending viscerosensory signals). The brain solves this problem by means
80 of an [internal model](#) [\[9\]](#). Psychologists refer to the internal model, including

interoception, by many terms, including memory [10], belief [11], perceptual inference [12], unconscious inference [13], embodied simulation [14], concepts and categories [15], controlled hallucination [16], and prediction [17, 18]. Regardless of what it is called, the brain is hypothesized to construct a dynamic model of its body in the world [19, 20]. In this paper we will use the terms prediction, simulation, and concept.

The process of building and refining an internal model based on viscerosensory signals does not, in and of itself, accomplish the brain’s most basic task. This task is to maximize the energy efficiency of bodily functions, to “anticipate changing needs, evaluate priorities, and prepare the organism to satisfy them *before* they lead to errors” (page 4, Sterling [21]), a process called [allostasis](#) (for further discussion on allostasis, see Sterling and Laughlin [22], Schulkin and Sterling [23]). Concurrent evolutionary [24] and neuroanatomical [19, 25, 15] evidence suggests that exteroceptive sensory signals, and the internal models anticipating them, contextualize and support motor control [9]. In a similar way, viscerosensory signals provide online feedback for allostasis, and interoceptive internal modeling subserves allostatic visceromotor control [19, 25, 26, 15]. Many lines of evidence suggest the same conclusion: the brain is predictively regulating the body, which is a problem of motor control rather than of perceiving the world. It is a problem of regulating the body along a desired trajectory to achieve efficiency.

Existing formal models of interoception and body regulation (such as those reviewed by Hulme et al. [27] and Petzschner et al. [28], as well as recent recent works such as Unal et al. [29]) have either formulated allostasis as a prospective decision-making problem (without considering how those decisions are enacted) or as a motor control problem (without considering where motor commands come from). Additionally, rather than treat metabolic efficiency as the objective, they discuss homeostasis, the regulation of bodily variables to fixed [set points](#) with fixed tolerances for error. While many interpretations allow for regulation to take place preemptively (see Carpenter [30]), homeostasis is still assumed to correct deviations from a fixed set-point [31]. In addition, homeostasis is not well suited to deal with variation in demand on bodily systems across contexts and time, variation that has now been well-documented (e.g. [32, 33, 34, 35]). This paper aims to fill this gap by proposing an initial formal model of allostatic regulation. In the process, it will connect existing accounts of motor control based on [internal models](#) [36, 37, 9] and accounts of brain function based on [feedback control](#) [38, 39, 40] to the brain’s regulation of the body’s internal environment.

119 This paper’s formal model of allostasis draws from [control theory](#), a dis-
120 cipline widely employed in both systems biology and engineering. Control
121 theory deals with driving dynamical systems to move (approximately) along
122 a certain desired trajectory, despite physical disturbances to those systems
123 that might drive it off that trajectory. Control theory also makes explicit the
124 question of what the desired trajectory is, how the trajectory might be physi-
125 cally realized, and how one system can drive another to follow a more desired
126 trajectory rather than a less desired one. This paper describes an approach
127 to formally modeling regulation of the body that retains compatibility with
128 previous empirical (e.g. Kleckner et al. [26], Young et al. [41]) and theoret-
129 ical (e.g. Pezzulo et al. [42], Corcoran and Hohwy [43], Petzschner et al. [28])
130 investigations, while building upon control theory from first principles.

131 Four sections in this paper connect interoception to allostasis. Section [2](#)
132 establishes how interoception enables the brain to estimate the physiological
133 efficiency of the body in the present moment, which is precisely what it needs
134 to know to evaluate and refine actions. Section [3](#) then introduces control the-
135 ory and explains its applications in physiology, motor control, and decision
136 making; these provide the conceptual tools for modeling how interoception
137 informs allostasis. Section [4](#) applies the principles of control theory to de-
138 rive a novel formal model of how the brain might estimate the desirability
139 of physiological trajectories and make prospective regulatory decisions. Fi-
140 nally, Section [5](#) synthesizes the previous three sections to explore the direct
141 implications of the proposed formalism. Appendix [A](#) provides a glossary of
142 terms; Appendix [B](#) provides mathematical details related to Section [3](#); and
143 Appendix [C](#) provides mathematical details related to Section [4](#).

144 2. Interoception: modeling the body, estimating its efficiency

145 This section takes up the question of how interoception offers performance
146 metrics for visceromotor regulation. Many interoceptive modalities consist
147 of viscerosensory signals whose values must remain within specific ranges
148 conducive to efficient bodily function and survival (making these signals dif-
149 ferent from exteroceptive sensory signals in this regard). A core assumption
150 is that the brain, as part of allostasis, estimates how efficiently physiologi-
151 cal processes can enable or support needed changes in resource levels [23].
152 Towards that end, Section 2.1 differentiates two types of viscerosensory vari-
153 ables: those that represent *quantities* of resources (called regulated resources)
154 and those that represent *rates*¹ at which processes act (called controlled pro-
155 cesses)². Section 2.2 applies these concepts to the well-studied controlled
156 process of the carotid baroreflex, which the brain must modulate by cen-
157 tral command to meet oncoming demand for the oxygen, glucose, etc. in the
158 blood. This subsection suggests that the brain predicts ongoing fluctuations
159 in physiological efficiency. Section 2.3 considers a more complex regulatory
160 setting, in which several physiological processes act on a common metabolic
161 resource in different ways, and generalizes the proposed notion of physio-
162 logical efficiency estimation to this more common case. Finally, Section 2.4
163 discusses how efficiency estimation in interoception could enable the brain to
164 constructively evaluate a rich variety of predicted bodily conditions without
165 requiring a modular, purpose-specific “reward” system.

166 The discussion of control theory in Section 3 then will use the concepts
167 described here. In Section 4 these concepts will undergird a mathematical
168 formalism for allostatic decision making.

169 2.1. Regulated resources and controlled processes in physiology

170 **Regulated resources** are kept relatively stable over time. Examples in-
171 clude blood glucose and core body temperature. By contrast, a non-regulated
172 resource such as blood alcohol (ethanol) does not have its level stabilized by
173 the body in most contexts. Insofar as a regulated resource like blood glucose
174 represents a physical quantity³ or a substance (like glucose), its quantity can-

¹Quantities of change per period of time

²The terminology of regulated resources and controlled processes comes from Kotas and Medzhitov [35] and Cabanac [33].

³which has physical units independent of time

175 not change instantaneously. Regulation of resources does not have to move
176 levels towards a specific [set point](#)⁴, and in fact, many can freely vary over a
177 range of possible values without regulatory response. Such a range is called
178 a [settling range](#), since the level of a resource might settle anywhere in the
179 range without provoking a regulatory response.

180 A regulated resource remains (relatively) stable over time thanks to the
181 adaptive change of one or more [controlled processes](#). Controlled process rates
182 are the rates at which physiological processes operate. These regulatory
183 processes contribute to the relative stability or change in regulated resources
184 over time. Examples of controlled processes include sweating and shivering,
185 while an example of a physiological change that is not a controlled process
186 is when body temperature increases as a result of the body being in direct
187 sunlight. The rates of controlled processes can speed up or slow down within
188 a broad range by altering energy expenditure. Where a controlled process
189 falls within its operating range will determine its effect upon the regulated
190 resource to which it is coupled. Controlled processes do not have to affect
191 the underlying regulated resource directly by a single causal mechanism; they
192 can have their effect via other controlled processes.

⁴A concept proposed by Cannon [\[44\]](#)

Illustration by example

Returning to the dodgeball example, the full range of physiological processes maintaining a person's ability to play would include the metabolic necessities and byproducts carried in the blood itself: oxygen, glucose, and carbon dioxide chief among them. These can be viewed in terms of the functional categories delineated above. The levels of oxygen, glucose, and carbon dioxide in the blood, at any given moment, are called **regulated resources**. The demand for metabolic inputs by the muscles then can be considered a **controlled process**. In the specific case of the muscles, their metabolic uptake changes the circulating levels of oxygen, glucose, and carbon dioxide. In addition, blood pressure is a controlled process which subserves the maintenance or replenishment of the regulated resources. The heart rate and levels of autonomic activation (in both branches of the autonomic nervous system) then also function as **controlled processes**, modulated to indirectly keep the regulated resources in the desired range.

193

194 In evolutionary terms, controlled processes contribute to the fitness of
195 the organism by responding to changes in the relevant underlying regulated
196 resource to keep that variable within a viable range. In mathematical terms,
197 changes in controlled process rates can be modeled as functions of regulated
198 resource levels. However, those controlled processes also themselves have lim-
199 ited ranges of possible action. The limited ranges of both regulated resources
200 and controlled processes can be predicted and modeled in terms of *capacity*
201 *curves*, which are the topic of the next section.

202 2.2. Predicting and modeling the ranges of regulated and controlled processes

203 Both the heart rate and blood pressure must increase during aerobic exer-
204 cise, as noted in the dodgeball example. If you were to try to play dodgeball
205 at a resting level of blood flow, your muscles quickly would become fatigued
206 and you would be unable to move (for a more detailed discussion, see Ster-
207 ling and Laughlin [22]). Your brain must therefore direct the sympathetic
208 branch of your autonomic nervous system to increase its outflow, including
209 increasing blood pressure via vasoconstriction ⁵. Under resting conditions,
210 the baroreceptor-heart rate reflex would normally counter any rise in blood

⁵Typically with a concomitant withdrawal of parasympathetic activity [45, 46].

211 pressure by slowing the heartbeat. However, with exertion, your blood pres-
212 sure and heart rate both must increase to support the needed increase in
213 blood flow required by your exercising muscles. To accomplish this specific
214 change, your brain modulates the response of your baroreceptor-heart rate
215 reflex [47, 48], shifting the entire function relating a change in your blood
216 pressure to a change in your heart rate [49]. The alterations enable redistri-
217 bution of blood to meet the new demand so that you can run to avoid the
218 ball or throw the ball at someone else.

219 However, blood pressure is a controlled process, not a regulated resource.
220 It must shift in order to stabilize the regulated resources of oxygen, glucose,
221 and carbon dioxide concentrations in the blood. It therefore lacks a set
222 point to which the brain will regulate the baroreceptor-heart rate reflex,
223 the heartbeat, or other variables affecting the blood pressure. Although
224 the controlled process regulating the blood pressure can and does shift its
225 rate with time, that rate can only rise or fall so far before reaching physical
226 limits, after which further modulation of the baroreceptor or the heartbeat
227 will have no additional significant effect. The baroreflex’s responsive range
228 can be defined as the range between where the controlled process (the blood
229 pressure) effectively cannot decrease further (the [threshold value](#) point) and
230 where it cannot increase further (the [saturation value](#) point ⁶).

231 Threshold and saturation points partly define curves that are derived from
232 functions which physiologists commonly use to model the connection between
233 perturbations and regulatory responses, usually naming them [response curves](#)
234 (e.g., Ogoh et al. [51]) or [transfer functions](#). The term [capacity curves](#) will be
235 used to emphasize the fact that while such curves can shift over time, in any
236 one instant they represent the current range of limited regulatory resources
237 available to an organism. The terms *threshold* and *saturation* will also be
238 used for the levels of the regulatory responses (plotted on the vertical axis)
239 of a capacity curve, rather than the levels of the perturbing stimulus (plotted
240 on the horizontal axis).

241 Figure 1 shows an example capacity curve for afferent activity in human
242 baroreceptors. The left tick marker shows the threshold value, and the right
243 tick marker shows the saturation value. A parameter called the [gain](#) specifies
244 the relative slope of the curve throughout its range, determining where the
245 threshold and saturation values will fall. The mean arterial blood pressure

⁶Terminology adapted from McDowall and Dampney [50])

Figure 1: Capacity curve for baroreceptor afferent firing, taken as a pedagogical example from Heesch [52]. As the curve flattens in either direction, the baroreflex can no longer respond proportionally to changes in blood pressure. The tick markers show the threshold value (the fifth percentile of response) and the saturation value (the 95th percentile of response) on the horizontal axis.

(horizontal axis) is a controlled process, and so the baroreflex activation (vertical axis) is also a controlled process, one which only affects the underlying regulated resources (e.g., blood glucose, blood oxygen, etc.) indirectly.

Mathematically, an ideal small change in the blood pressure will lead to a certain ideal small change in baroreflex activation ⁷. The **operating point** is where this potential response is greatest. For a symmetrical capacity curve such as that of the baroreceptor-heart rate reflex above, the operating point will lie in the center of the curve. Figure 2 depicts the operating point with a diamond marker and the potential response around that point as a yellow dotted line. Physiologists often employ the sigmoidal form displayed here for a capacity curve because it provides a good empirical fit to data (see

⁷Mathematically educated readers will recognize this as a derivative.

Figure 2: Capacity curve from Figure 1 above (blue), with the linearized response (orange) around the operating point (blue diamond marker). The diamond marker denotes the point of optimal responsiveness, or operating point. Responsiveness is optimal when the tangent line has maximal slope around the current blood pressure. Regulating to optimal responsiveness requires either keeping current blood pressure near the operating point, or relaxing the baroreflex’s gain to widen the curve. The latter sacrifices performance (slope) at the operating point but provides greater resilience against uncertainty and perturbations. Note that the operating point refers to the point on the horizontal axis.

²⁵⁷ McDowall and Dampney [50], Dampney [53]).

The capacity curve in Figure 1 has mathematical form

$$y(x; \mu, k, R, B) = \frac{R}{1 + \exp(-k(x - \mu))} + B, \quad (1)$$

and its parameters will take values according to the figure. These values include the response range R (from lower to upper asymptote), the lower boundary B on the response, the operating point $x = \mu$, and the gain k . The variable x on the figure’s horizontal axis represents the mean arterial blood pressure, while the variable y on the figure’s vertical axis represents the baroreflex activation as a percentage of the resting mean. For the figure,

the parameters have the values

$$\begin{aligned} R &= 200, & B &= 0, \\ \mu &= 100, & k &= 15. \end{aligned}$$

Equation 1 defined y , the baroreflex activation, as a function of x , the mean arterial blood pressure. Elementary algebra allows the equation to be solved for x or y as a function of an intermediate quantity $u \in (0, 1)$ called the quantile,

$$u(x; \mu, k) = \frac{1}{1 + \exp(-k(x - \mu))}, \quad (2)$$

$$y(u; R, B) = Ru + B, \quad (3)$$

$$x(u; \mu, k) = \mu + \frac{1}{k} \log \left(\frac{u}{1 - u} \right). \quad (4)$$

258 These equations outline the form of a [generative model](#): a procedure for
 259 probabilistically predicting observed variables in terms of unobserved vari-
 260 ables. μ, k, R, B, u serve as unobserved variables, which are sampled from
 261 a prior probability distribution not dependent on data. These variables are
 262 plugged into the equations to generate predictions for the observed variables:
 263 mean blood pressure x and baroreflex afferent activation y . Together, the
 264 prior and the likelihood can form a posterior probability distribution, which
 265 defines the probabilities of different values for the unobserved variables given
 266 the observed ones. In the brain, any internal generative model with similar
 267 structure to the above would likely obtain its probability densities for the
 268 unobserved variables from its general knowledge of the body in the world,
 269 rather than starting with an uninformed prior.

270 The proposal here has an unusual feature: the equations for x (the mean
 271 arterial blood pressure) and y (the baroreflex activation as a percent of base-
 272 line) are in terms of the quantile variable u . The quantile variable uniformly
 273 represents the relationship between the blood pressure and the baroreflex
 274 activation, irrespective of changes in the capacity curve’s operating point μ
 275 and gain k . The quantile depends only on the functional form of the capac-
 276 ity curve, not on the parameters. The distance $u(x) - u(\mu)$ (i.e., the relative
 277 distance between the current value of x and the operating point) therefore
 278 provides a time-independent performance metric for the regulatory task of
 279 the baroreceptor-heart reflex. Capacity curves change all the time due to

variation in their underlying physiological systems (see plot of arterial pressure over time in Bevan et al. [54], reprinted in Sterling [21]), but quantiles will retain the same meaning no matter the current parameter values. This supports high regulatory flexibility, a concept often proposed by physiologists as an adjustable set point [33].

Any point on any capacity curve can be written in terms of quantiles, because capacity curves represent physical responses with finite ranges. Insofar as controlled process responses have the bounded form described above, they can potentially be described in terms of capacity curves, with mathematical description similar to that given above (although usually more complex in the details). Insofar as this remains empirically true, interoceptive internal modeling [19] could be described, mathematically, as *estimating capacity curves over time*. The brain could potentially model those capacity curves in terms of quantiles without loss of generality, and those quantiles would have a clear regulatory interpretation.

Overall, if the brain’s internal model were to infer capacity curves as a part of interoception, then a variety of sites in the brain would have to generate predictions, and integrate prediction errors, regarding both regulated resources and controlled processes. These sites would have to receive afferent viscerosensory signals to which to compare efferent predictions. The brain would have to generate efferent predictions for each capacity curve’s key parameters (e.g., operating point, gain, boundary, and range), and combine those parameters with an efferent prediction of the present state’s quantile representation. These would generate interoceptive predictions of viscerosensory stimuli: the regulated resources and controlled processes related by capacity curves. Afferent viscerosensory signals would confirm or correct these predictions, and thus correct or confirm the estimated performance metric $u(x) - u(\mu)$.

The process of correcting and/or confirming predictions will usually entail spending a sizable amount of energy just on neural firing to update the various predictions [55]. On top of that, the brain also will have to spend energy to reconsider and re-plan current behavior. Imagine an internal chest pain during a game of dodgeball, when you know you haven’t been hit: it could be heartburn, or it could be a heart attack. Whatever the cause, the brain will have a metric of physiological efficiency with which to determine how to spend resources on updating predictions and behavior, so as to optimally keep regulated resources within the responsive ranges of their corresponding controlled processes.

318 *2.3. Modeling the viable ranges of multiple controlled processes to support*
 319 *multi-system regulation and coordinated action*

320 Unlike the simple regulatory relationship between blood pressure and
 321 heart rate, many regulated resources in the body cannot be tightly con-
 322 trolled by a small number of effectors. To return to the dodgeball example,
 323 the full range of physiological processes maintaining a person’s ability to
 324 play would include the metabolic necessities and byproducts carried in the
 325 blood itself: oxygen, glucose, and carbon dioxide are chief among them. In
 326 the specific case of the muscles, their metabolic uptake changes circulating
 327 levels of oxygen, glucose, and carbon dioxide, which are [regulated resources](#)
 328 that must remain within viable ranges. Heart rate, blood pressure, and the
 329 level of activation in both branches of the autonomic nervous are [controlled](#)
 330 [processes](#) that the brain modulates in service to the maintenance or replen-
 331 ishment of the regulated resources. Next, we focus on blood glucose as a
 332 regulated resource, with glucagon as the controlled process enabling secre-
 333 tion of glucose into the blood and insulin as the controlled process enabling
 334 removal of glucose from the blood.

335 Emerging theoretical [56, 57] and experimental [58] evidence suggests
 336 that blood glucose levels are not actively defended at a biologically hard-
 337 coded set point any more than heart rate is. Instead, glucagon and insulin
 338 activity balance each other’s effects to bring the blood glucose to a point
 339 within its settling range (a [settling point](#)) with glucose entering the blood
 340 after the person ingests food and glucose then crossing from the blood into
 341 other bodily tissues to support their function. Recent evidence suggests that
 342 when glucagon stimulates insulin production in β -cells in the pancreas, it
 343 acts to suppress overshoot of the blood glucose level [59]. This suggests that
 344 uptake of glucose into the blood from ingested food could plausibly act as
 345 the passive variable of a settling-point regulation model [60]. The rate of
 346 glucose uptake into tissues from the blood (which leads to insulin secretion)
 347 is plausibly a function of glucose availability in the blood.

348 Settling-point dynamics require either that a controlled process be regu-
 349 lated to decline proportionally to the current level of the regulated resource,
 350 or that outputs be regulated to increase proportionally to the current level
 351 of the regulated resource. Speakman et al. [60] give the example of a water
 352 reservoir, with water as the resource and outflow from the reservoir as the
 353 controlled process. If the depth of the reservoir grows higher due to rain, so
 354 will the volume of outflow. The depth of the reservoir stabilizes when the in-
 355 coming rain and the outgoing outflow over a period of time equal each-other.

356 In the body, both of these forms of regulation can and do happen: they
 357 are the job of the brain [61]. The brain may operate as an additional hi-
 358 erarchical level of control, actively balancing and minimizing the necessary
 359 metabolic control effort by preemptively regulating intake and uptake of glu-
 360 cose through behavior. Thus, uptake of glucose by the muscles during a game
 361 of dodgeball results in a fall in blood glucose and a corresponding increase
 362 in secretion of glucagon [62]. Glucagon acts to cause the release of glucose
 363 into the blood (from liver cells). In the event of glucose overshoot (i.e., ex-
 364 cess levels in blood), insulin will be secreted to restore blood glucose into
 365 its settling range. The brain also registers a “cost” of the glucagon release
 366 because it required energy expenditure (both in synthesis and secretion), an
 367 expenditure that could instead have been spent on the dodgeball game, had
 368 the glucose level been more actively maintained. The reverse can occur when
 369 the blood contains a surfeit of glucose stock, which then must be taken up
 370 into other tissues for storage or usage.

371 Mathematical modeling studies suggest that the inflection (operating)
 372 points on the generalized sigmoidal curves ([generalized capacity curves](#)) for
 373 glucagon and insulin can be found at 3.01 millimolar (mM) and 8.6 millimolar
 374 (mM) respectively, quite close to the lower and upper limits of normal human
 375 blood glucose [63]. Figure 3 shows plots of the resulting functions, which can
 376 be interpreted as *capacity curves*. The inflection points are shown by the
 377 diamond markers, and by definition, each inflection point is a local optimum
 378 of regulatory responsiveness. If glucagon and insulin have their greatest
 379 responsiveness at the limits for hypoglycemia and hyperglycemia respectively,
 380 then regulation of glucose by behavior will aim to avoid straining either of
 381 the two hormones’ capacity curves, effectively keeping blood glucose in the
 382 normoglycemic range.

383 Evidence from existing studies [61, 64, 65] suggests that chemo-sensory
 384 cells in the circumventricular organs and certain nuclei of the hypothalamus,
 385 which detect glucose, may be well-described as predictively modeling such
 386 generalized capacity curves. This kind of functionality may extend not only
 387 to insulin and glucagon, but to other paired (tandem) controlled processes in
 388 the endocrine system, such as leptin and ghrelin [64]. Afferent hypothalamic
 389 firing can be interpreted as sending prediction errors to other parts of the
 390 forebrain [66, 64] thereby signaling an unanticipated metabolic change, and
 391 potentially updating the brain’s internal model as described by these capac-
 392 ity curves. Each such capacity curve and the current point upon it could then
 393 motivate some mode of behavior: a predictable mixture of glucagon, ghrelin,

Figure 3: Capacity curves for glucagon (blue) and insulin (orange) responses to glucose levels in the blood, measured in millimolars. The diamond markers show the respective operating points (3.01 mM for glucagon, 8.6 mM for insulin) of the curves, and the space between those two markers denotes the potential settling range for blood glucose content. Derived from König et al. [63].

and other similar signals could be well-described as motivating consumption behaviors (e.g., a shift in autonomic activity towards greater relative sympathetic nervous system activation, exploration of the environment, etc.), while a predictable mixture of insulin, leptin, etc. could be well-described as motivating satiety behaviors (e.g., a shift towards relatively greater parasympathetic activity, reduced motor activity, etc.).

Similar paired controlled processes seem to appear in a variety of regulatory “modalities” throughout the body, ranging from autonomic activity on the heart [45] to blood glucose [61] (as above) to adiposity [60] (in the form of leptin and ghrelin). Evidence also suggests that the brain can combine these signals when they operate in tandem to regulate common behaviors [67]. Allostasis may employ a generalized control motif of having paired peripheral controlled processes, which sometimes work together to drive regulatory

behavior in one direction (e.g., a reciprocal mode of sympathetic increase and parasympathetic decrease which both drive the heart rate to increase), but can also “antagonize” each other’s effects (such as when both the sympathetic and parasympathetic branches coactivate, producing a heart rate that is the sum of these two countervailing forces, each driving the heart rate in a different direction). Interoceptive modeling in the brain also may employ these motifs. Such general motifs in interoceptive processing could provide a domain-general mechanism for quantifying regulatory imperatives in interoceptive internal models. This may provide greater flexibility in both physiological regulation and behavior than a centrally enforced set point can provide, as well as suffering less error against challenges in each direction (see [56, 57]).

2.4. *Viable ranges and capacities could obviate a modular “reward system”*

Standard accounts of allostatic regulation describe it chiefly on the physiological level of analysis, attributing allostatic control in the central nervous system to reinforcement learning. As Sterling [21] writes,

The central representation of “reward” is a brief burst of spikes in neurons of the ventral midbrain that release a pulse of dopamine to the nucleus accumbens and prefrontal cortex. The precise correspondence between a “feeling” and a specific neuro-transmitter is difficult to establish and is probably oversimplified, since many chemicals change in concert. Yet, one imagines that the dopamine pulse evokes momentary relief from flagellating anxiety and a brief sense of satisfaction/pleasure – at last, the carrot.

The picture of “rewards” painted here suggests a modular “reward center” or “reward system” in the ventral midbrain, one whose specialized role is to perform apples-to-oranges comparisons in service to allostasis. However, insofar as the ventral midbrain would function as a “reward center”, the “reward” signals sent to the rest of the brain would not carry contextual information about the bodily needs to which they refer. More recent evidence shows that there is no unique, localized “reward center” or “reward system” in the brain: broad cortical and subcortical brain networks play various roles in reward as a *construct* [68] or an *abstract concept* in experiments.

Since there is no single brain site that specifically encodes appetitive or aversive reinforcement value, it is useful to reframe discrete “reward” and “decision” systems as a domain-general allostatic control system. Abundant

443 empirical evidence supports such a reframing [69, 70], particularly analyses
444 of the *default-mode network* and the *saliency network* and their subcortical
445 connections [19, 26]. Computationally, these networks could implement a
446 formal model similar to what we introduce later in this paper, or they could
447 translate interoceptive information into a teaching signal for a reinforcement
448 learning system, as in Keramati and Gutkin [71]. The domain-generalty of
449 interoception provides further theoretical support for the idea that we do not
450 need “mental modules” or “faculty psychology concepts” to understand how
451 a brain works [72, 73].

452 If interoceptive processes operate to estimate parameters analogous to
453 operating points and tolerances, then those processes should be able to con-
454 vey sufficient information to the brain for purposes of regulating the body.
455 The ideas of Section 2.2 and Section 2.3 can be usefully combined here
456 when considering controlled processes that regulate the same underlying re-
457 source. From this perspective, interoceptive prediction errors, in the context
458 of decision-making experiments, can be interpreted as learning about “re-
459 wards” via “reward prediction errors”. Movement towards an operating point
460 then can be considered a “reward”, and movement away from an operating
461 point a “cost”. Each such movement can be weighted according to the same
462 capacity curve’s gain or inverse-tolerance; this would convey the (momentary,
463 estimated) relative “worth” of adapting to a load on one innervated organ
464 system versus another. Conceived of in a high-dimensional space, such move-
465 ments can be viewed as “towards” and “away from” trajectories of changing
466 operating points.

467 Within a view of brain function not based on a modular reward system,
468 the neurotransmitters produced by unmyelinated and lightly myelinated in-
469 teroceptive nerve fibers (see Carvalho and Damasio [74]) could play a role in
470 signaling capacity curves. These neurotransmitters, which include dopamine
471 and serotonin, are commonly thought to act as teaching signals for action [75].
472 In support of this suggestion, dopaminergic neurons in the ventral tegmental
473 area in mice [76, 77, 78] have been modeled using a class of mathematical
474 functions that include our capacity curves.

475 A mathematically sound and biologically plausible account of allostatic
476 control does not require a modular or separate “reward system” in the cen-
477 tral nervous system. Rather, it simply requires a brain and a viscerosensory
478 peripheral nervous system to behave as if parameters for capacity curves (de-
479 scribing how adaptable any given state *would be* to unexpected disturbances)
480 were signaled alongside the location of current physiological states on the

481 corresponding capacity curves. Different physiological needs (say, core body
482 temperature versus blood glucose levels) could then be added, subtracted,
483 compared, etc. by comparing the distance of the current state from the op-
484 erating point in any given dimension, scaled by the capacity curve’s gain.
485 Our formal model later will make this idea more precise, providing a way to
486 put numbers to such “distances” and “movements”.

487 2.5. *Summary*

488 This section outlined a proposal for how the nervous system could poten-
489 tially function to coordinate and control organ systems across timescales to
490 provide allostatic regulation of the body. Section 2.2 considered the move-
491 ment of physiological systems’ response curves (such as the example shown
492 in Figure 1) as signifying their capacity to adapt to challenge; the idea of
493 matching actual system loads to operating points (along the lines of Fig-
494 ure 2) provides a foundation for allostatic regulation. Section 2.3 extended
495 the idea of capacity curves to systems based on more behavioral, settling-
496 point regulation; in such systems the brain functions as the top level of
497 a hierarchical control scheme, regulating the lower-level controllers. In Sec-
498 tion 2.4 we then reasoned that movement toward or away from the responsive
499 range of a capacity curve can be treated as “reward” or “cost”, respectively,
500 and suggested that this could potentially obviate the need for a dedicated
501 neural circuit or module that specifically calculates the behavioral constructs
502 of “reward” and “cost”. Evidence from cognitive neuroscience supports the
503 view that the brain lacks such modules, suggesting that we may gain empiri-
504 cal and theoretical traction by investigating decision-making constructs from
505 an allostatic point of view. Section 3 will build on these ideas by introducing
506 control theory, and consider the use of control theory in physiological regu-
507 lation, motor control, and decision-making, as well as discuss the potential
508 for control theory to unify previously disparate views on bodily regulation.
509 Section 4 will build upon these foundations to propose a formalism for allo-
510 static decision making. Embodied decision making includes all three of the
511 forms of uncertainty to which allostatic regulation is subject: uncertainty
512 about what is physiologically efficient, uncertainty about the consequences
513 of movements, and uncertainty about the external world.

514 3. Control theory: A unifying lens for physiology, motor control, 515 and decision making

516 Section 2 hypothesized that allostatic regulation can be understood in
517 terms of controlled processes’ responsiveness to perturbation. It introduced
518 capacity curves, responsive ranges, regulated resources, and controlled pro-
519 cesses as ways to describe aspects of physiological regulation; it also sketched
520 the functional form of a generative model that could infer capacity curves
521 as latent properties of related interoceptive variables. It then suggested that
522 moving actual physiological states towards the operating points of maximum
523 adaptability, with each movement weighted by the relative gain (inverse-
524 tolerance) of the response capacity, can formalize the functional dynamic of
525 allostatic regulation. However, interoception is *perception* of the innervated
526 body: it can include sensing allostatic responsiveness of present states of
527 the body, as achieved by past actions, but it cannot produce present and
528 future actions in and of itself. The latter is the role of visceromotor control
529 processes. To investigate how the brain accomplishes visceromotor control,
530 some additional theoretical tools are required.

531 This section introduces concepts from engineering control theory, and
532 then reviews its applications in the life sciences. These include physiology
533 (Section 3.1), skeletomotor movement (Section 3.2), and decision making
534 (Section 3.3). Section 3.4 will connect future interoceptive states to present
535 movements, illuminating what makes allostatic regulation more energy effi-
536 cient than homeostatic regulation. The next section will build off the account
537 of control used in physiology to suggest how interoception supports allostasis.

538 3.1. Control theory for physiology: A reliable body built from unreliable parts

539 Broadly, [control theory](#) deals with driving a physical system towards a
540 desired trajectory, even when the system is built from unreliable or unpre-
541 dictable parts. Control theorists call the driven system the [plant](#), and its
542 desired trajectory the [reference trajectory](#)⁸. Generally a plant must be made
543 to conform to its reference trajectory by a driving system called the [controller](#).
544 In controls engineering, these systems are typically thought to be separate

⁸Generally there is assumed to be a collection of possible desirable trajectories, so that a system can compensate for severe disturbances to its original trajectory by picking another acceptable trajectory from the collection. Theorists have labeled this quality “meta-stability”.

Figure 4: Functional block diagram of a model-based control system. The “plant” (orange) is the object or system whose motion or other behavior is controlled. The “controller” (purple) sends signals (“controls”, solid black arrows) that change how the plant moves, and signals the expected outcome (“predictions”, solid yellow arrow) to the “state estimator” (yellow). The task behavior of the plant is prescribed to the controller by an engineer or machine operator (transparent box), in the form of the reference signal (solid purple arrow). Measurements of the plant output (dashed yellow arrow) feed back to the state estimator to yield updated estimates (dashed purple arrow), which the controller compares to the reference signal to adjust the controls.

545 physical entities with connections between them, and along these connections
 546 the systems transmit signals to each other. The “reference signal” that spec-
 547 ifies the reference trajectory goes into the controller, and signals that leave
 548 the controller and enter the plant are called [controls](#). Controls affect the
 549 [state](#) of the plant over time. The reference signal is thought to derive from a
 550 source that is external to the system, such as an engineer or a machine oper-
 551 ator. Figure 4 shows an example “block diagram” for an engineered control
 552 system.

553 A controller functions to steer the plant along its reference trajectory,
 554 adapting to external [disturbances](#) that would push the plant away from the
 555 reference trajectory. From the standpoint of a brain controlling a body,
 556 “disturbances” might be thought of as uncontrolled changes in the workings
 557 of the body’s internal systems. There is an important distinction between an
 558 unpredictable event and a disturbance: unpredictable events can either push
 559 a system away from its reference trajectory or towards it, but a disturbance,
 560 which may or may not be surprising, is always an event that pushes the
 561 system away from its reference trajectory. Thus, a disturbance is always
 562 relative to the reference trajectory.

Illustration by example

The dodgeball example illustrates the distinction between an unpredicted event and a disturbance. If a player looks across the field and sees a ball heading straight for her, her brain knows and predicts (by means of past experience) when and where the ball is likely to arrive. If she sees another player throw the ball, it will not be surprising (i.e., unpredicted) when the ball heads her way and hits her body. It will, nonetheless, be a disturbance; insofar as her brain estimated the movement of the ball and prepared her body to dodge, if she was unable to move fast enough, she deviated from her reference trajectory. By the same token, if a ball hits her from behind as she is standing still, her brain has made no estimate of its trajectory, nor has it prepared her body to dodge, but the hit remains a disturbance. For the same reason, if she positions herself in a way that enables her to catch, by luck or accident, a ball thrown at her, she has followed her reference trajectory, even though her brain will only register this after processing the ensuing prediction errors.

563

564 The whole point of a control system is to adapt to disturbances, and a
565 system can attain much greater robustness and adaptability by using sensors
566 to measure the plant's actual behavior over time. Control theorists call these
567 **measurements** the feedback for the controller, and the use of feedback to
568 adjust control outputs is called **feedback control**. Physiologists recognize
569 feedback control as a ubiquitous feature of bodily function [30, 79], with
570 endocrine control of blood glucose being a well-studied example [56, 57].
571 Feedback control is essential because no controller is ever perfect: neither all
572 forms of noise nor all external disturbances can ever be fully accounted for.
573 If the feedback loop (seen in Figure 4 as the arrows flowing from controller
574 to plant, plant to state estimator, and state estimator to controller) is cut,
575 the controller can no longer receive any information from the plant. The
576 control signals calculated under such circumstances are called *open-loop* or
577 *feedforward* controls, or even (somewhat idealistically) *plans*.

578 The concepts of control theory can illuminate the anatomy of the barore-
579 flex, the example physiological system described above. For simplicity's sake,
580 the baroreflex and its components are considered the system in question. Its
581 plant is the organs of the cardiovascular system as innervated by the auto-
582 nomic nervous system (ANS). Its controller is a comparator circuit in the

583 midbrain, specifically in the nucleus tractus solitarius (NTS) [80]. Its refer-
584 ence signal comes from two sources: over the short term, top-down signaling
585 from the cerebral cortex (which is outside the controller itself), and over the
586 long term, midbrain structures rostral to the NTS, based upon (as yet not
587 fully understood) endocrine signals [81]. These are compared to the actual
588 blood pressure measured by the carotid and aortic baroreceptors (feedback
589 state estimator). The baroreflex (the controller) adjusts cardiovascular vari-
590 ables to align the measured blood pressure (as sensed by the baroreceptors)
591 with the reference trajectory of operating points (as specified by short-term
592 signaling from the forebrain or longer-term endocrine signaling), by means
593 of the sympathetic and parasympathetic branches of the autonomic nervous
594 system (as control signals).

595 For a controller to perform well, it must contain some sort of copy or mir-
596 ror of the plant’s expected behavior, which is referred to as an [internal model](#)
597 [82, 83]. However, inaccuracies in the internal model limit controller perfor-
598 mance (as does an absence of feedback in open-loop control). Internal models
599 serve a dual purpose: to infer past trajectories⁹ (including their control sig-
600 nals) on the basis of present or even counterfactual measurements (such as
601 the reference signal) and to estimate future states and measurements¹⁰ in the
602 plant on the basis of control signals. Figure 4 divides these two functions
603 into the components of a control system that they inform: the controller
604 (purple) infers controls from state estimates to track the reference signal¹¹,
605 and the state estimator (yellow) predicts future measurements on the basis of
606 present control signals, refining those [state estimates](#) using measurements¹².
607 Together, internal models predict the future and infer the past in the plant
608 even when plant behavior is subject to [process noise](#) and measurements are
609 subject to [measurement noise](#).

610 Internal models play an important and specific role in control theory [84,
611 85]. State estimation of future measurements based on present control signals
612 allows the difference between predicted and actual measurements, which is
613 called the [prediction error](#), to be used to measure the accuracy and precision
614 of online control. The state estimator (yellow) can use prediction errors to
615 send updated state estimates (dashed purple arrow) back to the controller.

⁹As an [inverse model](#)

¹⁰As a [forward model](#)

¹¹Inverse modeling

¹²Forward modeling

616 Comparing the updated state estimates to the reference signal then yields
617 a quantity called the [control error](#), on the basis of which the controller can
618 refine the control signals. The process of calculating prediction errors and
619 control errors online, and using them to improve control signals, is called
620 [feedback control](#). Model-based feedback control is widely used and well-
621 known for driving control error to zero over time (a property called [stability](#)),
622 particularly when prediction errors are also driven to zero by refining the
623 internal model. This is a property that open-loop planning cannot enjoy,
624 since an open-loop control system does not measure the control error, which
625 amounts to (falsely) assuming it to be zero. An internal model alone only
626 suffices for open-loop, feedforward planning, while control requires feedback.
627 Controllers also can couple to each other hierarchically: a higher-level
628 controller can send a control signal to a lower-level controller, which functions
629 as the reference signal for that lower-level controller. In turn, the lower-level
630 controller may send a control error signal up to the higher-level controller.
631 The higher-level and lower-level controller also each may have their own
632 state estimators based on their own internal models. The next subsection
633 will address hierarchical control in human motor control.

634 *3.2. Moving the body: The referent control hypothesis*

635 The *referent control hypothesis* [86, 87] describes the skeletomotor system
636 in terms of a hierarchy of controllers, with higher-level controllers in the brain
637 prescribing reference trajectories to the lower-level reflexes in the spinal cord.
638 These reflexes then compare the actual length of the muscle, as signaled by
639 afferent proprioceptor neurons, to the reference length sent down by the
640 brain, and contracts the muscle to bring the two into agreement [88, 89].
641 Effectively, higher level controllers tell lower ones what trajectory to visit,
642 and the lower ones figure out how to track it successfully, a phenomenon
643 beginning to be considered in engineered (i.e., non-biological) control systems
644 [90].

645 Cortical regions involved in skeletomotor control (e.g., primary motor
646 cortex, premotor cortices, etc.) also send a copy of the downward-flowing
647 reference signals to somatosensory cortices (called an efferent copy), thereby
648 providing prior predictions to somatosensory regions. These “prior” signals
649 literally change the firing of neurons in somatosensory cortices, preparing
650 them to receive incoming signals from the world based on upcoming skeleto-
651 motor movements. This dynamic takes place across all levels of the neural
652 hierarchy, allowing the nervous system to use somatosensory prediction errors

as feedback to confirm or correct movement performance; it also allows the nervous system to distinguish **reafferent** (self-caused) from **exafferent** (externally caused) sensory signals¹³. Through this lens, the brain is considered to act as both a controller (in its visceromotor and skeletomotor functions) and a corresponding state estimator (in its perceptual and simulation functions). Since under the referent control hypothesis, skeletomotor “commands” take the form of descending signals specifying desired lengths and tensions for proprioceptive measurements, the brain’s state estimation machinery in sensory areas therefore can simulate the somatosensory consequences of those descending control signals.

The brain is hypothesized to exploit its internal model of both the body and the local external environment to predictively construct populations of reference trajectories as embodied simulations [14]. These simulated populations of referent trajectories can also be thought of as action concepts [92, 93] (for similar ideas see also Wolpert et al. [94] and Hickok [95]). In the brain, motor areas are hypothesized to implement feedforward control with action concepts by decompressing low-dimensional referent trajectories from higher in the neural hierarchy into higher-dimensional referent trajectories lower in the neural hierarchy¹⁴. Decompressive prediction by the brain eventually produces referent coordinates in the highest-dimensional, most redundant system of coordinates: references for individual peripheral stretch reflexes [86]. Each such stretch reflex circuit compares the ascending stimulus from its proprioceptor neuron to its centrally commanded reference, and activates its motoneuron to suppress the difference between the two. The motor system thus can be imagined as a hierarchy of controllers, with higher-level controllers specifying the reference signals for lower-level controllers as their output control signals. In addition to converting signals from the conceptual reference frame of the behavioral task to the concrete reference frames of individual limbs and muscles, this hierarchical structure finesses neural signaling delays to provide fast, accurate control [98].

Figure 5 shows an elegant “outside-in” view¹⁵ of the motor control problem, imagined along these lines. Here, an experimenter directs a participant to perform a task, who constructs an action concept from the instructions.

¹³See Straka et al. [91] for a thorough review.

¹⁴A finding discussed commonly in the predictive processing [96] and neuro-robotics [97] literatures.

¹⁵Language taken from Buzsáki [99].

686 This action concept seeds the construction of sensorimotor prediction popula-
687 tions in the skeletomotor decision controller. The skeletomotor controller fur-
688 ther unpacks the action concept into an actual body posture and its attendant
689 reference coordinates for muscles, as well as predictions for the somatosen-
690 sory state estimator. The movements created by the reference coordinates
691 are themselves stabilized by fast proprioceptive feedback at the level of
692 the individual stretch reflexes in the spinal cord. The hypothesis expressed
693 by the figure assumes that all systems in the body, both somatomotor and
694 visceromotor pathways, serve an externally driven behavioral task.

695 However, the structure of a motor-control experiment has only limited
696 overlap with the actual structure of motor control as it unfolds in the nat-
697 ural world. An experiment on reaching behaviors involves an experimenter
698 prescribing the reaching task to their participants. A participant’s brain
699 does not function specifically to follow instructions from an experimenter,
700 but rather to regulate their own body. The actual organization of the central
701 nervous system accords better with an “inside-out” view¹⁶ of motor function:
702 movement of the body (the somatomotor pathways) serves regulation of the
703 body. For example, in a game of dodgeball, if you unexpectedly step hard
704 on a sharp rock, you (usually) do not purposely impale yourself to stabilize
705 your posture. Rather, you recoil in pain, and the unplanned disturbance of
706 any tissue damage requires you to make a decision about what to do next:
707 excuse yourself to nurse your foot, or play through the pain.

Terminology

The terms for directions of neural signaling, depending on the implied origin of the signal, are efferent, afferent, and reafferent. Motor neuroscientists refer to efferent signals (flowing from somatomotor cortex down to somatosensory areas, the midbrain, and the peripheral nervous system) as feedforward signals. They then refer to afferent signals, particularly reafferent somatosensory signals, as feedback signals. Since the usage from motor neuroscience agrees with that from control theory, this paper follows that usage.

708

¹⁶Language again taken from Buzsáki [99].

Figure 5: Functional block diagram for an experimental psychologist’s task-oriented view of motor control. The diagram shows a formal logical structure here, at a conceptual level; the boxes and arrows do not map onto the anatomy of the brain or nervous system. In contrast to Figure 4, this diagram differentiates between skeletomotor (brain) and peripheral (stretch reflex) controllers and between sensory state estimators (brain) and peripheral sense organs (sensory surfaces). The diagram shows an engineering perspective on a psychology experiment, in which the experimenter prescribes a task or behavior to participants, and a participant’s brain then acts as control system to achieve the prescribed behavior. Systems that maintain the body therefore serve systems that move the body, which in turn serve a prescribed behavior.

3.3. Making decisions: constructing future reference trajectories

So far control theory has been presented as it applies to physiology and motor neuroscience. These control mechanisms have largely been local, in the sense that they only drive neural outputs to impact a small, narrow domain: blood vessels and baroreceptors modulating autonomic outflow to reduce the heart rate; proprioceptors in individual muscle spindles driving the stretch reflex. In the stretch reflex, the reference is set by top-down signaling from the brain as part of voluntary movement, and therefore varies widely. The baroreflex displays a similar structure, receiving parameters for its capacity curve via central command.

Reactive control requires reference signals before physiological needs are fully known. In contrast, the brain controls the body predictively, both viscerally and through overt somatomotor behavior. Predictive neural control of the viscera must take into account biological processes that change the body (such as those occurring with development) and cyclical routines such as the wake-sleep cycle. The brain also must coordinate across a variety of

725 physiological demands in the moment, each with its own capacity curve that
726 may change over time. At the same time, the brain is subject to uncertainty
727 from both sensory/measurement noise and process noise in the innervated
728 tissues. Decision making in the brain must therefore operate according to
729 control principles that take into account competing demands over time.

730 A control theory formulation meeting these criteria is [stochastic optimal](#)
731 [control](#) (SOC). Stochastic optimal control uses probability distributions over
732 states $\mathbf{x}$ to model the effects of both process noise and measurement noise.
733 The goal then is to optimize the probabilistic expectation of an [objective](#)
734 [function](#) $\mathcal{J}$ summed over the indefinite future. This expectation of a sum,
735 accounting for both uncertainty and time, is called the [value function](#), and
736 is defined recursively by the [Bellman equation](#). A control strategy is optimal
737 when it maximizes the value function. Since objective functions include terms
738 that quantify the relative tolerance for regulatory error, it supports a wide
739 variety of approximate solution methods to find “good enough” controls,
740 which drive the plant close to its reference trajectory even when noise and
741 disturbances prevent exact reference tracking.

742 Objective functions generalize the simple comparators often used in clas-
743 sical control; they compare a reference state with an actual state to generate
744 a real number as output. Generally the output value should be monotonically
745 related to the degree of agreement or disagreement between the reference and
746 the actual state. Objective functions with a fixed notional reference state and
747 a fixed tolerance for deviation from that state can model fixed physiological
748 set-points and tolerances, like those found in homeostatic theories of regula-
749 tion.

Terminology

Predictive homeostasis is the hypothesized mode of regulation in which anticipatory decision-making mechanisms maintain regulated resources at fixed set-points with fixed tolerances. Neurally, a homeostatic regulator would consist of a comparator circuit without an incoming signal modulating its reference. Since most existing computational models of visceromotor control and interoception (e.g., [28, 100]) fall into this camp, we consider them to model allostasis as predictive homeostasis, designed for settings in which long-run set-points and tolerances define the chief control mechanism. This family of models includes certain active inference approaches [42, 101, 102] and homeostatic reinforcement learning [71].

Sometimes in stochastic optimal control, the objective function is not known *a priori*, and must be learned from samples. This special case is called reinforcement learning (abbreviated as RL). Neuroscientists have studied reinforcement learning [103] as a model of how the brain might make decisions over time. Modeling midbrain phasic dopamine signaling with RL has led to a popular approach in computational neuroscience (the reward prediction error hypothesis [104, 105]), and continues to yield novel findings today [106]. Notably, these successes rely on a specific way to approximate the value function for behavior over time by comparing predictions generated in the brain to the actual sensory effects of behavior.

Stochastic optimal control, at first glance, might seem unnecessarily complex. Taking an expectation of a sum over time seems almost repetitive: why should the brain not just wait until it arrives to a future bodily state, and then compare it to the corresponding reference using the objective function? This is, after all, precisely how reactive control mechanisms work. What problem is allostatic decision-making in the brain solving that a homeostatic reflex in the body cannot? Reflexive, reactive control under uncertainty carries a hidden assumption: that uncertainty in the moment is equivalent to uncertainty over time, and thus if control mechanisms can compensate for errors in the moment, they can compensate for all errors in the future. When that assumption holds, predictive and reactive control will be equivalent. When that assumption does not hold, stochastic optimal control can yield vastly better regulation than reactive control.

764 That assumption is called [ergodicity](#), and it amounts (very roughly) to
765 modeling time as having no effect on probability distributions. [B.1](#) discusses
766 ergodicity in detail, including its implications for experimental design. [B.2](#)
767 also discusses a paradigm for studying the brain as a whole that does assume
768 ergodicity, allowing it to bridge perceptual processing with decision making
769 and motor control. The following material places more emphasis on situations
770 that are *not* ergodic, which include those with periodic structure, or with
771 irreversible changes over time. Real life is filled with non-ergodic situations
772 in which one must make decisions: the cycles of day and night (and circadian
773 rhythms tracking them) are non-ergodic; processes of development (from
774 childhood to adulthood) are non-ergodic; events like injury and death are
775 non-ergodic. The brain accounts for these non-ergodic realities of life in
776 decision-making processes [\[107\]](#).

777 One final alternative way to think of SOC is in terms of what a controller
778 needs to model, and which signals into a controller may be subject to noise or
779 disturbance. In a typical control problem, the control engineer “trusts” that
780 she can provide an exact reference signal, while building the controller to be
781 robust to noise and disturbances in the plant. In SOC and RL, the control
782 engineer may not be able to specify the reference signal exactly, but she
783 can provide the controller with encouragement (rewards) or discouragement
784 (costs) for following observed trajectories. “Rewards” thus count as evidence
785 in favor of the plant’s recent behavior following the reference trajectory, while
786 “costs” count as evidence for behavior deviating from the reference trajectory
787 [\[108\]](#). This implies that an optimal controller, viewed from the perspective of
788 [Figure 4](#) or [Figure 5](#), has two sources of uncertainty that require two separate
789 state-estimation processes: one for the state of the plant, and another for the
790 reference signal. The next subsection will describe how interoception works
791 alongside somatosensory perception to reduce both of these uncertainties.

792 *3.4. Allostatic control: Motivating movements with an interoceptive model*

793 The brain allostatically regulates the body. Accordingly, there ought
794 to be a description of the brain in terms analogous to inferring capacity
795 curves; these can be transformed into objective functions and projected into
796 the future to construct a value function. This value function would allow the
797 brain to take account of the future when deciding what to do now. When the
798 value function successfully predicted instantaneous allostatic capacity curves,
799 it would seem as if the future body prescribed a reference trajectory to the
800 present brain. Since moving the body would predictably change allostatic

Figure 6: Functional block diagram for a control-theoretic view of allostasis. In contrast to Figure 5, this diagram shows a closed-loop control system design for autonomous regulation of the body. An experimenter’s desired “task behavior” is replaced by the allostatic capacity estimator, which sends predictions of capacity curves to the interoceptive state estimator and receives prediction errors with which to update its estimates. The updated estimates are issued as a reference signal to the visceromotor controller. This diagram shows a formal logical structure, at a conceptual level; the hypothesis depicted is constrained by the inferred anatomical structures in Barrett [15] but the boxes and arrows do not map one-to-one onto the anatomy of the brain or nervous system [109].

capacities (i.e., operating points and tolerances) in the future, a “circular causation” of self-organization and autonomy would emerge.

Figure 6 diagrams allostatic regulation using the language of control theory. This figure employs the same visual vocabulary of plants (orange), controllers (purple), internal models (yellow), feedforward signals (solid arrows), and feedback signals (dashed arrows) as the earlier Figures 4 and 5. Systems that move the body are hypothesized to be in the service of systems that maintain the body (i.e., movement is in service of allostasis).

Allostatic regulation contains homeostatic regulation as a special case. Allostasis consists of regulating a system’s state to track a reference trajectory, one which fully allows for system states to change over time. Homeostasis consists of regulating a system’s state towards reference points, independent of time. Thus, an allostatic controller can implement homeostatic control by prescribing a reference trajectory as a single, unchanging point, while a homeostatic controller cannot implement allostatic control. This is because homeostasis does not really deal with context. The allostatic trajec-

817 tory approach allows for the possibility that past behaviors affect the present
818 state and reference. It allows the effect of past perturbations or interoceptive
819 conditions to influence what happens next.

Illustration by example

A return to the dodgeball example will ground these ideas. During a game of dodgeball, muscles will demand greater amounts of oxygen and glucose than they had needed during rest. Successfully throwing the ball at an opposing player will require mobilizing both the skeletomotor musculature (“soma”, plant) as well as the internal bodily systems such as the cardiovascular system (“viscera”, plant) across a timescale of tens of seconds to minutes via the visceromotor and somatomotor controllers (purple). We hypothesize that a functional equivalent to an allostatic capacity estimator (yellow) anticipates the need, altering the reference trajectory conveyed as a prediction to the visceromotor controller (purple). The visceromotor controller must then mobilize the cardiovascular system to supply those metabolic necessities via the blood. In this instance, among other adjustments, the visceromotor controller shifts and flattens the baroreflex’s capacity curve [110], allowing both vasoconstriction and an increased heart rate to work in tandem to supply more blood flow to the muscles.

820

821 Translating the example into the language of Figure 6, the allostatic ca-
822 pacity estimator (yellow) signals the anticipated demand as a reference tra-
823 jectory (solid purple arrow) to the visceromotor decision controller. The
824 visceromotor controller decompresses this low-dimensional reference trajec-
825 tory into higher-dimensional reference trajectories (solid purple arrows) for
826 the skeletomotor decision controller (purple) and peripheral reflex controllers
827 (purple). Simultaneously, the visceromotor decision controller sends efferent
828 copies (downward solid yellow arrows) to the viscerosensory (“interoceptive”)
829 and somatosensory exteroceptive state estimators (yellow) to generate sen-
830 sory predictions (e.g., predict that the heart rate will be close to the ref-
831 erence). The skeletomotor controller decompresses its own reference trajec-
832 tory into even higher-dimensional reference trajectories (solid purple arrow)
833 for peripheral reflex controllers (purple); these convey motor signals about
834 where your hands and feet should be, how bent your knees should be, etc.
835 The skeletomotor controller also emits efferent copies to the somatosensory
836 model, which generates sensory predictions regarding the sense data coming

837 from sensory surfaces: the strain of bending the knees, the thump of the heart
838 against the chest, and so on. Finally, the local reflex controllers enact motor
839 controls (solid black arrow) that move the body. Basing sensory predictions
840 on the efference copies enables the resultant measurements at the sensory
841 surfaces (dotted yellow arrow) to generate sensory prediction errors (dotted
842 yellow arrows), which serve as feedback on the timescale of tens or hundreds
843 of milliseconds. This feedback flows through the state estimators to update
844 their state estimates, and these estimates are then signalled to controllers as
845 control feedback. At the far end, updates to interoceptive state estimates
846 generate prediction errors that update the estimated allostatic capacities,
847 thus closing the loop.

848 3.5. Summary

849 This section discussed three applications of control theory to studying
850 the body and brain. Section 3.1 described control theory as a whole and
851 discussed its applications in physiology, providing an example of a control
852 system in Figure 4. Section 3.2 then discussed the study of voluntary motor
853 control in the nervous system. The construct of reference trajectories in
854 control theory finds a close analogue in the *referent control hypothesis* of
855 somatomotor control, and its application in Figure 5 shows an engineer’s view
856 of motor control. Section 3.3 discussed the necessity for allostatic decision-
857 making in the brain to take account of changing bodily conditions over time,
858 constructing references based on physiological capacity curves. The next
859 section will apply the concepts from this section to describe a formal model
860 of allostatic regulation.

861 Before diving into formal modeling details below, it might be helpful to
862 compare and contrast the approach here with a prominent modeling frame-
863 work: active inference [111, 112]. Like many active inference models, the
864 formal model below takes the form of information-theoretic model-predictive
865 control (similar to work such as Williams et al. [113] and Nasiriany et al. [114]
866 in engineering) with a hierarchically-defined objective (similar to Smith et al.
867 [115], Pezzulo et al. [101]). Insofar as such efforts can be considered “active
868 inference”, the formal model outlined in the next section is also an active
869 inference model. Unlike most active inference models in the literature, how-
870 ever (with the exception of Millidge [116]), the material below considers an
871 indefinite-time or “infinite horizon” control setting. Insofar as the research
872 community prefers for the term “active inference” to refer specifically to

873 formal models derived from the free energy principle (see [B.2](#) and Kirch-
874 hoff et al. [117]), with its ergodic assumptions and its unique expected free
875 energy objective, the formal model below is distinct from these traditional
876 active inference models.

877 4. Allostasis as trajectory-tracking stochastic optimal control

878 The previous section overviewed and applied control theory to study phys-
879 iological reflexes, voluntary motor movements, and decision making in the
880 brain. The previous section described stochastic optimal control (SOC) the-
881 ory as a mathematical formalism capable of flexibly modeling decision making
882 under uncertainty. This section will describe our formal model of allostatic
883 decision making: the Allostatic Path-Integral Control (APIC) model. APIC
884 has a simple idea at its core: just as perceptual concepts serve as internal
885 models of the body’s sensory surfaces [15, 92, 14], action concepts also serve
886 as internal models of potential behaviors and their predicted outcomes. The
887 brain refines and selects sensory predictions derived from a concept based on
888 their fit to past and present sensory evidence; we suggest that it likewise re-
889 fines and selects the motoric reference configurations derived from an action
890 concept based on their present and future allostatic value.

891 Section 4.1 derives an SOC objective function from the mathematical
892 form by which Section 2.2 represented capacity curves. Section 4.2 marshals
893 behavioral and ecological evidence of how animals balance and meet their
894 needs over time into a long-run mathematical form. Section 4.3 then sketches
895 a formal model of how action concepts fit into SOC. Section 4.4 describes
896 how to exploit action concepts optimally in feedback control. These last two
897 subsections include the motivations for particular modeling choices.

898 4.1. Transforming capacity curves into objective functions

899 Section 3 pointed out that allostatic control poses a decision-making prob-
900 lem: the brain must predict the body’s future needs in the form of a reference
901 trajectory or [value function](#), and move to satisfy those needs before they be-
902 come acute. This subsection presents a mathematical treatment of capacity
903 curves as objective functions. The resulting objective functions have maxima
904 at the operating points, and have slopes away from the maxima correspond-
905 ing to tolerances for error. Since this derivation of objective functions applies
906 to arbitrary capacity curves, it can be applied to model a variety of intero-
907 ceptive modalities.

908 Since any given capacity curve (such as the one in Figure 1 above) reaches
909 a maximum value on the vertical axis, it can be divided by its maximum Y-
910 value to “normalize” it to range between zero and one. Once normalized in
911 this way, it can be interpreted as a cumulative distribution function (CDF)
912 from probability theory. This is in fact precisely what Srinivasan et al. [118]

Figure 7: The probability density function (PDF) corresponding to the capacity curve shown in Figure 1, with the tick-marks delineating threshold (left) and saturation (right) values. Probability density here corresponds to responsiveness to perturbation, not to an empirical frequency of events.

913 did to interpret the firing of retinal neurons in flies as a form of predictive
 914 coding (for an excellent example, see their Figure 1). The derivative (instantaneous slope) of a CDF yields a probability density function (PDF). This
 915 is the more familiar way of representing a probability distribution, where
 916 height on the vertical axis corresponds to likelihood, but for a PDF derived
 917 from a capacity curve, it represents relative responsiveness to perturbation.
 918 We will call such a distribution a [reference distribution](#).
 919

920 Figure 7 therefore shows the result of normalizing and differentiating the
 921 capacity curve in Figure 2 to obtain its corresponding PDF. The density function's graph clearly shows that the baroreflex's capacity to adapt to changes
 922 in mean arterial pressure (MAP) degrades the further MAP moves from the
 923 peak at 100 mmHg. The rate at which it degrades, and the response elicited,
 924 is governed by the baroreflex gain. The gain of a capacity curve thereby
 925 defines the relative importance of deviations from the operating point, and
 926

Figure 8: Logarithm of the probability density shown in Figure 7, with the operating point shown as a diamond marker and tick-marks delineating threshold (left) and saturation (right) values. Probability density here corresponds to responsiveness to perturbation.

thus corresponds neatly to precision in predictive coding. The greater the baroreflex gain, the more sharply curved the PDF around its operating point, and the greater the response mobilized by any deviation from the operating point.

Section 2.3 described how the capacity curves in settling-point physiological controllers will often have inflection points that lie somewhere other than the center, thus being horizontally asymmetrical. This property translates neatly to probability densities: PDF's need not be horizontally symmetrical either. Bodily responses such as inflammation or nociception could have highly asymmetrical capacity curves, with the operating point even potentially being located near a zero value of the PDF.

For analytical convenience, we take the practice from Todorov [119] of using a log-density (the logarithm of a PDF) as an objective function, rather than the PDF itself. The operating point then continues to appear as a local maximum, while the gain determines the cost of movement away from

the operating point, or the value of movement towards it. Figure 8 shows the objective function (log-density) corresponding to the original baroreflex capacity curve, and the formal model below will use such log-densities as objective functions for optimal control. Figure 9 further clarifies the relationship between the various forms of capacity curve by showing all three alongside each-other. Assuming that the brain’s internal model takes the form of a **generative model** (here with discrete time), the internal states of this generative model can be numbered as $\mathbf{x}_t$ for natural numbers (that is, discrete counting numbers) $t \in \mathbb{N}$, with some number L of levels of model hierarchy. We will later define a specific graphical model that suits these notations.

We define the PDF corresponding to a physiological capacity curve as the reference distribution for the value on the horizontal axis of the PDF. These reference distributions have parameters, which we name $\rho_t^{(0:L-1)}$, and $p(\mathbf{x}_t; \rho_t^{(0:L-1)})$ denotes their PDF’s. $\log p(\mathbf{x}_t; \rho_t^{(0:L-1)})$ will denote the corresponding log-density objectives.

Our allostatic decision-making model focuses on optimizing a construct called the *instantaneous capture rate* (given informally in Equation 5), which we write as $J(\mathbf{x}_t, \mathbf{x}_{t-1})$. Since the model here works with discrete time steps, the instantaneous capture rate is defined as a function of the transition between one time-step and another. The instantaneous capture rate consists of a rate of resource intake minus a rate of effort expenditure. We suggest that we can identify the instantaneous capture rate with a rate of movement along capacity curves, which we will write mathematically below in Equation 6.

$$J(\mathbf{x}_t, \mathbf{x}_{t-1}) = \text{Intake}(\mathbf{x}_t, \mathbf{x}_{t-1}) - \text{Effort}(\mathbf{x}_t, \mathbf{x}_{t-1}). \quad (5)$$

The single-step capture rate is formalized as

$$J(\mathbf{x}_t, \mathbf{x}_{t-1}) = \log p(\mathbf{x}_t; \rho_t^{(0:L-1)}) - \log p(\mathbf{x}_{t-1}; \rho_{t-1}^{(0:L-1)}). \quad (6)$$

Equation 6 shows the difference between the log-density objective across two consecutive time steps. Since the reference distributions defining the log-density terms (at the lowest level, $\rho_t^{(0)}$) correspond to physiological capacity curves, movement towards an operating point will contribute positively to this equation, while movement away from an operating point will contribute negatively. Increases in responsiveness of controlled processes that handle

(a) The baroreflex capacity curve

(b) Probability density implied by the baroreflex capacity curve

(c) Logarithm of the probability density implied by the baroreflex capacity curve

Figure 9: The complete set of functions employed to model the baroreflex response. The second comes from taking the derivative of the first and normalizing it, and the third comes from taking the logarithm of the second.

973 metabolic inflow (such as glucagon) can be interpreted as “intake”, while de-
 974 creases in responsiveness of controlled process that handle metabolic outflow
 975 (such as insulin) can be interpreted as “effort”.

976 Section 4.2 below will present evidence for how the instantaneous capture
 977 rate is aggregated over time in foraging behaviors, particularly in feeding
 978 behaviors and the energetic efforts undertaken to enact them.

979 4.2. *Optimal foraging theory suggests a functional form for allostatic control*

980 Mathematically, combining momentary reference signals into a value func-
 981 tion requires first writing those reference signals as objective functions, and
 982 then combining them into a long-run functional form. The subsection above
 983 derived a form of objective function that expresses movement towards and
 984 away from the operating point of a capacity curve. This subsection will iden-
 985 tify a long-run functional form for allostatic decision-making that is based
 986 upon experimental findings.

987 Neural and ecological evidence supports the hypothesis that animals’ in-
 988 take and outflow optimize a long-run functional form called the global capture
 989 rate [120, 121]. Ecologists such as Stephens and Krebs [120] define the global
 990 capture rate $\bar{J}$ as the sum of all metabolic intake minus all effort expended,
 991 over the total time devoted to a behavior. Meanwhile, neuroscientists have
 992 used the global capture rate to help relate dopaminergic neuronal signaling
 993 [122] and animal behavior [123] to the discounting of rewards in decision-
 994 making tasks.

This formalism assumes a given behavior has a finite length in time of T ,
 and that all intakes and efforts are zero at time $t = 0$. The global capture
 rate is then written

$$\bar{J}(\mathbf{x}_{1:T}) = \frac{\sum_{t=1}^T J(\mathbf{x}_t, \mathbf{x}_{t-1})}{T}. \quad (7)$$

995 The global capture rate is thus the average over time of the individual “intake
 996 minus effort” instantaneous capture rates of Equation 6 [124]. Since the
 997 global capture rate is defined by dividing by time, it exists for any length of
 998 time, long or short. In the real world, intakes or rewards often only accrue
 999 at the conclusion of a behavioral episode, while efforts or costs necessarily
 1000 accrue throughout the behavior as energy is spent. Averaging over time
 1001 treats rewards and costs equally whenever they occur during a behavioral
 1002 episode. The global capture rate helps make behaviors commensurable, even

1003 when they take different lengths of time or accumulate rewards at different
 1004 points in time.

1005 Defining the global capture rate with respect to an internal model $p(\mathbf{x}_{1:T} |$
 1006 $T)$ entails generating trajectories of length T , conditioned upon an initial
 1007 state $\mathbf{x}_0$ as context,

$$\tilde{J}(\mathbf{x}_0) = \lim_{T \rightarrow \infty} \mathbb{E}_{p(\mathbf{x}_{1:T}|\mathbf{x}_0)} \left[\frac{1}{T} \sum_{t=1}^T J(\mathbf{x}_t, \mathbf{x}_{t-1}) \right]. \quad (8)$$

1008 This is an indefinite-time form of the global capture rate that generalizes
 1009 Equation 7 from settings with a known beginning and end to arbitrarily long
 1010 time scales.

1011 4.3. Feedforward control with generative action concepts

1012 This subsection will consider a formal model in discrete time, with ar-
 1013 bitrary state-spaces and sensory observations. Specifically, we write sensory
 1014 observations as indexed by discrete timesteps, t in o_t , and do similarly for
 1015 other variables. The formal model here incorporates the hypothesis that the
 1016 brain’s internal model stretches across a hierarchy of time-scales [125]. How-
 1017 ever, we restrict our description to $L = 4$ hierarchical “levels” (matching the
 1018 sensory surfaces, somatosensory and exteroceptive state estimators, intero-
 1019 ceptive state estimators, and allostatic capacity estimators in Figure 6). We
 1020 do not hypothesize a one-to-one mapping between functional structure and
 1021 neuroanatomy, and so the structure we have given here may not match neu-
 1022 roanatomy – but, the graphical model is formulated with an eye to keeping
 1023 it consistent with the biological details of allostasis. We consider it a first
 1024 attempt to build a functionally coherent model that can then be critiqued
 1025 and refined to match anatomical and empirical evidence.

1026 The complete state of the generative model at a specific timestep t is
 1027 written symbolically as

$$\mathbf{x}_t = (o_t, s_t^{(1:L)}, a_t, \rho_t^{(0:L-1)}), \quad (9)$$

1028 under the assumption that outcomes o_t are observed, latent states $s_t^{(1:L)}$ track
 1029 the unobservable state of the body, and $\rho_t^{(0:L-1)}$ parameterize reference dis-
 1030 tributions used below in Equation 6. $\rho_t^{(0)}$ (as above) parameterizes a ref-
 1031 erence distribution for o_t , and a_t models the closed-loop control action of

1032 motor reflexes; $\rho_t^{(1:L-1)}$ parameterize reference distributions for $s_t^{(1:L)}$ respec-
 1033 tively. $s_t^{(L)}$, as the highest level latent state, has no reference distribution.
 1034 The factorization of control into hierarchical levels $\rho_t^{(0:L-1)}$ is based upon the
 1035 referent-configuration control scheme from Section 3.2, so as to match as
 1036 closely as possible what is known about somatomotor control.

1037 A behavioral trajectory is written as contextualized by (conditioned upon)
 1038 an initial state $\mathbf{x}_0$. This initial state corresponds to the beginning of a behav-
 1039 ioral episode, within which interoceptive outcomes will be considered com-
 1040 mensurable. The following states from time 1 until time T , sampled from a
 1041 generative model p_θ with parameters θ , are then written as

$$\mathbf{x}_{1:T} \sim p_\theta(\mathbf{x}_{1:T} \mid \mathbf{x}_0), \quad (10)$$

$$\bar{J} \approx \frac{1}{T} \sum_{t=1}^T J(\mathbf{x}_t, \mathbf{x}_{t-1}), \quad (11)$$

1042 along with the corresponding approximate global capture rate. The internal
 1043 model simulates a population of potential behavioral trajectories $\mathbf{x}_{1:T}$, and
 1044 uses them to estimate the global capture rate $\bar{J}$ for the simulated behavior.
 1045 This quantity will play a role in the feedback control formalism described
 1046 later. Figure 10 summarizes the proposed model structure as a [probabilistic](#)
 1047 [graphical model](#)¹⁷, with $L = 4$ for purposes of exposition.

1048 Table 1 summarizes the key notations, including the random-variable
 1049 names and time indices. C.1 specifies the various distributions for our graph-
 1050 ical model and control problem in detail.

1051 We do not claim that our graphical model accurately captures the anatomy
 1052 of the brain (or the function of the entire brain). It captures reference-based
 1053 sensorimotor control across a hierarchy of timescales, a single feature shared
 1054 with the brain [125]. There can be no one-to-one mapping between computa-
 1055 tional theories and neuroanatomical findings [127, 128]. We instead account
 1056 here for what we already know about the brain that goes into our model, with-
 1057 out assuming any one-to-one mapping of brain structure and function. In the
 1058 actual brain, the hierarchy of timescales is also a hierarchy of dimensionali-
 1059 ties: information is more concrete and higher-dimensional at the bottom of
 1060 the hierarchy, more abstract and lower-dimensional at the top [129]. Dimen-

¹⁷For a primer on graphical models see Koller and Friedman [126].

Figure 10: A hierarchical generative model capturing multiple timescales and reference distributions at each level. Without addressing empirical questions about neural hierarchies, here we employ a model with $L = 4$ levels to match Figure 6. For $l \in [1..L - 1]$ each $s_t^{(l)}$ node denotes an unobserved latent state, and each $\rho_t^{(l-1)}$ represents parameters of a reference distribution for $s_t^{(l-1)}$. o_t represents observed sensory outcomes, and a_t represents the closed-loop control actions generated by motor reflexes. Arrows between random variables denote conditional dependencies. Arrows stretch further to the right when they denote change over longer time scales.

1061 sionality reduction up the hierarchy entails compression up the hierarchy, so
 1062 that state estimates in the higher-level variables have greater precision. Evi-
 1063 dence from machine learning shows that learning to control a system in terms
 1064 of a low-dimensional compressed representation (also called a “latent space”)
 1065 provides greater performance with fewer episodes of experience compared to
 1066 controlling the same system in terms of its raw measurements [130, 131, 132].
 1067 Since the visceromotor, premotor, and motor cortical areas that implement
 1068 action concepts employ low-dimensional, compressed multimodal summaries

t	Discrete time-step index
o_t	Sensory observations (external and visceral)
$s_t^{(1:L)}$	Unobserved states constructed by the internal model
$\rho_t^{(0:L-1)}$	Parameters to a reference distribution
a_t	Control actions emitted by peripheral reflexes
$\mathbf{x}_t$	A complete model state for time t
p_θ	Feedforward generative model with parameters θ
q_ϕ	Feedback generative model with parameters ϕ

Table 1: Random variable names used in our APIC model

1069 rather than raw measurements [133], we conjecture that they may benefit
1070 from the sample efficiency of latent-space control approaches. Finally, com-
1071 putational modeling has been directly applied in experiments that supported
1072 the presence of hierarchical Bayesian models in interoception [134, 135, 136].

Terminology

The language of generative models in the brain is typically used in the context of predictive processing approaches to brain function. These approaches typically label as a “generative” model the processes producing efferent signals in sensory areas of the cortex, while suggesting that the processing of reafferent signaling in those same areas encodes a “recognition” model [137]. The efferent and reafferent signals themselves are typically then labeled “predictions” and “prediction errors”. We will continue to use “feedforward” for efferent signals (or computations) and “feedback” for reafferent signals (or computations).

1073

1074 Stochastic optimal control with a generative model requires (by defini-
1075 tion) solving or approximating the recursive optimization problem called the
1076 [Bellman equation](#). For a time-averaged problem, the Bellman equation will
1077 contain a term for the objective function $J(\mathbf{x}_t, \mathbf{x}_{t-1})$ at a particular state, a
1078 term $\bar{J}(\mathbf{x}_0)$ for the (estimated) global capture rate of a behavior as a whole,
1079 and a recursive term. It is written as a definition of the optimal value function

$$\tilde{V}^*(\mathbf{x}_t, \mathbf{x}_{t-1}) = J(\mathbf{x}_t, \mathbf{x}_{t-1}) - \bar{J}(\mathbf{x}_0) + \max_{p_\theta} \mathbb{E}_{\mathbf{x}_{t+1} \sim p_\theta(\mathbf{x}_{t+1}|\mathbf{x}_t)} \left[\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right]. \quad (12)$$

1080 However, this formalism only describes how to plan optimal behavior or
1081 learn an optimal policy (a mapping from states to actions) through forward
1082 simulation. This equation does not describe how to integrate afferent sensory
1083 information as control feedback; it also imposes the great computational
1084 difficulty of finding exact solutions to a recursive optimization problem. The
1085 following subsection will discuss a way to tackle both of these limitations,
1086 yielding a more neurally plausible form of stochastic optimal control.

1087 4.4. *Feedback control with generative action concepts*

1088 Standard theories of optimal decision-making ignore the variability of
1089 choice outcomes, while evidence shows that human behavior takes the level
1090 of risk into account [138, 139]. This makes normative sense from the per-
1091 spective of embodied action: the brain has to transform even seemingly clear
1092 and simple decisions (“reach to the left”) into noisy, high-dimensional motor
1093 movements (“change these and those muscle spindles’ referent lengths and
1094 sensitivities”). Risk and uncertainty remain part of a movement even once
1095 the body has actually enacted it, since the distal body and world remain only
1096 partially observable via the sensory surfaces. This subsection will consider
1097 what these facts imply for the combined problem of decision-making and
1098 motor control faced by the brain. After these considerations, this subsec-
1099 tion will give a combined formulation of risk-sensitive decision making and
1100 feedback-stabilized motor control.

1101 Neural firing contains stochastic noise [140], and so finding an exact max-
1102 imum of a difficult recursive problem quickly seems implausible. Even if the
1103 brain could find the exact maximum quickly, noise in the motor system im-
1104 poses its own cost factors into the decision value [141]. Experimentally, hu-
1105 man and animal behavior show meaningful variability [142] across every level
1106 of task behavior, while Equation 12 says that any given task ought to cor-
1107 respond to a unique optimal way to act. Behavior also displays meaningful
1108 variation attributable to predictive uncertainty in people’s internal models,
1109 both about the task goals themselves [143] and about appropriate sensorimo-
1110 tor strategies [144]. A risk-neutral theory can explain neither of these effects,
1111 since the Bellman equation uses expectations to average away model-based
1112 uncertainty.

1113 In fact, when sensory signals carry only partial and ambiguous informa-
1114 tion about the state of the body and environment (as they always do in the
1115 real world), Equation 12 will have to be solved anew every time an $\mathbf{x}_t$ is
1116 updated based upon sensory feedback. This is because the equation does

not take into account the difference between planning ahead with a [forward model](#) and [feedback control](#) based on reafferent sensory signals.

Here we propose instead a feedback-control formalism which takes this difference into account. This formalism quantifies the deviation of the closed-loop, feedback-stabilized behavior from the open-loop, feedforward action concept. It balances this deviation with the present and future allostatic value of an online behavior to determine what to do. The formalism begins by updating the state estimate for each time-step t as the sensory signals o_t become available. A probability model q_ϕ with parameters ϕ , conditioned upon the sensory observations and initial state, then acts as a feedback controller¹⁸ in updating state estimates and motor references

$$\mathbf{x}_{1:T} \sim q_\phi(s_{1:T}^{(1:L)}, \rho_{1:T}^{(0:L-1)} \mid o_{1:T}, \mathbf{x}_0). \quad (13)$$

The updated state estimates, being based upon the observations o_t , may significantly differ from the state estimates emitted by the feedforward action concept. This deviation imposes a penalty term in the objective function, called an information divergence (specifically, the Kullback-Leibler or KL divergence). The resulting function

$$\mathcal{J}(\mathbf{x}_t, \mathbf{x}_{t-1}) = J(\mathbf{x}_t, \mathbf{x}_{t-1}) - D_{\text{KL}}(q_\phi(\mathbf{x}_{t+1} \mid \mathbf{x}_t) \parallel p_\theta(\mathbf{x}_{t+1} \mid \mathbf{x}_t)), \quad (14)$$

trades off between the allostatic responsiveness objective of Equation 6 and adherence to the “planned” population of predictions given by the feedforward action concept. The first term measures success at physiological regulation, while the second term penalizes the feedback controller q_ϕ for deviating from the action concept p_θ or suffering sensory prediction errors. When movements result in sensory outcomes very close to those predicted under the feedforward action concept p_θ , a low information divergence, that action concept can be considered robust in the face of sensory feedback.

Readers familiar with predictive processing [18, 145, 146] and active inference [111, 137] will recognize the form of the above equation as a negative free-energy or a variational lower bound (see Millidge et al. [147] for discussion on the variety of such bounds). Such objectives can typically be written

¹⁸Or equivalently within this context, an approximate posterior distribution over actions

1145 out and interpreted in several equivalent ways, each of which can come with
 1146 its own intuitions. There are arguments for the computational [148] and
 1147 thermodynamic [149] efficiency of minimizing this specific divergence in the
 1148 course of neural processing, but to date the available evidence does not rule
 1149 out other, more complex information divergences for penalizing feedback cor-
 1150 rection of movements.

The long-run value function (Equation 12) for a control problem can be used to derive an equation for the optimal feedback controller. By treating the pre-planned action concept as a kind of probabilistic “prior belief” about the behavioral trajectory and the decision-value of the behavior as a “likelihood function” linking the action concept to the objective function in Equation 14, Bayes’ rule (see Appendix C for details) will yield the optimal feedback controller:

$$q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t) = \frac{\exp\left(\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t)\right) p_{\theta}(\mathbf{x}_{t+1} \mid \mathbf{x}_t)}{\mathbb{E}_{\mathbf{x}_{t+1} \sim p_{\theta}(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\exp\left(\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t)\right) \right]}. \quad (15)$$

1151 Equation 15 links pre-planned action concepts (the “priors” $p_{\theta}(\mathbf{x}_{t+1} \mid \mathbf{x}_t)$) to
 1152 online feedback (the “likelihood” $\exp\left(\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t)\right)$) in a probabilistically
 1153 optimal way (by treating them both as densities, multiplying, and normaliz-
 1154 ing the result). A system that can represent and simulate from this equation
 1155 can (approximately) predict the way that an “optimal agent” (with the as-
 1156 sumed objective function) ought to move.

1157 Taking a generative modeling perspective, the controller defined in Equa-
 1158 tion 15 treats the original action concept p_{θ} as a prior, and conditions on
 1159 the long-run value of the potential future state $\mathbf{x}_{t+1}$ as a selection criterion.
 1160 It is therefore risk-sensitive and information-seeking. The term that would
 1161 typically correspond to model evidence now corresponds to the expected ex-
 1162 ponentiated decision value of the present state $\mathbf{x}_t$.

Substituting the augmented objective ($\mathcal{J}$, Equation 14) and the analytical expression for the optimal feedback controller (above, Equation 15) into the Bellman equation (Equation 12) will yield

$$\tilde{V}^*(\mathbf{x}_t, \mathbf{x}_{t-1}) = \mathcal{J}(\mathbf{x}_t, \mathbf{x}_{t-1}) - \bar{J}(\mathbf{x}_0) + \mathbb{E}_{q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right]. \quad (16)$$

Now the intractable recursive term in the equation is $\mathbb{E}_{q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right]$. Finding a way to replace this term will yield a more (computationally)

tractable problem. The information divergence term in $\mathcal{J}$ provides just such a way, since it can be written as precisely the difference between the intractable “hard” maximum under q_ϕ^* and the more tractable “smooth maximum” under the preplanned action concept. In symbols

$$-D_{\text{KL}}(q_\phi^*(\mathbf{x}_{t+1} | \mathbf{x}_t) \| p_\theta(\mathbf{x}_{t+1} | \mathbf{x}_t)) = -\mathbb{E}_{q_\phi^*(\mathbf{x}_{t+1} | \mathbf{x}_t)} [\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t)] + \log \mathbb{E}_{p_\theta(\mathbf{x}_{t+1} | \mathbf{x}_t)} [\exp(\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t))], \quad (17)$$

and so the penalty’s first term, when substituted into Equation 16, will cancel the intractable recursion. Only the second term of the penalty will remain, yielding a smooth maximization problem across whole action trajectories. The equation can then be solved (once again, see Appendix C for details) to write the value function without recursion as

$$\tilde{V}^*(\mathbf{x}_0) = \log \mathbb{E}_{q_\phi^*(\mathbf{x}_{1:T} | \mathbf{x}_0)} \left[\exp \left(\sum_{t=1}^T \mathcal{J}(\mathbf{x}_t, \mathbf{x}_{t-1}) - \bar{J}(\mathbf{x}_0) \right) \right], \quad (18)$$

$$= \log \mathbb{E}_{p_\theta(\mathbf{x}_{1:T} | \mathbf{x}_0)} \left[\exp \left(\sum_{t=1}^T J(\mathbf{x}_t, \mathbf{x}_{t-1}) - \bar{J}(\mathbf{x}_0) \right) \right]. \quad (19)$$

This function contains an exponentiated sum over timesteps, where each timestep’s addend has the form of an [advantage function](#),

$$A(\mathbf{x}_t, \mathbf{x}_{t-1}; \mathbf{x}_0) = \mathcal{J}(\mathbf{x}_t, \mathbf{x}_{t-1}) - \bar{J}(\mathbf{x}_0). \quad (20)$$

1163 This measures the *relative* decision value of transitioning into state $\mathbf{x}_t$ relative
 1164 to an estimated global capture rate $\bar{J}(\mathbf{x}_0)$ for the initial context $\mathbf{x}_0$. Viewed
 1165 another way, it measures how well the ongoing feedback-controlled behavior
 1166 has performed, relative to the predicted average for the prior action concept.

Terminology

A sum of a function across timesteps, under a probabilistic expectation, is often called a “path integral”. Since the value function in APIC can be expressed in terms of a path integral, APIC falls within the class of SOC methods known as *path-integral control* [150].

1167

Jensen’s inequality, commonly employed in predictive processing, allows moving the logarithm inside the expectation, at the cost of yielding a lower

bound to the optimal value function rather than an expression for it. Doing so cancels the exponential function inside the expectation. The lower bound is written in terms of an arbitrary feedback controller q_ϕ and the sum of advantage values obtained over the course of the behavior,

$$\tilde{\mathcal{V}}_{\theta,\phi}^* = \mathbb{E}_{q_\phi(\mathbf{x}_{1:T}|\mathbf{x}_0)} \left[\sum_{t=1}^T A(\mathbf{x}_t, \mathbf{x}_{t-1}; \mathbf{x}_0) \right] \leq \tilde{V}^*(\mathbf{x}_0). \quad (21)$$

It may seem obvious that any behavioral returns are less than or equal to the best possible returns. The implication of having a proper lower bound, however, is that any process for maximizing the value lower bound $\tilde{\mathcal{V}}_{\theta,\phi}^*$ does in fact maximize the optimal value function $\tilde{V}^*(\mathbf{x}_0)$ by proxy. This includes the kind of computations¹⁹ which predictive processing theorists posit the brain can in fact perform [151, 146] to incrementally improve its action concepts p_θ and its feedback controllers q_ϕ .

This concludes the description of the Allostatic Path-Integral Control (APIC) model. Since APIC employs action concepts and context states $\mathbf{x}_0$, it only requires planning and adjusting behavior in context (e.g. from timesteps 1 to T), despite optimizing a “global” (indefinite) capture rate (Equation 8). Since it employs “smooth” maximization rather than a “hard” recursive maximization, it can accommodate the sensitivity of behavior to risk. APIC follows in the tradition of modelers such as Belousov et al. [152], Mitchell et al. [153], and Piray and Daw [154], who consider the specific problem of making embodied decisions when sensory feedback contains noise and a ground-truth model of task dynamics is not available.

Summary. This subsection has detailed a formal model, called the Allostatic Path-Integral Control (APIC) model, for how the brain can realistically achieve allostatic regulation of the body in an online setting. APIC assumes that the brain starts with an action concept describing a potential behavior, and tries to maximize the allostatic returns on that behavior while keeping the online (feedback-stabilized) behavior close to the original plan. Incorporating an action concept, and penalizing deviation from it, provides an explicit expression for the optimal feedback controller. The infinite-horizon, average-objective setting for this stochastic optimal control model captures the time-averaging behind the global capture rate (i.e. Equation 7). This

¹⁹Gradient updates, to be precise.

1195 model can take advantage of neural stochasticity to optimize an objective
1196 function defined over a hierarchy of scales of space and time, allowing for
1197 both high-level and low-level behavioral control.

1198 *4.5. Summary*

1199 This section derived an objective function, formal problem setting, and
1200 formalism for allostatic control based upon the paradigm of path-integral con-
1201 trol. Section 4.1 connected physiology’s ever-shifting capacity curves (such as
1202 Figure 1) to probability density functions (shown in Figure 7), and used that
1203 connection to define an objective function. Section 4.2 described a problem
1204 setting for decision making that accounts for much available evidence about
1205 how animals make allostatic decisions in ecological settings. Section 4.3 then
1206 sketched the formal definitions needed to apply path-integral control to a
1207 generative model; gave a generative model (Figure 10) with the hierarchi-
1208 cal structure that the literature suggests is found in the brain; and derived a
1209 stochastic optimal control formalism based on simulations of potential futures
1210 by an internal model. Section 4.4 then shifted the perspective to feedback
1211 control, obtaining a risk-sensitive and tractable formalism.

1212 The next section will conclude the paper by discussing the implications
1213 of applying stochastic optimal control theory to allostatic physiological regu-
1214 lation in general, and the specific hypotheses for the brain implied by APIC.

5. Discussion: Interoception stabilizes action and constructs allostatic references

If the brain is an allostatic regulator, then its most basic task is to anticipate the body’s physiological needs and prepare to meet them before they arise. This paper provided a unified interpretation of allostasis in terms of brain-body interaction and neural computations. This unified interpretation is built around stable circular interactions between the brain and body, and control theory has provided language with which to describe the mechanisms stabilizing those interactions.

Each section of this paper addressed a particular circular interaction, and below we connect each with its implications for psychological investigation. Section 5.1 will discuss what the Allostatic Path-Integral Control (APIC) model in Section 4 implies about brain function. Section 5.2 will discuss what control theory brought to the study of physiology, motor control, and allostatic decision-making (with reference to Section 3). Section 5.3 will discuss interoception and capacity curves in light of Section 2’s overview of interoception. Each subsection will end with a paragraph giving specific predictions made by our view of allostasis in the brain and body.

5.1. Viewing the brain as an allostatic optimal controller

The Allostatic Path-Integral Control (APIC) model in Section 4 implies a number of specific hypotheses, beyond those generic to stochastic optimal control. This subsection will situate those commitments within the broader literature on formal modeling of motor and decision functions. We will first describe commitments shared with other modeling approaches, then describe less common commitments, and finally we will describe several commitments that are unique to the APIC model.

On the theoretical side, the APIC model shares a number of modeling choices with active inference models, perhaps enough for APIC to be considered an active inference model of sorts. APIC employs (normalized) probability densities as its objective function to provide a “common currency” for different “rewards” and “costs” [155, 64, 156, 157, 147, 158]. As in active inference models based on the free-energy principle (such as Stephan et al. [159]), the objective in our APIC model is a function of (among other things) *precision* terms, which specify the relative worth of a change in one variable versus another. APIC and active inference both optimize motor behaviors via a variational lower bound on a long-run objective, and APIC is likely

compatible with the neural process theories [160] that have been forwarded to ground active inference in the brain. These commitments place APIC alongside active inference, in contrast to most reinforcement learning models of decision-making (e.g., Niv [104]).

APIC also shares a number of features common to other predictive processing paradigms, beyond active inference. Most importantly, it uses a probabilistic, generative internal model. There is a broad family of predictive processing approaches to neural function, and a proposal for a neural implementation for APIC could depend on the neurocomputational details of any of them (e.g. Boerlin et al. [161], Spratling [162], Kadmon et al. [163]). These details are beyond the present scope, but, some important features can be highlighted. Predictive processing models must construct potential bodily movements from some distribution of possibilities, and neural evidence [164, 165] suggests that neural representations may not summarize probability distributions in a small, fixed number of parameters. Therefore, to understand how movements are constructed, we must understand the entire distribution, and not just its summary statistics. APIC can accommodate this, and is therefore compatible with neural process theories of predictive processing not based on sufficient statistics²⁰.

Unlike most work in predictive processing, APIC is designed from first principles to solve a combined problem of both decision making and motor control, rather than to address perceptual or cognitive problems²¹. Unlike other predictive processing models, in APIC the objective is optimized by making capacity curve precisions small (implying greater resilience to challenge) rather than large (implying reduced resilience to challenge). The most elementary of the APIC model’s first principles are: the body, when considered as a set of coordinated systems, can only perform one action at a time, whereas the brain can imagine many possible actions. The study of embodied decision-making in neuroscience begins from these assumptions [38] and an emerging body of evidence suggests that the brain controls the body by keeping many possible actions “in the running” until sensory feedback forces an irreversible decision [167, 168, 99, 24]. The APIC model works this way by default, thanks to its roots in path-integral control [150].

²⁰See Sanborn [166] for a review of sampling and variational approaches to approximate inference in the brain.

²¹See Ramstead et al. [137] for a counterexample.

Finally, the APIC model implies some hypotheses that are (to our knowledge) unique. These derive largely from the APIC model’s combination of the neuroanatomy of Barrett [15] and the feedback control formalism of Thijssen and Kappen [169]. Active inference modeling typically interprets q_ϕ as an approximate posterior belief. APIC re-conceptualizes q_ϕ as a feedback controller. This interpretation comes from the path-integral control (PIC) literature [169, 170, 171], whereby the same mechanism can both stabilize movements (keeping q_ϕ close to the original action concept p_θ) and maximize decision value. APIC’s unification of inference with feedback control handles uncertainty about the consequences of movements (optimal feedback control); uncertainty about the state of the body and the world (predictive processing); and uncertainty about the reference or objective of action (reinforcement learning) in a single framework. The first term of APIC’s objective function is physiological: it represents the responsiveness of controlled processes in the body, as inferred through interoception. The second term of APIC’s objective function separately quantifies the difference between planned and actual movements (including visceromotor movements).

How would an implementation of APIC map onto the brain? We did not map functional descriptions in control theory terms to specific anatomic or functional assemblies in the brain, but nevertheless, research does suggest that brain systems can be described as enacting internal modeling, feedback control, and decision making. For example, the hippocampal-entorhinal complex is often considered to construct a predictive “cognitive map” [172] out of sequential episodes [10], which helps evaluate a value function for a control problem [173]. The cerebellum may help exploit reafferent and exafferent sensory prediction errors as online corrective feedback for movements [174, 175], although the role of the cerebellum in visceromotor control requires more study to understand the mechanisms involved. Regions in the brain’s default mode network are thought to help construct state estimates (e.g., Buckner [176], Barrett [15]) and implement allostatic control [26]. Evidence for common functional gradients across the cerebral cortex, cerebellum, and hippocampus suggests that future empirical studies should apply common computational paradigms to different brain regions [177].

Predictions. The APIC objective function proposes a theory of how behavioral reinforcement and allostatic control interact. This theory has implications for how a primary (unlearned) reinforcer would arise. According to APIC, a primary reinforcer arises when the dynamics of a controlled

physiological process align with those of its capacity curve (as a function of an underlying regulated resource). If a primary reinforcer arose directly from a regulated resource, then the specific objective-function formulation we have suggested would be falsified; however, this would not necessarily falsify APIC as a computational unification of decision making and motor control for action concepts. Likewise, only the objective-function would be falsified if future work identified a modular “reward” or reinforcement systems in the brain—i.e. a reward system that operates separately from decision making and motor control. Alternatively, empirical tests for controlled variables [178, 179] could determine which physiological parameters (if any) elicit global, behavioral responses to disturbance.

5.2. The body and brain through the lens of control theory

Section 3 introduced the concepts of control theory by describing their applications in peripheral physiology, motor control, and decision-making. This subsection first summarizes the theoretical “point of view” obtained from control theory, and then reviews several of the key ways in which control theory can be used to clarify brain-body interactions.

Control theory provides a way to conceptualize how physiological systems can function reliably as a whole, despite being built from unreliable parts. The dominant mechanism used by control theorists to ensure such stability is *feedback*, in which measurements flow from the underlying system being regulated (the plant) to the system doing the regulating (the controller). This constant flow of information from the plant to the controller helps the controller move the plant according to a desired trajectory, expressed physically as a reference signal. Figure 6 shows the process of allostatic control in the body using the language of control theory.

Voluntary skeletomotor movements are controlled by central commands from the brain. The brain signals movements and bodily postures in terms of proprioceptive referent configurations, which then constrain the motor reflexes to minimize the difference between the reference and the actual proprioceptive signal. This form of *referent control* extends all the way up through the brain, and in the form of posterior state estimates (see Figure 6) provides an explanation for reafferent connections from sensory to motor areas of the cerebral cortex. We have hypothesized in this paper that referent control is shared between skeletomotor and visceromotor modalities²². Visceromotor

²²Note that this does *not* imply that people will have a sense of agency in controlling

reflexes, such as the carotid baroreflex, are responsible for controlling blood pressure in response to central commands from the brain. These visceromotor reflexes may use feedback control mechanisms that are similar to those used by skeletomotor movements. This means that visceromotor reflex circuits, like the baroreflex, may only provide a single kind of feedback: the stabilization of movements, i.e., via the short loop between distal physiology, sensory surfaces, and peripheral reflex controllers (as in stretch reflex proprioceptors). Figure 6 shows how visceral sense data ascend, in effect, to become exafferent interoceptive prediction errors, which combine with efferent interoceptive predictions to generate posterior state estimates. From the posterior interoceptive state estimates, the brain can estimate the allostatic capacities of the many controlled processes in the body. Optimizing the responsiveness of these allostatic capacities then finally “closes the loop” and provides a reference signal, constraining the central command signals from the brain to the motor reflexes.

Predictions. Above we have hypothesized that the visceral nervous system (including visceromotor efferents, interoceptive afferents, and the autonomic nervous system) operates, in its motor operations, via referent control. This would imply that interoception can be anatomically segregated into motoneuron-interneuron-interoceptor reflexes that stabilize visceral “movements” (as in somatosensory modalities), and non-motor interoceptive endpoints, which serve another functional role. Rather than stabilizing a (visceral) “movement” that is constrained from the top down by central command, we suggest these visceral sense data may instead constrain the central commands themselves. Alternatively, empirical tests for controlled variables [178, 179] could determine which physiological parameters elicit local, reflexive responses to disturbance.

5.3. Interoception and capacity curves

Section 2 introduced much of the physiology fundamental to this paper. It provided the distinction between regulated resources and controlled processes, and then described interoception as the predictive internal modeling of such variables. The section then described how capacity curves can quantify both the relations between variables and the objective of allostatic regulation. This subsection first clarifies the distinction between the actual state

their autonomic nervous systems, as they do in skeletomotor control.

of the body and the interoceptive predictions constructed by the brain, and then elaborates on how that distinction affects the behavioral construct of “reward” in psychology.

This paper maintained a careful distinction between visceral sense data, interoceptive predictions, and (re)afferent viscerosensory signals. The distinction must be carefully navigated, because it is not yet known where, neurophysiologically, the distinction between raw viscerosensory data and prediction errors can be made. Most previous work on interoceptive predictive processing has assumed that peripheral interoceptive neurons fire in an entirely stimulus-driven way, without any modulation by descending predictions; this is because, to date, the large majority of research studying peripheral predictive coding has focused on the retina in certain model organisms (e.g., Srinivasan et al. [118], Hosoya et al. [180], Liu et al. [181]). However, some theories and at least one experiment take the other side of the issue. Theories of peripheral predictive coding [22, 182] reason that a neuron’s most metabolically “cheap” responses should represent the predictable stimuli, while “expensive” responses are reserved for the most surprising stimuli. Barrett [15] has suggested that descending visceromotor prediction signals modulate viscerosensory data as it ascends through the brainstem and midbrain. Dworkin [183] demonstrated that certain peripheral interoceptors in humans reduced their afferent firing rate to zero under a constant stimulus; since the stimulus remained the same while neural firing changed, the interoceptors in that experiment could not be entirely stimulus driven. We hope that interoception researchers will invest future experimental effort in interoceptive modalities to differentiate the effect of repetition suppression from the function of predictive coding.

The question of where sense data are converted to prediction errors remains open, but it remains very likely that the brain issues interoceptive predictions, which are constrained and corrected by viscerosensory sense data. The predictive processes taking place in the brain form a model of the innervated viscera, but they can only directly affect the viscera through visceromotor signals. The basic relation between the brain and the world outside itself applies in the interoceptive realm as it does in exteroceptive modalities. The brain can constrain (or influence) the body using motor signals, and actual bodily sense data constrain (or influence) the interoceptive contents of the brain’s internal model.

Assuming that neither noise nor nervous dysfunction prevent the brain from accurately predicting the visceral sense data, we have hypothesized

that the brain’s interoceptive representations contain both state estimates of the viscera and a functional analogue of allostatic capacity curves. The movement of estimated states along capacity curves, and the movement of the capacity curves themselves, then determines the brain’s estimate of allostatic responsiveness. We have hypothesized that increases in such responsiveness could function as “rewards”, positively reinforcing behavioral trajectories, while decreases in responsiveness could function as “costs”, negatively reinforcing behavioral trajectories. Such a hypothesis may account for *alliesthesia* [184] in decision making and behavior, in which the relation between an exogenous stimulus and the body’s internal state determines whether that stimulus positively or negatively reinforces skeletomotor and visceromotor action [185, 15]. This can include actions which alter the capacity curves themselves, such as relaxing the baroreflex gain during exercise to accommodate rising heart rate and blood pressure.

Note that rewards are not always experienced as pleasant (e.g., the removal of something unpleasant to strengthen a behavior, called negative reinforcement) and not all stressors are experienced as unpleasant (e.g., exercise), consistent with the suggestion that approach/avoid features of behavior and pleasant/unpleasant affective features of experience are not always aligned. The degree to which behavioral and experiential “valences” align with each other in a given situation, and even the choice of which specific variables within behavioral and experiential processes to compare, remain largely open questions.

Finally, this paper has taken a physiologist’s view of the innervated body, delineating regulated resources from controlled processes. The “ideal type” capacity curve treats the action of a controlled process as a function of an underlying regulated resource, and both variables receive peripheral interoceptive innervation. For example, glucose, glucagon, and insulin levels in the circulating blood are sensed separately. However, in many cases, the body employs hierarchies of regulatory mechanisms, making some controlled processes act as functions of other controlled processes. The baroreflex falls into this category. Thus, despite being a canonical interoceptive modality for experimentalists, the carotid baroreflex may actually be functionally atypical. The physiological differences between the baroreflex and chemosensing may provide some functional hints. The baroreflex forms an entire reflex circuit, whose visceromotor function under central command may resemble that of a somatomotor reflex (as suggested in Section 5.2 above).

Predictions. Physiological questions about regulated resources versus controlled processes may be amenable to anatomical investigation: where a reflexive control circuit is found, it may reflect the monitoring of a controlled process. Where peripheral interoceptors report “raw” visceral sensations, without an intervening motor reflex circuit, that may reflect the sensing of a regulated resource. The open question then would be how the central nervous system determines which controlled processes act as functions of which regulated resources. Speculating somewhat further, if an anatomical difference existed between viscerosensing of regulated resources and controlled processes, it would also enable inferences about the constraints provided for the brain by those respective viscerosensory endpoints. Interoception of regulated resources *in relation to* controlled processes would then constrain the brain’s representations of the allostatic responsiveness of behavior, its reference signal for movements. Interoception of controlled processes in relation to signaled top-down references would provide feedback to the brain that is analogous to what happens in somatosensory modalities: feedback stabilization of (visceral) movements.

5.4. Conclusions

Let’s return to our amateur dodgeball player. Under the hypothesis laid out in this paper, afferent sensory signaling (including viscerosensory afferents) conveys prediction errors, which update a predictive internal model in the brain of the body on the dodgeball field. Motor processes exploit the updated model contents to improve the player’s performance. Playing dodgeball, and improving one’s performance, requires anticipating the physiological and metabolic demands of the game (which capacity curves can capture) and mobilizing the body’s internal visceromotor systems to meet those demands before they arise.

Of course, within an allostatic view of the brain, the desire to play dodgeball does not arise *ex nihilo*, any more than a “desire” exists to maintain blood glucose within certain ranges. These desires (or “motivations”) are thought to begin as abstract conceptualizations [15] that predict allostatic capacities and interoceptive states (and can be described with low dimensional features, such as reward, stress, threat, and so on), monitoring the controlled processes that constitute the body’s physiology. The APIC model’s regulatory targets consist of ever-moving capacity curves, and so the hypothetical dodgeball player’s brain can both predict the accuracy with which their skeletomotor

1501 movements aim at other players, and (paraphrasing Klein [[186](#)]) be moti-
1502 vated to improve their aim. We posit that this integration of interoception
1503 and allostatic control, played out across nested temporal and spatial scales,
1504 enables the brain to make the decisions and enact the movements that enable
1505 us to score well in dodgeball.

1506 *Acknowledgements.* The authors gratefully acknowledge the support of the
1507 Army Research Institute (grant W911NF-16-1-0191, principal investigator
1508 Quigley), the National Institutes of Health-National Institute of Mental Health
1509 (NIH-NIMH, grant R01MH113234, principal investigators Barrett and Isaac-
1510 cowitz), the National Institutes of Health-National Cancer Institute (NIH-
1511 NCI, grant 1U01CA193632-01A1, principal investigator Barrett), and the
1512 National Science Foundation (NSF-NCS, grant 1835309, principal investiga-
1513 tor van de Meent). The authors would also like to thank Deniz Erdogmus
1514 for elucidating discussions on control theory.

References

- [1] A.D. Craig, How do you feel? Interoception: the sense of the physiological condition of the body, *Nature Reviews Neuroscience* 3 (2002) 655–666. URL: <https://doi.org/10.1038/nrn894><http://10.0.4.14/nrn894>.
- [2] A. D. Craig, *How do you feel?: an interoceptive moment with your neurobiological self*, Princeton University Press, 2015.
- [3] K. S. Quigley, S. Kanoski, W. M. Grill, L. F. Barrett, M. Tsakiris, Functions of Interoception: From Energy Regulation to Experience of the Self, *Trends in Neurosciences* 44 (2021) 29–38. URL: <https://doi.org/10.1016/j.tins.2020.09.008>. doi:10.1016/j.tins.2020.09.008.
- [4] A. D. Craig, How do you feel - now? The anterior insula and human awareness, *Nature Reviews Neuroscience* 10 (2009) 59–70. doi:10.1038/nrn2555. [arXiv:1511.04103](https://arxiv.org/abs/1511.04103).
- [5] W. G. Chen, D. Schloesser, A. M. Arensdorf, J. M. Simmons, C. Cui, R. Valentino, J. W. Gnadt, L. Nielsen, C. S. Hillaire-Clarke, V. Spruance, T. S. Horowitz, Y. F. Vallejo, H. M. Langevin, The Emerging Science of Interoception: Sensing, Integrating, Interpreting, and Regulating Signals within the Self, *Trends in Neurosciences* 44 (2021) 3–16. URL: <https://doi.org/10.1016/j.tins.2020.10.007>. doi:10.1016/j.tins.2020.10.007.
- [6] R. Dantzer, Neuroimmune interactions: From the brain to the immune system and vice versa, *Physiological Reviews* 98 (2018) 477–504. doi:10.1152/physrev.00039.2016.
- [7] I. E. de Araujo, M. Schatzker, D. M. Small, Rethinking Food Reward, *Annual Review of Psychology* 71 (2020) 1–26. doi:10.1146/annurev-psych-122216-011643.
- [8] P. A. Muller, F. Matheis, M. Schneeberger, Z. Kerner, V. Jové, D. Mucida, Microbiota-modulated CART+ enteric neurons autonomously regulate blood glucose, *Science* 370 (2020) 314–321. doi:10.1126/science.abd6176.

- [9] D. McNamee, D. M. Wolpert, Internal Models in Biological Control, *Annual Review of Control, Robotics, and Autonomous Systems* 2 (2019) 339–364. doi:[10.1146/annurev-control-060117-105206](https://doi.org/10.1146/annurev-control-060117-105206).
- [10] G. Buzsáki, D. Tingley, Space and Time: The Hippocampus as a Sequence Generator, *Trends in Cognitive Sciences* 22 (2018) 853–869. URL: <https://doi.org/10.1016/j.tics.2018.07.006>. doi:[10.1016/j.tics.2018.07.006](https://doi.org/10.1016/j.tics.2018.07.006).
- [11] P. Schwartenbeck, T. H. FitzGerald, R. Dolan, Neural signals encoding shifts in beliefs, *NeuroImage* 125 (2016) 578–586.
- [12] N. C. Aggelopoulos, Perceptual inference, *Neuroscience and Biobehavioral Reviews* 55 (2015) 375–392. URL: <http://dx.doi.org/10.1016/j.neubiorev.2015.05.001>. doi:[10.1016/j.neubiorev.2015.05.001](https://doi.org/10.1016/j.neubiorev.2015.05.001).
- [13] H. Von Helmholtz, *Treatise on physiological optics* vol. iii (1867).
- [14] L. W. Barsalou, Simulation, situated conceptualization, and prediction, *Philosophical Transactions of the Royal Society B: Biological Sciences* 364 (2009) 1281–1289. doi:[10.1098/rstb.2008.0319](https://doi.org/10.1098/rstb.2008.0319).
- [15] L. F. Barrett, The theory of constructed emotion: an active inference account of interoception and categorization, *Social cognitive and affective neuroscience* 12 (2017) 1–23. doi:[10.1093/scan/nsw154](https://doi.org/10.1093/scan/nsw154). [arXiv:scan/nsw154](https://arxiv.org/abs/1704.04730).
- [16] R. Grush, The emulation theory of representation: Motor control, imagery, and perception, *Behavioral and brain sciences* 27 (2004) 377.
- [17] M. Bar, Predictions: A Universal Principle in the Operation of the Human Brain, *Philosophical Transactions of the Royal Society B: Biological Sciences* 364 (2009) 1181–1182. doi:[10.1093/acprof:oso/9780195395518.001.0001](https://doi.org/10.1093/acprof:oso/9780195395518.001.0001).
- [18] K. Friston, S. Kiebel, Predictive coding under the free-energy principle, *Philosophical Transactions of the Royal Society B: Biological Sciences* 364 (2009) 1211–1221. URL: <http://rstb.royalsocietypublishing.org/cgi/doi/10.1098/rstb.2008.0300>. doi:[10.1098/rstb.2008.0300](https://doi.org/10.1098/rstb.2008.0300).

- 1579 [19] L. F. Barrett, W. K. Simmons, Interoceptive predictions in the
1580 brain, *Nature Reviews Neuroscience* 16 (2015) 419–429. URL:
1581 <http://dx.doi.org/10.1038/nrn3950>
1582 <http://dx.doi.org/10.1038/nrn3950>.
1583 doi:10.1038/nrn3950. arXiv:15334406.
- 1584 [20] J. B. Hutchinson, L. F. Barrett, The Power of Predictions: An Emerg-
1585 ing Paradigm for Psychological Research, *Current Directions in Psy-*
1586 *chological Science* (2019). doi:10.1177/0963721419831992.
- 1587 [21] P. Sterling, Allostasis: A model of predictive regulation, *Physiology*
1588 *and Behavior* 106 (2012) 5–15. URL: <http://dx.doi.org/10.1016/j.physbeh.2011.06.004>.
1589 doi:10.1016/j.physbeh.2011.06.004.
- 1590 [22] P. Sterling, S. Laughlin, *Principles of neural design*, MIT Press, 2015.
- 1591 [23] J. Schulkin, P. Sterling, Allostasis: A Brain-Centered, Predictive
1592 Mode of Physiological Regulation, *Trends in Neurosciences* 42 (2019)
1593 740–752. URL: <http://dx.doi.org/10.1016/j.tins.2019.07.010>.
1594 doi:10.1016/j.tins.2019.07.010.
- 1595 [24] P. Cisek, Resynthesizing behavior through phylogenetic refine-
1596 ment, *Attention, Perception, and Psychophysics* (2019). doi:10.3758/
1597 s13414-019-01760-1.
- 1598 [25] L. Chanes, L. F. Barrett, Redefining the Role of Limbic Areas in Cor-
1599 tical Processing, *Trends in Cognitive Sciences* 20 (2016) 96–106. URL:
1600 <http://dx.doi.org/10.1016/j.tics.2015.11.005>. doi:10.1016/j.
1601 tics.2015.11.005.
- 1602 [26] I. R. Kleckner, J. Zhang, A. Touroutoglou, L. Chanes, C. Xia, W. K.
1603 Simmons, K. S. Quigley, B. C. Dickerson, L. Feldman Barrett, Evi-
1604 dence for a large-scale brain system supporting allostasis and intero-
1605 ception in humans, *Nature Human Behaviour* 1 (2017). doi:10.1038/
1606 s41562-017-0069.
- 1607 [27] O. J. Hulme, T. Morville, B. Gutkin, Neurocomputational theo-
1608 ries of homeostatic control, *Physics of Life Reviews* 1 (2019) 1–
1609 19. URL: <https://doi.org/10.1016/j.plrev.2019.07.005>. doi:10.
1610 1016/j.plrev.2019.07.005.

- [28] F. H. Petzschner, S. N. Garfinkel, M. P. Paulus, C. Koch, S. S. Khalsa, Computational Models of Interoception and Body Regulation, *Trends in Neurosciences* 44 (2021) 63–76. URL: <https://doi.org/10.1016/j.tins.2020.09.012>. doi:10.1016/j.tins.2020.09.012.
- [29] O. Unal, O. C. Eren, G. Alkan, F. H. Petzschner, Y. Yao, K. E. Stephan, Inference on homeostatic belief precision, *Biological Psychology* 165 (2021) 108190. URL: <https://www.sciencedirect.com/science/article/pii/S0301051121001836>. doi:<https://doi.org/10.1016/j.biopsycho.2021.108190>.
- [30] R. H. S. Carpenter, Homeostasis: a plea for a unified approach, *AJP: Advances in Physiology Education* 28 (2004) 180–187. URL: <http://ajpadvan.physiology.org/cgi/doi/10.1152/advan.00012.2004>. doi:10.1152/advan.00012.2004.
- [31] P. Sterling, Homeostasis vs allostasis: implications for brain function and mental disorders, *JAMA psychiatry* 71 (2014) 1192–1193.
- [32] N. Mrosovsky, *Rheostasis: the physiology of change.*, Oxford University Press, 1990.
- [33] M. Cabanac, Adjustable set point: To honor Harold T. Hammel, *Journal of Applied Physiology* 100 (2006) 1338–1346. doi:10.1152/japplphysiol.01021.2005.
- [34] S. C. Woods, D. S. Ramsay, Homeostasis: Beyond Curt Richter, *Appetite* 49 (2007) 388–398. doi:10.1016/j.appet.2006.09.015.
- [35] M. E. Kotas, R. Medzhitov, Homeostasis, Inflammation, and Disease Susceptibility, *Cell* 160 (2015) 816–827. URL: <http://dx.doi.org/10.1016/j.cell.2015.02.010>. doi:10.1016/j.cell.2015.02.010.
- [36] K. P. Körding, D. M. Wolpert, Bayesian decision theory in sensorimotor control, *Trends in Cognitive Sciences* 10 (2006) 319–326. doi:10.1016/j.tics.2006.05.003.
- [37] R. B. Gillespie, A. H. Ghasemi, J. Freudenberg, Human Motor Control and the Internal Model Principle, *IFAC-PapersOnLine* 49 (2016) 114–119. doi:10.1016/j.ifacol.2016.10.471.

- [38] G. Pezzulo, P. Cisek, Navigating the Affordance Landscape: Feedback Control as a Process Model of Behavior and Cognition, *Trends in Cognitive Sciences* 20 (2016) 414–424. URL: <http://dx.doi.org/10.1016/j.tics.2016.03.013>. doi:10.1016/j.tics.2016.03.013.
- [39] G. Pezzulo, F. Donnarumma, P. Iodice, D. Maisto, I. Stoianov, Model-based approaches to active perception and control, *Entropy* 19 (2017). doi:10.3390/e19060266.
- [40] R. S. Maeda, T. Cluff, P. L. Gribble, J. A. Pruszynski, Feedforward and feedback control share an internal model of the arm’s dynamics, *Journal of Neuroscience* 38 (2018) 10505–10514.
- [41] H. A. Young, C. M. Gaylor, D. de Kerckhove, H. Watkins, D. Benton, Interoceptive accuracy moderates the response to a glucose load: a test of the predictive coding framework, *Proceedings of the Royal Society B* 286 (2019) 20190244.
- [42] G. Pezzulo, F. Rigoli, K. Friston, Active Inference, homeostatic regulation and adaptive behavioural control, *Progress in Neurobiology* 134 (2015) 17–35. URL: <http://dx.doi.org/10.1016/j.pneurobio.2015.09.001>. doi:10.1016/j.pneurobio.2015.09.001.
- [43] A. W. Corcoran, J. Hohwy, Allostasis, interoception , and the free energy principle : Feeling our way forward, in: *Interoceptive basis of the mind*, 2017, pp. 1–16. URL: psyarxiv.com/zbqnx. doi:10.17605/OSF.IO/ZBQNX.
- [44] W. B. Cannon, Organization for physiological homeostasis, *Physiological reviews* 9 (1929) 399–431.
- [45] G. G. Berntson, J. T. Cacioppo, K. S. Quigley, Autonomic determinism: The modes of autonomic control, the doctrine of autonomic space, and the laws of autonomic constraint, *Psychological Review* 98 (1991) 459–487. doi:10.1037/0033-295X.98.4.459.
- [46] L. B. Rowell, D. S. O’Leary, D. L. Kellogg, Integration of Cardiovascular Control Systems in Dynamic Exercise, *Comprehensive Physiology* (1996) 770–838. doi:10.1002/cphy.cp120117.

- 1673 [47] J. T. Potts, X. R. Shi, P. B. Raven, Carotid baroreflex responsiveness
1674 during dynamic exercise in humans, *American Journal of Physiology -*
1675 *Heart and Circulatory Physiology* 265 (1993). doi:[10.1152/ajpheart.](https://doi.org/10.1152/ajpheart.1993.265.6.h1928)
1676 [1993.265.6.h1928](https://doi.org/10.1152/ajpheart.1993.265.6.h1928).
- 1677 [48] J. T. Potts, X. Shi, P. B. Raven, Cardiopulmonary baroreceptors mod-
1678 ulate carotid baroreflex control of heart rate during dynamic exercise
1679 in humans, *American Journal of Physiology - Heart and Circulatory*
1680 *Physiology* 268 (1995). doi:[10.1152/ajpheart.1995.268.4.h1567](https://doi.org/10.1152/ajpheart.1995.268.4.h1567).
- 1681 [49] S. Ogoh, W. L. Wasmund, D. M. Keller, A. O. Yurvati, K. M. Gal-
1682 lagher, J. H. Mitchell, P. B. Raven, Role of central command in carotid
1683 baroreflex resetting in humans during static exercise, *Journal of Phys-*
1684 *iology* 543 (2002) 349–364. doi:[10.1113/jphysiol.2002.019943](https://doi.org/10.1113/jphysiol.2002.019943).
- 1685 [50] L. M. McDowall, R. A. Dampney, Calculation of threshold and satura-
1686 tion points of sigmoidal baroreflex function curves, *American Journal of*
1687 *Physiology - Heart and Circulatory Physiology* 291 (2006) 2003–2007.
1688 doi:[10.1152/ajpheart.00219.2006](https://doi.org/10.1152/ajpheart.00219.2006).
- 1689 [51] S. Ogoh, J. P. Fisher, E. A. Dawson, M. J. White, N. H. Secher, P. B.
1690 Raven, Autonomic nervous system influence on arterial baroreflex con-
1691 trol of heart rate during exercise in humans, *Journal of Physiology* 566
1692 (2005) 599–611. doi:[10.1113/jphysiol.2005.084541](https://doi.org/10.1113/jphysiol.2005.084541).
- 1693 [52] C. M. Heesch, Reflexes that control cardiovascular function, *American*
1694 *Journal of Physiology - Advances in Physiology Education* 22 (1999)
1695 234–244.
- 1696 [53] R. A. Dampney, Central neural control of the cardiovascular system:
1697 Current perspectives, *Advances in Physiology Education* 40 (2016)
1698 283–296. doi:[10.1152/advan.00027.2016](https://doi.org/10.1152/advan.00027.2016).
- 1699 [54] A. Bevan, A. Honour, F. Stott, Direct arterial pressure recording in
1700 unrestricted man., *Clinical Science* 36 (1969) 329–344.
- 1701 [55] J. E. Theriault, L. Young, L. F. Barrett, The sense of should: A
1702 biologically-based framework for modeling social pressure, *Physics of*
1703 *Life Reviews* 36 (2021) 100–136. URL: [https://doi.org/10.1016/j.](https://doi.org/10.1016/j.plrev.2020.01.004)
1704 [plrev.2020.01.004](https://doi.org/10.1016/j.plrev.2020.01.004). doi:[10.1016/j.plrev.2020.01.004](https://doi.org/10.1016/j.plrev.2020.01.004).

- 1705 [56] P. T. Saunders, J. H. Koeslag, J. A. Wessels, Integral
1706 Rein Control in Physiology, Journal of Theoretical Biology
1707 194 (1998) 163–173. URL: [https://www.sciencedirect.com/
1708 science/article/pii/S0022519398907469?via%3Dihub%0Ahttp:
1709 //linkinghub.elsevier.com/retrieve/pii/S0022519398907469%
1710 5Cnpapers2://publication/doi/10.1006/jtbi.1998.0746](https://www.sciencedirect.com/science/article/pii/S0022519398907469?via%3Dihub%0Ahttp://linkinghub.elsevier.com/retrieve/pii/S0022519398907469%5Cnpapers2://publication/doi/10.1006/jtbi.1998.0746).
1711 doi:10.1006/JTBI.1998.0746.
- 1712 [57] P. T. Saunders, J. H. Koeslag, J. a. Wessels, Integral Rein Control in
1713 Physiology II: A General Model, Journal of Theoretical Biology (2000)
1714 1–13.
- 1715 [58] J. W. Sohn, W. K. Ho, Cellular and systemic mechanisms for glu-
1716 cose sensing and homeostasis, Pflugers Archiv European Journal of
1717 Physiology 472 (2020) 1547–1561. doi:10.1007/s00424-020-02466-2.
- 1718 [59] I. Garzilli, S. Itzkovitz, Design principles of the paradoxical feedback
1719 between pancreatic alpha and beta cells, Scientific Reports 8 (2018)
1720 1–12. doi:10.1038/s41598-018-29084-4.
- 1721 [60] J. R. Speakman, D. A. Levitsky, D. B. Allison, M. S. Bray, J. M. De
1722 Castro, D. J. Clegg, J. C. Clapham, A. G. Dulloo, L. Gruer, S. Haw,
1723 J. Hebebrand, M. M. Hetherington, S. Higgs, S. A. Jebb, R. J. Loos,
1724 S. Luckman, A. Luke, V. Mohammed-Ali, S. O’Rahilly, M. Pereira,
1725 L. Perusse, T. N. Robinson, B. Rolls, M. E. Symonds, M. S. Westerterp-
1726 Plantenga, Set points, settling points and some alternative models:
1727 Theoretical options to understand how genes and environments com-
1728 bine to regulate body adiposity, DMM Disease Models and Mechanisms
1729 4 (2011) 733–745. doi:10.1242/dmm.008698.
- 1730 [61] B. M. Filippi, M. A. Abraham, J. T. Yue, T. K. Lam, Insulin and
1731 glucagon signaling in the central nervous system, Reviews in En-
1732 docrine and Metabolic Disorders 14 (2013) 365–375. doi:10.1007/
1733 s11154-013-9258-4.
- 1734 [62] J. Hall, M. Hall, Guyton and Hall Textbook of Medical Physiology E-
1735 Book, Guyton Physiology, Elsevier Health Sciences, 2020. URL: [https:
1736 //books.google.com/books?id=H1rrDwAAQBAJ](https://books.google.com/books?id=H1rrDwAAQBAJ).

- 1737 [63] M. König, S. Bulik, H. G. Holzhütter, Quantifying the contribution
1738 of the liver to glucose homeostasis: A detailed kinetic model of human
1739 hepatic glucose metabolism, *PLoS Computational Biology* 8 (2012).
1740 doi:[10.1371/journal.pcbi.1002577](https://doi.org/10.1371/journal.pcbi.1002577).
- 1741 [64] T. Morville, K. Friston, D. Burdakov, H. R. Siebner, O. J. Hulme, The
1742 homeostatic logic of reward, *bioRxiv* (2018) 242974.
- 1743 [65] C. A. Zimmerman, Y. C. Lin, D. E. Leib, L. Guo, E. L. Huey, G. E.
1744 Daly, Y. Chen, Z. A. Knight, Thirst neurons anticipate the homeo-
1745 static consequences of eating and drinking, *Nature* 537 (2016) 680–
1746 684. URL: <http://dx.doi.org/10.1038/nature18950>. doi:[10.1038/](https://doi.org/10.1038/nature18950)
1747 [nature18950](https://doi.org/10.1038/nature18950).
- 1748 [66] Y. Chen, Z. A. Knight, Making sense of the sensory regulation of hunger
1749 neurons, *BioEssays* 38 (2016) 316–324. doi:[10.1002/bies.201500167](https://doi.org/10.1002/bies.201500167).
- 1750 [67] C. A. Zimmerman, Z. A. Knight, Layers of signals that regulate ap-
1751 petite, *Current opinion in neurobiology* 64 (2020) 79–88.
- 1752 [68] K. C. Berridge, M. L. Kringelbach, Pleasure Systems in the Brain,
1753 *Neuron* 86 (2015) 646–664. URL: [http://dx.doi.org/10.1016/j.](http://dx.doi.org/10.1016/j.neuron.2015.02.018)
1754 [neuron.2015.02.018](https://doi.org/10.1016/j.neuron.2015.02.018). doi:[10.1016/j.neuron.2015.02.018](https://doi.org/10.1016/j.neuron.2015.02.018).
- 1755 [69] L. F. Barrett, A. B. Satpute, Large-scale brain networks in affective
1756 and social neuroscience: towards an integrative functional architecture
1757 of the brain, *Current opinion in neurobiology* 23 (2013) 361–372.
- 1758 [70] L. M. Hackel, G. M. Larson, J. D. Bowen, G. A. Ehrlich, T. C. Mann,
1759 B. Middlewood, I. D. Roberts, J. Eyink, J. C. Fetterolf, F. Gonzalez,
1760 et al., On the neural implausibility of the modular mind: Evidence
1761 for distributed construction dissolves boundaries between perception,
1762 cognition, and emotion, *Behavioral and Brain Sciences* 39 (2016).
- 1763 [71] M. Keramati, B. Gutkin, Homeostatic reinforcement learning for in-
1764 tegrating reward collection and physiological stability, *eLife* 3 (2014).
1765 doi:[10.7554/eLife.04811.001](https://doi.org/10.7554/eLife.04811.001).
- 1766 [72] L. F. Barrett, The future of psychology: Connecting mind to brain,
1767 *Perspectives on psychological science* 4 (2009) 326–339.

- 1768 [73] K. A. Lindquist, L. F. Barrett, A functional architecture of the hu-
1769 man brain: emerging insights from the science of emotion, *Trends in*
1770 *cognitive sciences* 16 (2012) 533–540.
- 1771 [74] G. B. Carvalho, A. Damasio, Interoception and the origin of feel-
1772 ings: A new synthesis, *BioEssays* 43 (2021) 1–11. doi:[10.1002/bies.](https://doi.org/10.1002/bies.202000261)
1773 [202000261](https://doi.org/10.1002/bies.202000261).
- 1774 [75] Y. L. Boureau, P. Dayan, Opponency revisited: Competition and coop-
1775 eration between dopamine and serotonin, *Neuropsychopharmacology*
1776 36 (2011) 74–97. URL: <http://dx.doi.org/10.1038/npp.2010.151>.
1777 doi:[10.1038/npp.2010.151](https://doi.org/10.1038/npp.2010.151).
- 1778 [76] W. Dabney, M. Rowland, M. G. Bellemare, R. Munos, Distributional
1779 reinforcement learning with quantile regression, in: *Thirty-Second*
1780 *AAAI Conference on Artificial Intelligence*, 2018.
- 1781 [77] W. Dabney, Z. Kurth-Nelson, N. Uchida, C. K. Starkweather, D. Has-
1782 sabis, R. Munos, M. Botvinick, A distributional code for value in
1783 dopamine-based reinforcement learning, *Nature* 577 (2020) 671–675.
- 1784 [78] A. S. Lowet, Q. Zheng, S. Matias, J. Drugowitsch, N. Uchida, Dis-
1785 tributional reinforcement learning in the brain, *Trends in Neuro-*
1786 *sciences* 43 (2020) 980–997. URL: [https://www.sciencedirect.com/](https://www.sciencedirect.com/science/article/pii/S0166223620301983)
1787 [science/article/pii/S0166223620301983](https://www.sciencedirect.com/science/article/pii/S0166223620301983). doi:[https://doi.org/](https://doi.org/10.1016/j.tins.2020.09.004)
1788 [10.1016/j.tins.2020.09.004](https://doi.org/10.1016/j.tins.2020.09.004).
- 1789 [79] C. Cosentino, D. Bates, *Feedback Control in Systems Biology*,
1790 CRC Press, 2011. URL: [https://books.google.com/books?id=wk_](https://books.google.com/books?id=wk_RBQAAQBAJ)
1791 [RBQAAQBAJ](https://books.google.com/books?id=wk_RBQAAQBAJ).
- 1792 [80] B. S. Zanutto, M. E. Valentinuzzi, E. T. Segura, Neural set point
1793 for the control of arterial pressure: Role of the nucleus tractus soli-
1794 tarius, *BioMedical Engineering Online* 9 (2010) 1–13. doi:[10.1186/](https://doi.org/10.1186/1475-925X-9-4)
1795 [1475-925X-9-4](https://doi.org/10.1186/1475-925X-9-4).
- 1796 [81] J. W. Osborn, J. D. Foss, Renal nerves and long-term control of arterial
1797 pressure, *Comprehensive Physiology* 7 (2017) 263–320. doi:[10.1002/](https://doi.org/10.1002/cphy.c150047)
1798 [cphy.c150047](https://doi.org/10.1002/cphy.c150047).

- [82] R. C. Conant, W. R. Ashby, Every good regulator of a system must be a model of that system, *International Journal of Systems Science* 1 (1970) 89–97. doi:[10.1080/00207727008920220](https://doi.org/10.1080/00207727008920220).
- [83] B. A. Francis, W. M. Wonham, The internal model principle of control theory, *Automatica* 12 (1976) 457–465. doi:[10.1016/0005-1098\(76\)90006-6](https://doi.org/10.1016/0005-1098(76)90006-6).
- [84] D. M. Wolpert, M. Kawato, Multiple paired forward and inverse models for motor control, *Neural Networks* 11 (1998) 1317–1329. doi:[10.1016/S0893-6080\(98\)00066-5](https://doi.org/10.1016/S0893-6080(98)00066-5).
- [85] M. Kawato, Internal models for motor control and trajectory planning, *Current Opinion in Neurobiology* 9 (1999) 718–727. doi:[10.1016/S0959-4388\(99\)00028-8](https://doi.org/10.1016/S0959-4388(99)00028-8).
- [86] A. G. Feldman, Referent control of action and perception, 2015. doi:[10.1007/978-1-4939-2736-4](https://doi.org/10.1007/978-1-4939-2736-4).
- [87] M. L. Latash, Laws of nature that define biological action and perception, *Physics of Life Reviews* 36 (2021) 47–67. URL: <https://doi.org/10.1016/j.plrev.2020.07.007>. doi:[10.1016/j.plrev.2020.07.007](https://doi.org/10.1016/j.plrev.2020.07.007).
- [88] M. L. Latash, Motor synergies and the equilibrium-point hypothesis, *Motor Control* 14 (2010) 294–322. doi:[10.1123/mcj.14.3.294](https://doi.org/10.1123/mcj.14.3.294).
- [89] A. G. Feldman, Active sensing without efference copy: Referent control of perception, *Journal of Neurophysiology* 116 (2016) 960–976. doi:[10.1152/jn.00016.2016](https://doi.org/10.1152/jn.00016.2016).
- [90] J. Merel, M. Botvinick, G. Wayne, Hierarchical motor control in mammals and machines, *Nature Communications* 10 (2019) 1–12. URL: <http://dx.doi.org/10.1038/s41467-019-13239-6>. doi:[10.1038/s41467-019-13239-6](https://doi.org/10.1038/s41467-019-13239-6).
- [91] H. Straka, J. Simmers, B. P. Chagnaud, A New Perspective on Predictive Motor Signaling, *Current Biology* 28 (2018) R232–R243. URL: <https://doi.org/10.1016/j.cub.2018.01.033>. doi:[10.1016/j.cub.2018.01.033](https://doi.org/10.1016/j.cub.2018.01.033).

- 1830 [92] L. F. Barrett, B. L. Finlay, Concepts, goals and the control of survival-
1831 related behaviors, *Current opinion in behavioral sciences* 24 (2018)
1832 172–179.
- 1833 [93] A. Leshinskaya, M. F. Wurm, A. Caramazza, Concepts of Actions and
1834 their Objects, in: *The Cognitive Neurosciences*, 2020, pp. 757–765.
- 1835 [94] D. M. Wolpert, K. G. Pearson, C. P. Ghez, The Organization and
1836 Planning of Movement, in: *Principles of Neuroscience*, 5 ed., 2013, pp.
1837 743–767.
- 1838 [95] G. Hickok, The myth of mirror neurons: The real neuroscience of com-
1839 munication and cognition, WW Norton & Company, 2014.
- 1840 [96] A. Clark, Surfing uncertainty: Prediction, action, and the embodied
1841 mind, Oxford University Press, 2015.
- 1842 [97] J. W. Barter, H. H. Yin, Achieving natural behavior in a robot
1843 using neurally inspired hierarchical perceptual control, *iScience*
1844 24 (2021) 102948. URL: [https://www.sciencedirect.com/science/](https://www.sciencedirect.com/science/article/pii/S2589004221009160)
1845 [article/pii/S2589004221009160](https://www.sciencedirect.com/science/article/pii/S2589004221009160). doi:[https://doi.org/10.1016/](https://doi.org/10.1016/j.isci.2021.102948)
1846 [j.isci.2021.102948](https://doi.org/10.1016/j.isci.2021.102948).
- 1847 [98] Y. Nakahira, Q. Liu, N. Bernat, T. Sejnowski, J. Doyle, Theoretical
1848 foundations for layered architectures and speed-accuracy tradeoffs in
1849 sensorimotor control, *Proceedings of the American Control Conference*
1850 2019-July (2019) 809–814. doi:[10.23919/acc.2019.8814897](https://doi.org/10.23919/acc.2019.8814897).
- 1851 [99] G. Buzsáki, The brain from inside out, Oxford University Press, 2019.
- 1852 [100] X. Gu, T. FitzGerald, Interoceptive inference: homeostasis and
1853 decision-making, *Trends Cogn Sci* 18 (2014) 269–70.
- 1854 [101] G. Pezzulo, F. Rigoli, K. J. Friston, Hierarchical Active Inference: A
1855 Theory of Motivated Control, *Trends in Cognitive Sciences* 22 (2018)
1856 294–306. URL: <http://dx.doi.org/10.1016/j.tics.2018.01.009>.
1857 doi:[10.1016/j.tics.2018.01.009](https://doi.org/10.1016/j.tics.2018.01.009).
- 1858 [102] A. W. Corcoran, G. Pezzulo, J. Hohwy, From allostatic agents to
1859 counterfactual cognisers: active inference, biological regulation, and
1860 the origins of cognition, *Biology & Philosophy* 35 (2020) 1–45.

- 1861 [103] Richard S. Sutton and Andrew G. Barto, Reinforcement Learning, Sec-
1862 ond Edition, 2018.
- 1863 [104] Y. Niv, Reinforcement learning in the brain, *Journal of Mathematical*
1864 *Psychology* 53 (2009) 139–154.
- 1865 [105] M. Colombo, Deep and beautiful. the reward prediction error hypoth-
1866 esis of dopamine, *Studies in history and philosophy of science part C:*
1867 *Studies in history and philosophy of biological and biomedical sciences*
1868 45 (2014) 57–67.
- 1869 [106] A. S. Lowet, Q. Zheng, S. Matias, J. Drugowitsch, N. Uchida, Distri-
1870 butional reinforcement learning in the brain, *Trends in Neurosciences*
1871 (2020).
- 1872 [107] M. Mangalam, D. G. Kelty-Stephen, Point estimates, Simpson’s para-
1873 dox, and nonergodicity in biological sciences, *Neuroscience and Biobe-*
1874 *havioral Reviews* 125 (2021) 98–107. URL: [https://doi.org/10.](https://doi.org/10.1016/j.neubiorev.2021.02.017)
1875 [1016/j.neubiorev.2021.02.017](https://doi.org/10.1016/j.neubiorev.2021.02.017). doi:[10.1016/j.neubiorev.2021.](https://doi.org/10.1016/j.neubiorev.2021.02.017)
1876 [02.017](https://doi.org/10.1016/j.neubiorev.2021.02.017).
- 1877 [108] K. Friston, R. Adams, R. Montague, What is value-accumulated reward
1878 or evidence?, *Frontiers in Neurorobotics* 6 (2012) 1–25. doi:[10.3389/](https://doi.org/10.3389/fnbot.2012.00011)
1879 [fnbot.2012.00011](https://doi.org/10.3389/fnbot.2012.00011).
- 1880 [109] K. M. Lee, F. Ferreira-Santos, A. B. Satpute, Predictive processing
1881 models and affective neuroscience, *Neuroscience and Biobehavioral*
1882 *Reviews* 131 (2021) 211–228. URL: [https://doi.org/10.1016/j.](https://doi.org/10.1016/j.neubiorev.2021.09.009)
1883 [neubiorev.2021.09.009](https://doi.org/10.1016/j.neubiorev.2021.09.009). doi:[10.1016/j.neubiorev.2021.09.009](https://doi.org/10.1016/j.neubiorev.2021.09.009).
- 1884 [110] S. E. Dicarlo, V. S. Bishop, Onset of exercise shifts operating point
1885 of arterial baroreflex to higher pressures, *American Journal of Phys-*
1886 *iology - Heart and Circulatory Physiology* 262 (1992). doi:[10.1152/](https://doi.org/10.1152/ajpheart.1992.262.1.h303)
1887 [ajpheart.1992.262.1.h303](https://doi.org/10.1152/ajpheart.1992.262.1.h303).
- 1888 [111] K. J. Friston, J. Daunizeau, J. Kilner, S. J. Kiebel, Action and be-
1889 havior: A free-energy formulation, *Biological Cybernetics* 102 (2010)
1890 227–260. doi:[10.1007/s00422-010-0364-z](https://doi.org/10.1007/s00422-010-0364-z).

- 1891 [112] K. Friston, S. Samothrakis, R. Montague, Active inference and agency:
1892 Optimal control without cost functions, *Biological Cybernetics* 106
1893 (2012) 523–541. doi:[10.1007/s00422-012-0512-8](https://doi.org/10.1007/s00422-012-0512-8).
- 1894 [113] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, E. A. Theodorou,
1895 Information theoretic model predictive control: Theory and applica-
1896 tions to autonomous driving, *IEEE Transactions on Robotics* 34 (2018)
1897 1603–1622.
- 1898 [114] S. Nasiriany, V. H. Pong, A. Nair, A. Khazatsky, G. Berseth,
1899 S. Levine, DisCo RL: Distribution-Conditioned Reinforcement Learn-
1900 ing for General-Purpose Policies, in: *IEEE International Conference on*
1901 *Robotics and Automation*, 2021. URL: [http://arxiv.org/abs/2104.](http://arxiv.org/abs/2104.11707)
1902 [11707](http://arxiv.org/abs/2104.11707). [arXiv:2104.11707](https://arxiv.org/abs/2104.11707).
- 1903 [115] R. Smith, J. F. Thayer, S. S. Khalsa, R. D. Lane, The hierarchi-
1904 cal basis of neurovisceral integration, *Neuroscience and Biobehavioral*
1905 *Reviews* 75 (2017) 274–296. URL: [http://doi.wiley.com/10.1029/](http://doi.wiley.com/10.1029/2003GL019380)
1906 [2003GL019380](http://doi.wiley.com/10.1029/2003GL019380). doi:[10.1016/j.neubiorev.2017.02.003](https://doi.org/10.1016/j.neubiorev.2017.02.003).
- 1907 [116] B. Millidge, Deep active inference as variational policy gradients,
1908 *Journal of Mathematical Psychology* 96 (2020) 102348. URL: <https://doi.org/10.1016/j.jmp.2020.102348>. doi:[10.1016/j.jmp.2020.](https://doi.org/10.1016/j.jmp.2020.102348)
1909 [102348](https://doi.org/10.1016/j.jmp.2020.102348). [arXiv:1907.03876](https://arxiv.org/abs/1907.03876).
- 1911 [117] M. Kirchhoff, T. Parr, E. Palacios, K. Friston, J. Kiverstein, The
1912 markov blankets of life: Autonomy, active inference and the free energy
1913 principle, *Journal of the Royal Society Interface* 15 (2018). doi:[10.](https://doi.org/10.1098/rsif.2017.0792)
1914 [1098/rsif.2017.0792](https://doi.org/10.1098/rsif.2017.0792).
- 1915 [118] M. V. Srinivasan, S. B. Laughlin, A. Dubs, Predictive coding: A fresh
1916 view of inhibition in the retina, *Proceedings of the Royal Society of*
1917 *London - Biological Sciences* 216 (1982) 427–459. doi:[10.1098/rspb.](https://doi.org/10.1098/rspb.1982.0085)
1918 [1982.0085](https://doi.org/10.1098/rspb.1982.0085).
- 1919 [119] E. Todorov, Optimal control theory, *Bayesian brain: probabilistic*
1920 *approaches to neural coding* (2006) 268–298.
- 1921 [120] D. W. Stephens, J. R. Krebs, *Foraging theory*, Princeton University
1922 Press, 2019.

- 1923 [121] R. Shadmehr, A. A. Ahmed, Vigor: Neuroeconomics of Movement Con-
1924 trol, MIT Press, 2020.
- 1925 [122] S. Kobayashi, W. Schultz, Influence of reward delays on responses
1926 of dopamine neurons, *Journal of Neuroscience* 28 (2008) 7837–7846.
1927 doi:[10.1523/JNEUROSCI.1600-08.2008](https://doi.org/10.1523/JNEUROSCI.1600-08.2008).
- 1928 [123] N. D. Daw, D. S. Touretzky, Behavioral considerations suggests an
1929 average reward TD model of the dopamine system, *Neurocomputing*
1930 32-33 (2000) 679–684. doi:[10.1016/S0925-2312\(00\)00232-0](https://doi.org/10.1016/S0925-2312(00)00232-0).
- 1931 [124] R. Shadmehr, H. J. Huang, A. A. Ahmed, A representation of effort in
1932 decision-making and motor control, *Current biology* 26 (2016) 1929–
1933 1934.
- 1934 [125] S. J. Kiebel, J. Daunizeau, K. J. Friston, A Hierarchy of Time-
1935 Scales and the Brain, *PLoS Computational Biology* 4 (2008)
1936 e1000209. URL: [http://dx.doi.org/10.1371/journal.pcbi.](http://dx.doi.org/10.1371/journal.pcbi.1000209)
1937 [1000209](http://dx.doi.org/10.1371/journal.pcbi.1000209). doi:[10.1371/journal.pcbi.1000209](https://doi.org/10.1371/journal.pcbi.1000209).
- 1938 [126] D. Koller, N. Friedman, Probabilistic graphical models: principles and
1939 techniques, MIT press, 2009.
- 1940 [127] G. M. Edelman, J. A. Gally, Degeneracy and complexity in biological
1941 systems, *Proceedings of the National Academy of Sciences* 98 (2001)
1942 13763–13768. URL: [http://www.pnas.org/cgi/doi/10.1073/pnas.](http://www.pnas.org/cgi/doi/10.1073/pnas.231499798)
1943 [231499798](http://www.pnas.org/cgi/doi/10.1073/pnas.231499798). doi:[10.1073/pnas.231499798](https://doi.org/10.1073/pnas.231499798).
- 1944 [128] K. J. Friston, C. J. Price, Degeneracy and redundancy in
1945 cognitive anatomy, *Trends in Cognitive Sciences* 7 (2003)
1946 151–152. URL: [https://www.sciencedirect.com/science/](https://www.sciencedirect.com/science/article/pii/S1364661303000548)
1947 [article/pii/S1364661303000548](https://www.sciencedirect.com/science/article/pii/S1364661303000548). doi:[https://doi.org/10.1016/](https://doi.org/10.1016/S1364-6613(03)00054-8)
1948 [S1364-6613\(03\)00054-8](https://doi.org/10.1016/S1364-6613(03)00054-8).
- 1949 [129] B. L. Finlay, R. Uchiyama, Developmental mechanisms channel-
1950 ing cortical evolution, *Trends in Neurosciences* 38 (2015) 69–
1951 76. URL: [https://www.sciencedirect.com/science/article/pii/](https://www.sciencedirect.com/science/article/pii/S0166223614002124)
1952 [S0166223614002124](https://www.sciencedirect.com/science/article/pii/S0166223614002124). doi:[https://doi.org/10.1016/j.tins.2014.](https://doi.org/10.1016/j.tins.2014.11.004)
1953 [11.004](https://doi.org/10.1016/j.tins.2014.11.004).

- 1954 [130] M. Watter, J. T. Springenberg, J. Boedecker, M. Riedmiller, Embed
1955 to Control: a locally linear latent dynamics model for control from raw
1956 images, in: Advances in Neural Information Processing Systems, 2015.
- 1957 [131] K. Chua, R. Calandra, R. McAllister, S. Levine, Deep Reinforcement
1958 Learning in a Handful of Trials using Probabilistic Dynamics Models,
1959 in: Advances in Neural Information Processing Systems, volume 2018-
1960 Decem, 2018, pp. 4754–4765. [arXiv:1805.12114](https://arxiv.org/abs/1805.12114).
- 1961 [132] P. Becker-Ehmck, M. Karl, J. Peters, P. van der Smagt, Learning
1962 to Fly via Deep Model-Based Reinforcement Learning, in: Fourth
1963 Machine Learning in Planning and Control of Robot Motion Workshop
1964 at International Conference on Robotics and Automation, 2020, pp. 1–
1965 14. URL: <http://arxiv.org/abs/2003.08876>. [arXiv:2003.08876](https://arxiv.org/abs/2003.08876).
- 1966 [133] A. B. Satpute, K. A. Lindquist, The Default Mode Network’s Role
1967 in Discrete Emotion, Trends in Cognitive Sciences 23 (2019) 851–
1968 864. URL: <https://doi.org/10.1016/j.tics.2019.07.003>. doi:10.
1969 [1016/j.tics.2019.07.003](https://doi.org/10.1016/j.tics.2019.07.003).
- 1970 [134] R. Smith, R. Kuplicki, J. Feinstein, K. L. Forthman, J. L. Stewart,
1971 M. P. Paulus, T. . Investigators, S. S. Khalsa, A bayesian computa-
1972 tional model reveals a failure to adapt interoceptive precision estimates
1973 across depression, anxiety, eating, and substance use disorders, PLoS
1974 computational biology 16 (2020) e1008484.
- 1975 [135] R. Smith, R. Kuplicki, A. Teed, V. Upshaw, S. S. Khalsa, Confir-
1976 matory evidence that healthy individuals can adaptively adjust prior
1977 expectations and interoceptive precision estimates, in: International
1978 Workshop on Active Inference, Springer, 2020, pp. 156–164.
- 1979 [136] O. K. Harrison, L. Nanz, S. Marino, R. Lüchinger, F. Hennel, A. J.
1980 Hess, S. K. Fraessle, S. Iglesias, F. Vinckier, F. H. Petzschnier, et al.,
1981 Interoception of breathing and its relationship with anxiety, bioRxiv
1982 (2021).
- 1983 [137] M. J. Ramstead, M. D. Kirchhoff, K. J. Friston, A tale of two densities:
1984 active inference is enactive inference, Adaptive Behavior 28 (2020) 225–
1985 239. doi:[10.1177/1059712319862774](https://doi.org/10.1177/1059712319862774).

- 1986 [138] D. Braun, A. Nagengast, D. Wolpert, Risk-sensitivity in sen-
1987 sorimotor control, *Frontiers in Human Neuroscience* 5 (2011)
1988 1. URL: [https://www.frontiersin.org/article/10.3389/fnhum.](https://www.frontiersin.org/article/10.3389/fnhum.2011.00001)
1989 [2011.00001](https://www.frontiersin.org/article/10.3389/fnhum.2011.00001). doi:10.3389/fnhum.2011.00001.
- 1990 [139] Y. Niv, J. A. Edlund, P. Dayan, J. P. O’Doherty, Neu-
1991 ral prediction errors reveal a risk-sensitive reinforcement-
1992 learning process in the human brain, *Journal of Neuro-*
1993 *science* 32 (2012) 551–562. URL: [https://www.jneurosci.](https://www.jneurosci.org/content/32/2/551)
1994 [org/content/32/2/551](https://www.jneurosci.org/content/32/2/551). doi:10.1523/JNEUROSCI.5498-10.2012.
1995 [arXiv:https://www.jneurosci.org/content/32/2/551.full.pdf](https://www.jneurosci.org/content/32/2/551.full.pdf).
- 1996 [140] A. A. Faisal, L. P. Selen, D. M. Wolpert, Noise in the nervous system,
1997 *Nature reviews neuroscience* 9 (2008) 292–303.
- 1998 [141] S. G. Manohar, T. T.-J. Chong, M. A. Apps, A. Batla, M. Stamelou,
1999 P. R. Jarman, K. P. Bhatia, M. Husain, Reward pays the cost of noise
2000 reduction in motor and cognitive control, *Current Biology* 25 (2015)
2001 1707–1716.
- 2002 [142] J. P. Gallivan, C. S. Chapman, D. M. Wolpert, J. R. Flanagan,
2003 Decision-making in sensorimotor control, *Nature Reviews Neu-*
2004 *roscience* 19 (2018) 519–534. URL: [http://dx.doi.org/10.1038/](http://dx.doi.org/10.1038/s41583-018-0045-9)
2005 [s41583-018-0045-9](http://dx.doi.org/10.1038/s41583-018-0045-9). doi:10.1038/s41583-018-0045-9.
- 2006 [143] M. K. McBeath, D. M. Shaffer, M. K. Kaiser, How baseball out-
2007 fielders determine where to run to catch fly balls, *Science* 268 (1995)
2008 569–573. URL: [https://science.sciencemag.org/content/268/](https://science.sciencemag.org/content/268/5210/569)
2009 [5210/569](https://science.sciencemag.org/content/268/5210/569)[https://science.sciencemag.org/content/268/5210/](https://science.sciencemag.org/content/268/5210/569.abstract)
2010 [569.abstract](https://science.sciencemag.org/content/268/5210/569.abstract). doi:10.1126/science.7725104.
- 2011 [144] J. P. Scholz, G. Schöner, The uncontrolled manifold concept: Identify-
2012 ing control variables for a functional task, *Experimental Brain Research*
2013 126 (1999) 289–306. doi:10.1007/s002210050738.
- 2014 [145] A. Clark, Whatever next? Predictive brains, situated agents, and
2015 the future of cognitive science, *Behavioral and Brain Sciences*
2016 36 (2013) 181–204. URL: [http://www.journals.cambridge.org/](http://www.journals.cambridge.org/abstract_S0140525X12000477)
2017 [abstract_S0140525X12000477](http://www.journals.cambridge.org/abstract_S0140525X12000477). doi:10.1017/S0140525X12000477.
2018 [arXiv:0140-525X](http://www.journals.cambridge.org/abstract_S0140525X12000477).

- [146] R. Bogacz, A tutorial on the free-energy framework for modelling perception and learning, *Journal of Mathematical Psychology* 76 (2017) 198–211. URL: <http://dx.doi.org/10.1016/j.jmp.2015.11.003>. doi:10.1016/j.jmp.2015.11.003.
- [147] B. Millidge, A. Tschantz, C. L. Buckley, Whence the Expected Free Energy?, *Neural Computation* 33 (2021) 447–482. [arXiv:arXiv:2004.08128v3](https://arxiv.org/abs/2004.08128v3).
- [148] S. Chatterjee, P. Diaconis, The sample size required in importance sampling, *Annals of Applied Probability* 28 (2018) 1099–1135. doi:10.1214/17-AAP1326. [arXiv:1511.01437](https://arxiv.org/abs/1511.01437).
- [149] S. Still, D. A. Sivak, A. J. Bell, G. E. Crooks, Thermodynamics of prediction, *Physical Review Letters* 109 (2012) 1–5. doi:10.1103/PhysRevLett.109.120604. [arXiv:1203.3271](https://arxiv.org/abs/1203.3271).
- [150] H. J. Kappen, Path integrals and symmetry breaking for optimal control theory, *Journal of Statistical Mechanics: Theory and Experiment* 11011 (2005) 205–229. doi:10.1088/1742-5468/2005/11/P11011. [arXiv:0505066](https://arxiv.org/abs/0505066).
- [151] A. M. Bastos, W. M. Usrey, R. A. Adams, G. R. Mangun, P. Fries, K. J. Friston, Canonical microcircuits for predictive coding., *Neuron* 76 (2012) 695–711. doi:10.1016/j.neuron.2012.10.038.
- [152] B. Belousov, G. Neumann, C. A. Rothkopf, J. Peters, Catching heuristics are optimal control policies, in: *Advances in Neural Information Processing Systems*, 2016.
- [153] B. A. Mitchell, N. Lauharatanahirun, J. O. Garcia, N. Wymbs, S. Grafton, J. M. Vettel, L. R. Petzold, A minimum free energy model of motor learning, *Neural Computation* 31 (2019) 1945–1963. doi:10.1162/neco_a_01219.
- [154] P. Piray, N. Daw, Linear reinforcement learning: Flexible reuse of computation in planning, grid fields, and cognitive control (2019). doi:10.1101/856849.
- [155] K. Friston, F. Rigoli, D. Ognibene, C. Mathys, T. Fitzgerald, G. Pezzulo, Active inference and epistemic value, *Cognitive Neuroscience* 6

- (2015) 187–224. URL: <http://dx.doi.org/10.1080/17588928.2015.1020053>. doi:[10.1080/17588928.2015.1020053](https://doi.org/10.1080/17588928.2015.1020053).
- [156] M. Allen, A. Levy, T. Parr, K. J. Friston, In the body’s eye: the computational anatomy of interoceptive inference, *BioRxiv* (2019) 603928.
- [157] K. Kobayashi, M. Hsu, Common neural code for reward and information value, *Proceedings of the National Academy of Sciences* 116 (2019) 13061–13066.
- [158] A. Tschantz, L. Barca, D. Maisto, C. L. Buckley, A. Seth, G. Pezzulo, Simulating homeostatic, allostatic and goal-directed forms of interoceptive control using active inference, *bioRxiv* (2021).
- [159] K. E. Stephan, Z. M. Manjaly, C. D. Mathys, L. A. Weber, S. Paliwal, T. Gard, M. Tittgemeyer, S. M. Fleming, H. Haker, A. K. Seth, et al., Allostatic self-efficacy: a metacognitive theory of dyshomeostasis-induced fatigue and depression, *Frontiers in human neuroscience* 10 (2016) 550.
- [160] K. Friston, T. Fitzgerald, F. Rigoli, P. Schwartenbeck, G. Pezzulo, Active Inference: a Process Theory, *Neural Computation* 29 (2017) 1–49. URL: <http://arxiv.org/abs/1206.1275>. doi:[10.1162/NECO.arXiv:1206.1275](https://doi.org/10.1162/NECO.arXiv:1206.1275).
- [161] M. Boerlin, C. K. Machens, S. Deneve, Predictive Coding of Dynamical Variables in Balanced Spiking Networks, *PLoS Computational Biology* 9 (2013) 1–16. URL: <http://dx.doi.org/10.1371/journal.pcbi.1003258>. doi:[10.1371/journal.pcbi.1003258](https://doi.org/10.1371/journal.pcbi.1003258).
- [162] M. W. Spratling, A review of predictive coding algorithms, *Brain and Cognition* 112 (2017) 92–97. doi:[10.1016/j.bandc.2015.11.003](https://doi.org/10.1016/j.bandc.2015.11.003).
- [163] J. Kadmon, J. Timcheck, S. Ganguli, Predictive coding in balanced neural networks with noise, chaos and delays, in: *Advances in Neural Information Processing Systems*, NeurIPS, 2020, pp. 1–12. [arXiv:2006.14178](https://arxiv.org/abs/2006.14178).
- [164] J. P. Gallivan, L. Logan, D. M. Wolpert, J. R. Flanagan, Parallel specification of competing sensorimotor control policies for alternative

- 2082 action options, *Nature Neuroscience* 19 (2016) 320–326. doi:[10.1038/](https://doi.org/10.1038/nn.4214)
2083 [nn.4214](https://doi.org/10.1038/nn.4214).
- 2084 [165] L. F. Barrett, K. S. Quigley, P. Hamilton, An active inference theory of
2085 allostasis and interoception in depression, *Philosophical Transactions*
2086 *of the Royal Society B: Biological Sciences* 371 (2016). doi:[10.1098/](https://doi.org/10.1098/rstb.2016.0011)
2087 [rstb.2016.0011](https://doi.org/10.1098/rstb.2016.0011).
- 2088 [166] A. N. Sanborn, Types of approximation for probabilistic cogni-
2089 tion: Sampling and variational, *Brain and Cognition* 112 (2017)
2090 98–101. URL: [https://www.sciencedirect.com/science/article/](https://www.sciencedirect.com/science/article/pii/S0278262615300038)
2091 [pii/S0278262615300038](https://www.sciencedirect.com/science/article/pii/S0278262615300038). doi:[https://doi.org/10.1016/j.bandc.](https://doi.org/10.1016/j.bandc.2015.06.008)
2092 [2015.06.008](https://doi.org/10.1016/j.bandc.2015.06.008), perspectives on Human Probabilistic Inferences and the
2093 ‘Bayesian Brain’.
- 2094 [167] P. Cisek, Cortical mechanisms of action selection: The affordance com-
2095 petition hypothesis, *Philosophical Transactions of the Royal Society*
2096 *B: Biological Sciences* 362 (2007) 1585–1599. doi:[10.1098/rstb.2007.](https://doi.org/10.1098/rstb.2007.2054)
2097 [2054](https://doi.org/10.1098/rstb.2007.2054).
- 2098 [168] P. Cisek, J. F. Kalaska, Neural mechanisms for interacting with a
2099 world full of action choices, *Annual Review of Neuroscience* 33 (2010)
2100 269–298. doi:[10.1146/annurev.neuro.051508.135409](https://doi.org/10.1146/annurev.neuro.051508.135409).
- 2101 [169] S. Thijssen, H. J. Kappen, Path integral control and state-dependent
2102 feedback, *Physical Review E - Statistical, Nonlinear, and Soft*
2103 *Matter Physics* 91 (2015) 1–7. doi:[10.1103/PhysRevE.91.032104](https://doi.org/10.1103/PhysRevE.91.032104).
2104 [arXiv:1406.4026](https://arxiv.org/abs/1406.4026).
- 2105 [170] H. J. Kappen, H. C. Ruiz, Adaptive Importance Sampling for Control
2106 and Inference, *Journal of Statistical Physics* 162 (2016) 1244–1266.
2107 doi:[10.1007/s10955-016-1446-7](https://doi.org/10.1007/s10955-016-1446-7). [arXiv:1505.01874](https://arxiv.org/abs/1505.01874).
- 2108 [171] S. A. Menchón, H. J. Kappen, Learning effective state-feedback
2109 controllers through efficient multilevel importance samplers, *Inter-*
2110 *national Journal of Control* 92 (2019) 2776–2783. URL: [https://](https://doi.org/10.1080/00207179.2018.1459857)
2111 doi.org/10.1080/00207179.2018.1459857. doi:[10.1080/00207179.](https://doi.org/10.1080/00207179.2018.1459857)
2112 [2018.1459857](https://doi.org/10.1080/00207179.2018.1459857).
- 2113 [172] K. L. Stachenfeld, M. M. Botvinick, S. J. Gershman, Hip-
2114 pocampus as predictive map, *Nature neuroscience* 28 (2017)

2115 391–397. URL: <http://dx.doi.org/10.1038/nn.4650>. doi:10.1017/
2116 CB09781107415324.004. arXiv:arXiv:1011.1669v3.

2117 [173] N. D. Daw, S. J. Gershman, I. Momennejad, E. M. Russek, M. M.
2118 Botvinick, Predictive representations can link model-based reinforcement
2119 learning to model-free mechanisms, volume 13, 2017. doi:10.
2120 1371/journal.pcbi.1005768. arXiv:arXiv:1612.00429v2.

2121 [174] D. M. Wolpert, R. C. Miall, M. Kawato, Internal models in the cere-
2122 bellum, Trends in cognitive sciences 2 (1998) 338–347.

2123 [175] C. Hull, Prediction signals in the cerebellum: beyond supervised motor
2124 learning, Elife 9 (2020) e54073.

2125 [176] R. L. Buckner, The serendipitous discovery of the brain’s default
2126 network, NeuroImage 62 (2012) 1137–1145. URL: [https://www.](https://www.sciencedirect.com/science/article/pii/S1053811911011992)
2127 [sciencedirect.com/science/article/pii/S1053811911011992](https://www.sciencedirect.com/science/article/pii/S1053811911011992).
2128 doi:<https://doi.org/10.1016/j.neuroimage.2011.10.035>, 20
2129 YEARS OF fMRI.

2130 [177] Y. Katsumi, N. Kamona, J. Zhang, J. G. Bunce, J. B. Hutchinson,
2131 M. Yarossi, E. Tunik, K. S. Quigley, B. C. Dickerson, L. F. Barrett,
2132 Functional connectivity gradients as a common neural architecture for
2133 predictive processing in the human brain, BioRxiv (2021).

2134 [178] R. S. Marken, et al., Controlled variables: Psychology as the center
2135 fielder views it, American Journal of Psychology 114 (2001) 259–282.

2136 [179] H. H. Yin, Restoring purpose in behavior, in: Computational and
2137 robotic models of the hierarchical organization of behavior, Springer,
2138 2013, pp. 319–347.

2139 [180] T. Hosoya, S. A. Baccus, M. Meister, Dynamic predictive coding by
2140 the retina, Nature 436 (2005) 71–77. doi:10.1038/nature03689.

2141 [181] B. Liu, A. Hong, F. Rieke, M. B. Manookin, Predictive en-
2142 coding of motion begins in the primate retina, Nature Neuro-
2143 science 24 (2021) 1280–1291. URL: [http://dx.doi.org/10.1038/](http://dx.doi.org/10.1038/s41593-021-00899-1)
2144 [s41593-021-00899-1](http://dx.doi.org/10.1038/s41593-021-00899-1). doi:10.1038/s41593-021-00899-1.

- 2145 [182] N. Qian, J. Zhang, Neuronal Firing Rate As Code Length: a Hypoth-
2146 esis, *Computational Brain & Behavior* (2019) 34–53. doi:[10.1007/
2147 s42113-019-00028-z](https://doi.org/10.1007/s42113-019-00028-z).
- 2148 [183] B. R. Dworkin, *Learning and physiological regulation*, University of
2149 Chicago Press, 1993.
- 2150 [184] M. Cabanac, Physiological Role of Pleasure, *Science* (1971) 1103–1107.
- 2151 [185] L. F. Barrett, E. Bliss-Moreau, Affect as a psychological primitive,
2152 *Advances in experimental social psychology* 41 (2009) 167–218.
- 2153 [186] C. Klein, What do predictive coders want?, *Synthese* 195 (2018) 2541–
2154 2557.
- 2155 [187] O. Peters, The ergodicity problem in economics, *Nature*
2156 *Physics* 15 (2019) 1216–1221. URL: [http://dx.doi.org/10.1038/
2157 s41567-019-0732-0](http://dx.doi.org/10.1038/s41567-019-0732-0). doi:[10.1038/s41567-019-0732-0](https://doi.org/10.1038/s41567-019-0732-0).
- 2158 [188] K. Friston, The free-energy principle: a unified brain theory?, *Nature*
2159 *Reviews Neuroscience* 11 (2010) 127–138. doi:[10.1038/nrn2787](https://doi.org/10.1038/nrn2787).
- 2160 [189] S. Denève, C. K. Machens, Efficient codes and balanced networks,
2161 *Nature Neuroscience* 19 (2016) 375–382. doi:[10.1038/nn.4243](https://doi.org/10.1038/nn.4243).
- 2162 [190] Y. Huang, R. P. Rao, *Predictive coding*, Wiley Interdisciplinary Re-
2163 views: *Cognitive Science* 2 (2011) 580–593.
- 2164 [191] J. Zhang, O. Abiose, Y. Katsumi, A. Touroutoglou, B. C. Dicker-
2165 son, L. F. Barrett, Intrinsic Functional Connectivity is Organized as
2166 Three Interdependent Gradients, *Scientific Reports* 9 (2019) 1–14.
2167 URL: <http://dx.doi.org/10.1038/s41598-019-51793-7>. doi:[10.
2168 1038/s41598-019-51793-7](https://doi.org/10.1038/s41598-019-51793-7).
- 2169 [192] M. Andrews, The math is not the territory: navigating the
2170 free energy principle, *Biology & Philosophy* 36 (2021) 1–
2171 19. URL: <https://doi.org/10.1007/s10539-021-09807-0>. doi:[10.
2172 1007/s10539-021-09807-0](https://doi.org/10.1007/s10539-021-09807-0).

- 2173 [193] S. H. Yeo, D. W. Franklin, D. M. Wolpert, When Optimal Feedback
2174 Control Is Not Enough: Feedforward Strategies Are Required for Op-
2175 timal Control with Active Sensing, PLoS Computational Biology 12
2176 (2016) 1–22. doi:[10.1371/journal.pcbi.1005190](https://doi.org/10.1371/journal.pcbi.1005190).
- 2177 [194] R. A. Adams, S. Shipp, K. J. Friston, Predictions not com-
2178 mands: Active inference in the motor system, Brain Structure
2179 and Function 218 (2013) 611–643. doi:[10.1007/s00429-012-0475-5](https://doi.org/10.1007/s00429-012-0475-5).
2180 [arXiv:arXiv:1011.1669v3](https://arxiv.org/abs/1011.1669v3).
- 2181 [195] W. Penny, K. Stephan, A Dynamic Bayesian Model of Homeostatic
2182 Control, Adaptive and Intelligent Systems (2014) 60–69.
- 2183 [196] W. James, The principles of psychology, se - 2 vol ed.,
2184 Henry Holt and Company, New York, NY, USA, 1890. URL:
2185 [http://link.library.utoronto.ca/eir/EIRdetail.cfm?](http://link.library.utoronto.ca/eir/EIRdetail.cfm?Resources__ID=1117305&T=F&Vol=1)
2186 [Resources__ID=1117305&T=F&Vol=1](http://link.library.utoronto.ca/eir/EIRdetail.cfm?Resources__ID=1117305&T=F&Vol=1).
- 2187 [197] E. Todorov, Finding the most likely trajectories of optimally-controlled
2188 stochastic systems, IFAC Proceedings Volumes (IFAC-PapersOnline)
2189 44 (2011) 4728–4734. doi:[10.3182/20110828-6-IT-1002.01704](https://doi.org/10.3182/20110828-6-IT-1002.01704).

2190 A. Glossary

2191 **advantage function** A function quantifying the relative decision value of
2192 visiting a given state, compared to the (estimated) time-average decision
2193 value across the course of a behavior. 48

2194 **allostasis** Brain-centered predictive regulation in which the brain anticipates
2195 the needs of the body and attempts to meet those needs before they arise. 4

2196 **control error** The difference between a measured or actual state and the
2197 reference state. 24

2198 **controller** A physical system whose behavior transforms signals encoding a
2199 reference trajectory into signals encoding controls to drive a plant. 20

2200 **controls** A signal or signals that drive the behavior of a plant by changing
2201 its state. 21

2202 **control theory** The engineering discipline of driving a physical system to-
2203 wards a desired trajectory. 5, 20

2204 **disturbance** An outside factor that can disrupt the state of a physical sys-
2205 tem. 21

2206 **ergodicity** A condition of a stochastic system, or a probabilistic model of
2207 a system, in which averaging together a series of measurements taken over
2208 time is equivalent to averaging together the same number of measurements
2209 taken independently at the same time. 30

2210 **exafferent** Afferent sensory signals arising from causes external to the body
2211 or outside the physiological control of the nervous system. 25

2212 **exteroception** The set of processes by which the nervous system takes in,
2213 integrates, and infers the causes of signals from outside the body. 3

2214 **feedback control** The adjustment of controls in real time based on mea-
2215 surements of plant state fed back into the controller. 4, 22, 24, 46

2216 **flow** A variable in a (physiological) dynamical system that represents the
2217 rate of a process, counted in the ratio of physical units to temporal units.

2218 **controlled process** A particular flow variable that is manipulated by
2219 homeostatic or allostatic mechanisms to maintain the value of a regulated
2220 variable within a desired range. 7, 8, 14

2221 **generative model** A probability model of the joint distribution over unob-
2222 served decision variables, and observed variables given the unobserved deci-
2223 sion variables, from which novel instances can be sampled. 12, 38

2224 **internal model** A physical system that partially and approximately en-
 2225 codes the dynamical structure of the plant to be controlled, or potentially a
 2226 probability distribution over such structures. [3](#), [4](#), [23](#)

2227 **forward model** An internal model that makes predictions forwards in
 2228 time, on the basis of a present state. [23](#), [46](#)

2229 **inverse model** An internal model that makes predictions backwards in
 2230 time, on the basis of an estimated or desired future state. [23](#)

2231 **interoception** The set of processes by which the nervous system takes in,
 2232 integrates, and makes meaning of sensory signals originating within the body.
 2233 [3](#)

2234 **measurement** A signal fed back from the plant to inform the controller of
 2235 the plant state. [22](#)

2236 **noise** Randomness or uncertainty intrinsic to physical systems which cannot
 2237 be measured or controlled at an infinitely fine scale..

2238 **measurement noise** Randomness or uncertainty intrinsic to the instru-
 2239 ments or sensors that take measurements and transmit them to the con-
 2240 troller. [23](#)

2241 **process noise** Randomness or uncertainty intrinsic to the plant as a phys-
 2242 ical system. [23](#)

2243 **objective function** A function from actual or estimated plant states to
 2244 real numbers, in which higher numbers denote better agreement with the
 2245 reference trajectory. [28](#)

2246 **operating point** The point on a response curve at which the response avail-
 2247 able to a local perturbation is greatest, often but not necessarily found at
 2248 the center parent. [10](#)

2249 **plant** A physical system to be controlled. [20](#)

2250 **prediction error** The difference between a measurement or actual state,
 2251 and the predictive estimate of that state. [23](#)

2252 **probabilistic graphical model** A probability model in which the nodes and
 2253 arrows of a graph express the conditional independence structure between
 2254 random variables see. [42](#)

2255 **reafferent** Afferent sensory signals arising from the consequences of move-
 2256 ments or self-caused physiological changes. [25](#)

2257 **reference distribution** A probability distribution whose probability mass
 2258 or density at each point corresponds to the relative robustness of a physio-
 2259 logical controller to perturbations from that point. [36](#)

2260 **reference trajectory** The desired trajectory of evolution through a state-
 2261 space for plant behavior over time. [20](#)

2262 **response curve** The plotted curve, usually S-shaped, showing a response
 2263 to a stimulus as a function of the stimulus quantity or intensity. [9](#)

2264 **capacity curve** The response curve of a centrally regulated physiologi-
 2265 cal variable or reflex, which compels a central regulatory response when
 2266 dysregulated. [9](#)

2267 **generalized capacity curve** The capacity curve of a controlled physi-
 2268 ological flow, with an inflection point not identical to its center. [15](#)

2269 **gain** The parameter to a response curve determining its slope around the
 2270 central point. Greater gain implies a smaller distance to an asymptote and
 2271 a narrower adaptive range. [9](#)

2272 **saturation value** The value of an input variable (such as a regulated vari-
 2273 able) for which a process (such as a controlled variable's underlying process)
 2274 can yield no further increase in response. [9](#)

2275 **set point** Fixed, specific points in the quantitative state-space of physiolog-
 2276 ical variable to which regulatory systems work to return that variable. [4](#),
 2277 [7](#)

2278 **settling point** The point at which a stock-and-flow system's stock reser-
 2279 voir settles for any given level of the passively unregulated inflow or outflow
 2280 variables. [14](#)

2281 **settling range** A range throughout which a regulated variable can settle
 2282 freely under physiologically unregulated (but perhaps behaviorally regulated)
 2283 inflow and outflow, without triggering an active physiological response. *see*
 2284 [settling point](#), [7](#)

2285 **stability** A property of the coupled dynamics of a controller and a plant,
 2286 under which control error will eventually shrink arbitrarily close to zero after
 2287 a disturbance.. [24](#)

2288 **state** A vector of real numbers that determine how plant behavior will evolve
 2289 over time. [21](#)

2290 **state estimate** Estimates of plant state based on measurements, including
 2291 measurements accumulated over time. [23](#)

2292 **stochastic optimal control** A variety of control theory in which we model
 2293 all forms of noise and uncertainty using probability distributions, and write

2294 the reference trajectory in terms of finding the maximum of a function of the
 2295 estimated plant state. 28
 2296 **stock** A variable in a (physiological) dynamical system that represents a
 2297 quantity, counted in physical units.
 2298 **regulated resource** A particular stock variable that is maintained at or
 2299 near a stable level by homeostatic or allostatic mechanisms in the body. 6,
 2300 8, 14
 2301 **threshold value** The value of an input variable (such as a regulated variable)
 2302 for which a process (such as a controlled variable's underlying process) can
 2303 yield no further reduction in response. 9
 2304 **transfer function** The ratio of a system output to its input, written in the
 2305 time-independent frequency domain via a mathematical transformation. 9
 2306 **value function** The probabilistic average, over estimated plant states, of
 2307 the sum of objective function values into the indefinite future. 28, 35, 53
 2308 **Bellman equation** Equation that defines the optimal value function re-
 2309 cursively over timesteps. 28, 44
 2310 **viscerosensory signaling** Afferent sensory signaling from the innervated
 2311 viscera. 3

2312 B. How ergodicity interacts with control

2313 B.1. Ergodicity: whether uncertainty and noise make a difference over time

The concept of *ergodicity* formalizes the equivalence between present and future uncertainty or noise. In an ergodic system, averaging together a series of measurements taken over time is equivalent to averaging together the same number of measurements taken simultaneously with separate instruments. If we were to assume an objective function $f(x)$ under an ergodic probability model $p(x_t)$, we would write that

$$\lim_{T \rightarrow \infty} \mathbb{E}_p \left[\frac{1}{T} \sum_{t=1}^T f(x_t) \right] = \mathbb{E}_{p(x_t)} [f(x_t)]. \quad (\text{B.1})$$

2314 This equation says that taking individual measurements $f(x_t)$ over any hy-
 2315 pothetical period of time T and averaging them together (the left side of the
 2316 equation) will yield the same result as just calculating the average measure-
 2317 ment immediately (the right side of the equation).

2318 Conceptually, if we would like to know the average speed at which a
2319 particular kind of car travels on a highway, we could either track a single such
2320 vehicle over an extended period of time, or measure the speed of multiple
2321 vehicles (of the same kind) at the same time. In either case, as long as we took
2322 sufficiently many measurements, we would obtain the same result, whether
2323 we calculate by dividing by time or by the number of vehicles. However,
2324 the very obscurity of this thought experiment should hint to us that many
2325 real-world measurements are *not* ergodic.

2326 Situations that are not ergodic include those with periodic structure, or
2327 with irreversible changes. Our real lives are, therefore, filled with non-ergodic
2328 decisions to make: the cycles of day and night are non-ergodic; our develop-
2329 ment and aging processes are non-ergodic; injury and death are non-ergodic.
2330 Insofar as our internal environment has its own cycles (breathing, heartbeat,
2331 eating and drinking, etc.) it too is non-ergodic. The earliest brains con-
2332 tained body-clocks that synchronized bodily cycles to environmental cycles
2333 [23]; they were non-ergodic as well. This is why we invested the extra ef-
2334 fort above of detailing a variety of control theory able to model non-ergodic
2335 decision-making.

2336 In contrast, many typical behavioral experiments involve their partici-
2337 pants in ergodic decision-making problems: we choose as experimenters to
2338 treat averaging within-participant measurements across time as equivalent to
2339 averaging between-participant measurements. In fact, an experimental par-
2340 ticipant can become fatigued, bored, excited, hungry, or thirsty through the
2341 course of the experiment. If given food and drink, they can cycle between
2342 hunger, satiation, and then hunger again. Assuming ergodicity in experi-
2343 mental design implies ignoring all the bodily and environmental factors that
2344 we do not design, and which do not conform to our assumption.

2345 This has practical implications for decision-making processes in the brain,
2346 as well. Only in an ergodic problem can the brain rightfully equate averag-
2347 ing over the history of a real behavior with averaging over some population
2348 of mental simulations. This is why decision-making often involves mental
2349 simulations over the entire course of a possible behavior, allowing the brain
2350 to account for cyclical or irreversible events. In fact, when we give people
2351 non-ergodic decision making problems in experiments, the evidence suggests
2352 subjects may solve them “the right way”, imagining the temporal course of
2353 possible futures [187].

2354 *B.2. Formulating perception and control ergodically leads to the Free Energy*
2355 *Principle (FEP)*

2356 The brain’s internal model allows it to regulate its body in a world
2357 full of uncertainty by means of visceromotor commands, obtaining feedback
2358 by anticipating the commands’ viscerosensory consequences as interoceptive
2359 predictions. These interoceptive predictions implement “top-down” sensory
2360 simulations, fashioned from past experiences, that are continually compared
2361 against information received from the viscerosensory surfaces about the ac-
2362 tual state of the body in the world.

2363 These hypotheses can be integrated into an emerging theoretical con-
2364 sensus, dubbed predictive processing [17, 188, 145, 189, 20]. “Top down”
2365 sensory predictions continuously anticipate events within the body and out-
2366 side it that are sensed via its sensory surfaces. Sensory prediction signals
2367 cascade across multiple gradients within the brain, including across the cere-
2368 bral cortex [190, 191], cerebellum [174, 175], and hippocampus [10], as well
2369 as across all levels of the neuraxis, involving hypothalamus, basal ganglia,
2370 superior colliculus and various midbrain and brainstem structures [26]. In
2371 fact, it is possible that every sensory neuron, in effect, receives some form
2372 of prediction signalling from some of the neurons projecting to it, and sends
2373 prediction errors²³ to other neurons via its own projections [189]. “Bottom-
2374 up” information coming from the sensory surfaces (such as the retina, the
2375 olfactory bulb, and the lamina I spinothalamic tract) acts as sensory pre-
2376 diction error signals and is visualized as the dashed arrows in Figure 6. In
2377 this way, the brain continuously maintains a simulation of the sensory en-
2378 vironment inside the body, as well as the relevant aspects of the sensory
2379 environment outside the body, and updates that running simulation accord-
2380 ing to computed error. The updated simulations constitute approximately
2381 optimal posterior inferences about the likely causes of sensory events.

Among predictive processing hypotheses of brain function, the free-energy
principle (FEP) has enjoyed a particular popularity [188, 111, 192]. Under
the ergodic assumptions of the Free-Energy Principle as formulated by Friston et al. [111], organisms are hypothesized to minimize the entropy of their

²³Formally, score functions, the gradients of a log-likelihood.

sensory outcomes

$$H(o) = \lim_{T \rightarrow \infty} -\frac{1}{T} \int_0^T dt \log p(o(t)) \quad (\text{B.2})$$

$$= -\log \int ds p(o(t), s(t)) , \quad (\text{B.3})$$

2382 where $p(o(t), s(t))$ is interpreted as a generative model of sensory outcomes
 2383 $o(t)$ with hidden variables $s(t)$, evolving in continuous time. Note that the
 2384 first integral is a time average (denoted by dt) and the second an ensemble
 2385 average (denoted by ds). The authors of Friston et al. [111] propose that
 2386 organisms carry out this task by minimizing an upper bound to the sen-
 2387 sory entropy, called the variational free-energy, which also appears in related
 2388 predictive-coding models of perception.

2389 The equivalence of the time-average in Equation B.2 and the sample av-
 2390 erage in Equation B.3 is precisely the assumption of ergodicity discussed in
 2391 B.1. This assumption renders the time-average equivalent to the instanta-
 2392 neous sample-average, making any one moment in time as good as any other
 2393 for purposes of optimizing the free energy (or the entropy). Ergodicity also
 2394 implies there exists a long-run stationary distribution, with no dependence
 2395 on the start state or the particular sequencing of states $s(t)$ across time.

Under such assumptions, we can formulate an extraordinarily elegant control formalism just by writing the sensory outcomes as functions of action $a(t)$, yielding an objective function

$$H(o) = -\log \int ds p(o(t), a(t), s(t)) ,$$

and a variational bound in terms of approximate posterior beliefs $q(s(t))$ and their “free energy” $\mathcal{F}$,

$$-\log p(o(t)) \leq \mathcal{F} [q(s(t)), a(t)] .$$

2396 The variational bound can then be used as a proxy for the true objective,
 2397 and implies that action should descend the gradient of free energy to min-
 2398 imize long-run sensory entropy. As a side-effect of giving action the same
 2399 objective functional as sensory prediction, this *active inference* formulation
 2400 under the free-energy principle implies a (partial, contextual) equivalence
 2401 between precision and value. Accumulating evidence becomes the only true

good, while losing it becomes the only true ill, and precision is the parameter whose optimization most easily affords increasing evidence. Precision here is mathematical equivalent (up to some constant factors) to the gain or inverse-tolerance we described as a parameter to capacity curves in Section 2.2. We employ the separate terminology to make sure it is understood that precision measures the confidence of a probabilistic estimate of some quantity, while the gain or tolerance measures the steepness or shallowness of a regulatory response.

Within this paradigm, active inference modeling has taken two general forms. *Motor* active inference (mActIf) tends to assume (local) ergodicity, so that local hill-climbing on the steady-state density becomes the best control strategy. In terms of classical control theory, when the reference signal is a set-point and disturbances are independently random across time, reactive behavior and homeostatic reflex arcs will suffice to maintain optimal regulation of physiology. In contrast, *decision* active inference (dActInf) maintains the appearance of ergodic homeostasis within the internal environment by writing a homeostatic distribution or “prior preferences” $p(o \mid m)$ (the notation of conditioning on m to write prior preferences comes from Friston et al. [112]) and then using predictive decision-making to make internal states conform to that distribution.

Under decision active inference, mental simulations “try out” potential policies, considering their effects over a finite time horizon, and policies thereby compete to control action based on their expected free energy. dActInf thus implements what control engineers would call model-predictive control with an information-theoretic objective [113]. Similarly, the “prior preferences” function as what we would call homeostatic set-points and tolerances, as addressed in Stephan et al. [159], Corcoran and Hohwy [43], Pezzulo et al. [101], and many similar works. Under ergodic assumptions, time-averaging may be invoked to interpret all physiological changes as (survivable) deviations from fixed homeostatic set-points with fixed tolerances.

Ergodicity yields an especially elegant formulation of control because it collapses the distinction between information from the past and information from the future. The precise reason that stochastic optimal control problems are hard to solve is that computing the true optimal decision require knowledge of the future – not estimates, knowledge. Both mathematically and intuitively, you cannot really know you have chosen the best course of action without knowing the course of future events ahead of time. Without that knowledge, the best that any real controller, or brain, can do is to estimate

the necessary quantities using all information available, make corrections as new information reveals itself, and wait until the notional end of a behavioral episode to evaluate performance retrospectively. Even small errors or misestimates compound, the further into the future you try to predict, because they cannot be corrected online by present sense-data. Unfortunately, a stochastic optimal control problem of the sort implied by the time-averaging we discussed in Section 4.2 requires prediction indefinitely far into the future.

Of course, this is what the brain’s internal model is for: refining information contained in past sensory signals into the optimal estimate of the future. The optimal estimate of the future will probably still be off by some amount, but it nonetheless represents the best possible estimate that can be made given the information available. The actual brain never makes a truly optimal estimate, of course: it approximates the optimal estimate, effectively approximating an approximation of the unavailable future information. Even so, this approximately optimal estimation enables feedforward, prospective control that performs drastically better than waiting to react until the body is actually challenged or harmed [193]. If you are walking on a dark road, and you see something that looks vaguely like a car careening towards you, you don’t wish for a more optimal estimate, you get out of the way. It is only in this sense of “prediction”, prediction without real-time correction, that predictive processing can be suitable for modeling sequential decision problems.

This has implications for the low-level motor control strategy employed in regulating the viscera. Previous motor active inference accounts of visceromotor action have noted that visceromotor cortical areas mostly lack integration of ascending prediction errors, and hypothesized that they “function more like deterministic models of actions that are to be executed” [19]. They then extend the analysis of Adams et al. [194] to the visceromotor domain: they conceptualize visceromotor efferent signaling as specifying visceral reference coordinates rather than “predictions” to be revised in a Bayesian sense [25, 115]. Continuing the analogy, they suggest that the uncertainty of visceromotor predictions specifies physiological tolerances [195, 159]. However, this account of visceromotor active inference implicitly relied on the ergodicity assumption embedded in motor active inference as a theory. Our model will relax this ergodicity assumption, while providing an alternative explanation for the anatomical observations motor active inference can explain. In short, we could consider our model a probabilistic way of encoding hierarchical referent control [89] and the ideomotor principle [196]. We gratefully

2478 thank investigators such as Corcoran et al. [102] who have considered al-
 2479 lostasis explicitly for setting the stage for our work, and emphasize that the
 2480 lens of control theory is what enables us to separate the idea of a reference
 2481 coordinate from that of a prediction.

2482 C. APIC model derivation details

2483 C.1. A hierarchical generative model with reference distributions

The APIC model assumes a hierarchical generative model designed to explain sensory outcomes o_t using hidden states $s_t^{(1)}, \dots, s_t^{(L)}$ across some number $L \in [2, \infty)$ of levels. At each level of hierarchy $l \in [1..L]$, the variables $\rho_t^{(0)}, \dots, \rho_t^{(L-1)}$ parameterize a reference distribution, written

$$\begin{aligned} & p(o_t; \rho_t^{(0)}), \\ & p(s_t^{(1)}; \rho_t^{(1)}), \\ & \dots \\ & p(s_t^{(L-1)}; \rho_t^{(L-1)}). \end{aligned}$$

The parameters for reference distributions themselves come from multiple levels of *policies*, distributions π over lower-level goals or actions, conditioned upon higher-level states and goals:

$$\begin{aligned} & \pi(a_t \mid o_t, \rho_t^{(0)}), \\ & \pi(\rho_t^{(0)} \mid s_t^{(1)}, \rho_t^{(1)}), \\ & \dots \\ & \pi(\rho_t^{(L-1)} \mid s_t^{(L)}). \end{aligned}$$

At the bottom level, a joint likelihood predicts the exteroceptive, interoceptive, and somatosensory observations together, based upon the lowest-level hidden states and reference signals,

$$p_\theta(o_t \mid s_t^{(1)}, \rho_t^{(0)}).$$

The prior densities for each hierarchical level $l \in [1 \dots L]$ are written

$$p_\theta(s_1^{(L)}) \qquad p_\theta(s_1^{(l)} \mid s_1^{(l+1)}),$$

and transition densities for each hierarchical level $l \in [1 \dots L]$ are written

$$p_\theta(s_t^{(L)} \mid s_{t-1}^{(L)}) \quad p_\theta(s_t^{(l)} \mid s_{t-1}^{(l)}, \rho_{t-1}^{(l)}, s_t^{(l+1)}).$$

2484 Figure 10 displays the resulting complete graphical model for $L = 4$.

The above prior densities and likelihood imply that by Bayes' rule there exists a posterior distribution

$$p_\theta(s_t^{(1)}, \rho_t^{(0)} \mid o_t) = \frac{p_\theta(o_t \mid s_t^{(1)}, \rho_t^{(0)})p_\theta(s_t^{(1)}, \rho_t^{(0)})}{p_\theta(o_t)},$$

2485 which APIC allows approximating by more-or-less any means. Since policies
2486 π also took the form of distributions above, Bayesian inference implies a
2487 posterior distribution over reference signals ρ and actions a_t as well.

The *complete* state of the generative model can be abbreviated as

$$\begin{aligned} s_t^{(1:L)} &= (s_t^{(1)}, \dots, s_t^{(L)}) \\ \rho_t^{(1:L-1)} &= (\rho_t^{(1)}, \dots, \rho_t^{(L-1)}), \\ \mathbf{x}_t &= (a_t, o_t, s_t^{(1:L)}, \rho_t^{(1:L-1)}) \\ \mathbf{x}_{1:T} &= (\mathbf{x}_1, \dots, \mathbf{x}_T). \end{aligned}$$

The complete reference distribution, across all hierarchical levels, can be written

$$p(\mathbf{x}_t; \rho_t^{(1:L-1)}) = p(o_t; \rho_t^{(0)}) \prod_{l=1}^{L-1} p(s_t^{(l)}; \rho_t^{(l)}), \quad (\text{C.1})$$

and the complete transition dynamics as

$$\begin{aligned} p_\theta(\mathbf{x}_t \mid \mathbf{x}_{t-1}) &= p_\theta(o_t \mid s_t^{(1)}, \rho_t^{(0)})\pi(a_t \mid s_t^{(1)}, \rho_t^{(0)})p_\theta(s_t^{(L)} \mid s_{t-1}^{(L)})\pi(\rho_t^{(L-1)} \mid s_t^{(L)}) \\ &\quad \prod_{l=1}^{L-1} p_\theta(s_t^{(l)} \mid s_{t-1}^{(l)}, \rho_{t-1}^{(l)}, s_t^{(l+1)})\pi(\rho_t^{(l)} \mid s_t^{(l+1)}, \rho_t^{(l+1)}), \end{aligned}$$

where

$$p_\theta(s_t^{(1:L)} \mid s_{t-1}^{(1:L)}, \rho_t^{(1:L-1)}) = p_\theta(s_t^{(L)} \mid s_{t-1}^{(L)}) \prod_{l=L-1}^1 p_\theta(s_t^{(l)} \mid s_{t-1}^{(l)}, \rho_{t-1}^{(l)}, s_t^{(l+1)}).$$

2488 *C.2. Full derivation of optimal value function and transition dynamics*

The differential Bellman equation is the defining principle of stochastic optimal control, an equation for the long-run value function. In the infinite-time, average-reward setting, assuming a global capture rate of $\bar{J}(\mathbf{x}_0)$ for a whole behavior, it can be written

$$\tilde{V}^*(\mathbf{x}_t, \mathbf{x}_{t-1}) = \mathcal{J}(\mathbf{x}_t, \mathbf{x}_{t-1}) - \tilde{J}(\mathbf{x}_0) + \max_{p_\theta(\mathbf{x}_{t+1}|\mathbf{x}_t)} \mathbb{E}_{p_\theta(\mathbf{x}_{t+1}|\mathbf{x}_t)} \left[\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right]$$

with the maximization taking place with respect to achievable (by varying the actions) transition distributions. $\tilde{V}^*(\mathbf{x}_t, \mathbf{x}_{t-1})$ then denotes the best achievable value for any transition between individual states. Combining that expression with the transition dynamics as a Boltzmann or “softly maximizing” distribution will yield a generative model of the optimal transition dynamics, as seen in Equation 15:

$$q_\phi^*(\mathbf{x}_{t+1} | \mathbf{x}_t) = \frac{\exp \left(\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right) p_\theta(\mathbf{x}_{t+1} | \mathbf{x}_t)}{\mathbb{E}_{\mathbf{x}_{t+1} \sim p_\theta(\mathbf{x}_{t+1}|\mathbf{x}_t)} \left[\exp \left(\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right) \right]}.$$

2489 This transition distribution optimally trades off between probable transitions
 2490 and valuable transitions. The optimal transition distribution’s mode, when it
 2491 can be calculated or approximated, corresponds to the most likely trajectory
 2492 for an optimal agent (with the assumed goals) to follow [197]. Writing it with
 2493 a q_ϕ rather than a p_θ denotes that it will be used as an observation-driven
 2494 feedback controller, or more formally an importance sampling proposal. Just
 2495 as the optimal transition distribution provides a generative model of an opti-
 2496 mally controlled system, when used as an observation-driven proposal it also
 2497 provides the optimal feedback controller.

Substituting the optimal feedback controller back into the differential Bellman equation will yield

$$\tilde{V}^*(\mathbf{x}_t, \mathbf{x}_{t-1}) = \mathcal{J}(\mathbf{x}_t, \mathbf{x}_{t-1}) - \tilde{J}(\mathbf{x}_0) + \mathbb{E}_{q_\phi^*(\mathbf{x}_{t+1}|\mathbf{x}_t)} \left[\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right],$$

which can be further expanded by substituting Equation 14 for $\mathcal{J}$:

$$\begin{aligned} \tilde{V}^*(\mathbf{x}_t, \mathbf{x}_{t-1}) = & J(\mathbf{x}_t, \mathbf{x}_{t-1}) - D_{\text{KL}} \left(q_\phi^*(\mathbf{x}_{t+1} | \mathbf{x}_t) \| p_\theta(\mathbf{x}_{t+1} | \mathbf{x}_t) \right) - \\ & \tilde{J}(\mathbf{x}_0) + \mathbb{E}_{q_\phi^*(\mathbf{x}_{t+1}|\mathbf{x}_t)} \left[\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right]. \end{aligned}$$

First, however, the information divergence penalty must be expanded in order to justify Equation 17. First the divergence is written out in terms of its definition

$$\begin{aligned} D_{\text{KL}} \left(q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t) \parallel p_{\theta}(\mathbf{x}_{t+1} \mid \mathbf{x}_t) \right) &= \mathbb{E}_{q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\log \frac{q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t)}{p_{\theta}(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \right] \\ &= \mathbb{E}_{q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\log q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t) \right] - \mathbb{E}_{q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\log p_{\theta}(\mathbf{x}_{t+1} \mid \mathbf{x}_t) \right], \end{aligned}$$

and then the definition of q_{ϕ}^* is substituted into the expectation of the logarithm,

$$\begin{aligned} \mathbb{E}_{q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\log q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t) \right] &= \\ \mathbb{E}_{q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\log \frac{\exp \left(\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right) p_{\theta}(\mathbf{x}_{t+1} \mid \mathbf{x}_t)}{\mathbb{E}_{\mathbf{x}_{t+1} \sim p_{\theta}(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\exp \left(\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right) \right]} \right] \\ &= \mathbb{E}_{q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\log \frac{\exp \left(\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right) p_{\theta}(\mathbf{x}_{t+1} \mid \mathbf{x}_t)}{\mathbb{E}_{\mathbf{x}_{t+1} \sim p_{\theta}(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\exp \left(\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right) \right]} \right] \\ &= \mathbb{E}_{q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) + \log p_{\theta}(\mathbf{x}_{t+1} \mid \mathbf{x}_t) \right] - \log \mathbb{E}_{\mathbf{x}_{t+1} \sim p_{\theta}(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\exp \left(\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right) \right] \\ &= \mathbb{E}_{q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\log q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t) \right] = \mathbb{E}_{q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right] + \\ &\quad \mathbb{E}_{q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\log p_{\theta}(\mathbf{x}_{t+1} \mid \mathbf{x}_t) \right] - \log \mathbb{E}_{\mathbf{x}_{t+1} \sim p_{\theta}(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\exp \left(\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right) \right]. \end{aligned}$$

This substitutes into the equation for the divergence to imply that

$$\begin{aligned} D_{\text{KL}} \left(q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t) \parallel p_{\theta}(\mathbf{x}_{t+1} \mid \mathbf{x}_t) \right) &= \mathbb{E}_{q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right] + \\ &\quad \mathbb{E}_{q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\log p_{\theta}(\mathbf{x}_{t+1} \mid \mathbf{x}_t) \right] - \log \mathbb{E}_{\mathbf{x}_{t+1} \sim p_{\theta}(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\exp \left(\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right) \right] - \\ &\quad \mathbb{E}_{q_{\phi}^*(\mathbf{x}_{t+1} \mid \mathbf{x}_t)} \left[\log p_{\theta}(\mathbf{x}_{t+1} \mid \mathbf{x}_t) \right]. \end{aligned}$$

The second and last terms of this equation are identical, and so cancel. This leaves a divergence of

$$D_{\text{KL}}(q_{\phi}^*(\mathbf{x}_{t+1} | \mathbf{x}_t) || p_{\theta}(\mathbf{x}_{t+1} | \mathbf{x}_t)) = \mathbb{E}_{q_{\phi}^*(\mathbf{x}_{t+1} | \mathbf{x}_t)} \left[\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right] - \log \mathbb{E}_{\mathbf{x}_{t+1} \sim p_{\theta}(\mathbf{x}_{t+1} | \mathbf{x}_t)} \left[\exp \left(\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right) \right],$$

and therefore a (subtractive) divergence penalty as seen in Equation 17:

$$- D_{\text{KL}}(q_{\phi}^*(\mathbf{x}_{t+1} | \mathbf{x}_t) || p_{\theta}(\mathbf{x}_{t+1} | \mathbf{x}_t)) = -\mathbb{E}_{q_{\phi}^*(\mathbf{x}_{t+1} | \mathbf{x}_t)} \left[\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right] + \log \mathbb{E}_{\mathbf{x}_{t+1} \sim p_{\theta}(\mathbf{x}_{t+1} | \mathbf{x}_t)} \left[\exp \left(\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right) \right].$$

Substituting the above for the penalty in the differential Bellman equation will yield

$$\begin{aligned} \tilde{V}^*(\mathbf{x}_t, \mathbf{x}_{t-1}) &= J(\mathbf{x}_t, \mathbf{x}_{t-1}) - \tilde{J}(\mathbf{x}_0) + \cancel{\mathbb{E}_{q_{\phi}^*(\mathbf{x}_{t+1} | \mathbf{x}_t)} \left[\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right]} \\ &\quad - \cancel{\mathbb{E}_{q_{\phi}^*(\mathbf{x}_{t+1} | \mathbf{x}_t)} \left[\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right]} + \log \mathbb{E}_{\mathbf{x}_{t+1} \sim p_{\theta}(\mathbf{x}_{t+1} | \mathbf{x}_t)} \left[\exp \left(\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right) \right], \\ \tilde{V}^*(\mathbf{x}_t, \mathbf{x}_{t-1}) &= J(\mathbf{x}_t, \mathbf{x}_{t-1}) - \tilde{J}(\mathbf{x}_0) + \log \mathbb{E}_{\mathbf{x}_{t+1} \sim p_{\theta}(\mathbf{x}_{t+1} | \mathbf{x}_t)} \left[\exp \left(\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right) \right], \end{aligned} \tag{C.2}$$

2498 the smooth differential Bellman equation. This equation is smoothly maxi-
 2499 mizing instead of exactly maximizing, and thus more likely to be compatible
 2500 with neural stochasticity. Its recursive term makes the weak assumption
 2501 about the future that the trajectory will continue as planned according to
 2502 the generative action concept.

This equation can be shortened by applying the exponential function to both its sides, obtaining

$$\exp \left(\tilde{V}^*(\mathbf{x}_t, \mathbf{x}_{t-1}) \right) = \exp \left(J(\mathbf{x}_t, \mathbf{x}_{t-1}) - \tilde{J}(\mathbf{x}_0) \right) \cdot \mathbb{E}_{\mathbf{x}_{t+1} \sim p_{\theta}(\mathbf{x}_{t+1} | \mathbf{x}_t)} \left[\exp \left(\tilde{V}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right) \right],$$

and further shortened by defining the exponential of the value function as the *desirability* function,

$$\tilde{Z}^*(\mathbf{x}_t, \mathbf{x}_{t-1}) = \exp \left(\tilde{V}^*(\mathbf{x}_t, \mathbf{x}_{t-1}) \right),$$

to finally obtain

$$\tilde{Z}^*(\mathbf{x}_t, \mathbf{x}_{t-1}) = \exp \left(J(\mathbf{x}_t, \mathbf{x}_{t-1}) - \tilde{J}(\mathbf{x}_0) \right) \mathbb{E}_{\mathbf{x}_{t+1} \sim p_\theta(\mathbf{x}_{t+1}|\mathbf{x}_t)} \left[\tilde{Z}^*(\mathbf{x}_{t+1}, \mathbf{x}_t) \right].$$

In this form, the differential Bellman equation is linear, allowing its terms to be expanded recursively. For example, at time step $t = 1$

$$\begin{aligned} \tilde{Z}^*(\mathbf{x}_1, \mathbf{x}_0) &= \exp \left(J(\mathbf{x}_1, \mathbf{x}_0) - \tilde{J}(\mathbf{x}_0) \right) \mathbb{E}_{\mathbf{x}_2 \sim p_\theta(\mathbf{x}_2|\mathbf{x}_1)} \left[\tilde{Z}^*(\mathbf{x}_2, \mathbf{x}_1) \right] \\ \tilde{Z}^*(\mathbf{x}_2, \mathbf{x}_1) &= \exp \left(J(\mathbf{x}_2, \mathbf{x}_1) - \tilde{J}(\mathbf{x}_0) \right) \mathbb{E}_{\mathbf{x}_3 \sim p_\theta(\mathbf{x}_3|\mathbf{x}_2)} \left[\tilde{Z}^*(\mathbf{x}_3, \mathbf{x}_2) \right] \\ \tilde{Z}^*(\mathbf{x}_3, \mathbf{x}_2) &= \dots \end{aligned}$$

and so on. When the expectations are regrouped together and moved to the outermost level, the equation as a whole can be written over a trajectory of length T as

$$\tilde{Z}^*(\mathbf{x}_{1:T}; \mathbf{x}_0) = \mathbb{E}_{p_\theta(\mathbf{x}_{1:T}|\mathbf{x}_0)} \left[\exp \left(\sum_{t=1}^T J(\mathbf{x}_t, \mathbf{x}_{t-1}) - \tilde{J}(\mathbf{x}_0) \right) \right]. \quad (\text{C.3})$$

2503 This path-integral form for the value/desirability function gives its name to
 2504 the technique of *path-integral control*. Equation C.3 also provides exactly the
 2505 term required by Equation 15 for computing the optimal feedback controller
 2506 or transition distribution.