



How to count biological minds: symbiosis, the free energy principle, and reciprocal multiscale integration

Matthew Sims¹ 

Received: 11 March 2020 / Accepted: 14 September 2020 / Published online: 24 September 2020
© The Author(s) 2020, corrected publication 2020

Abstract

The notion of a physiological individuals has been developed and applied in the philosophy of biology to understand symbiosis, an understanding of which is key to theorising about the major transition in evolution from multi-organismality to multicellularity. The paper begins by asking what such symbiotic individuals can help to reveal about a possible transition in the evolution of cognition. Such a transition marks the movement from cooperating individual biological cognizers to a functionally integrated cognizing unit. Somewhere along the way, did such cognizing units simultaneously have cognizers as parts? Expanding upon the multiscale integration view of the Free Energy Principle, this paper develops an account of reciprocal integration, demonstrating how some coupled biological cognizing systems, when certain constraints are met, can result in a cognizing unit that is in ways greater than the sum of its cognizing parts. Symbiosis between *V. Fischeri* bacteria and the bobtail squid is used to provide an illustration this account. A novel manner of conceptualizing biological cognizers as gradient is then suggested. Lastly it is argued that the reason why the notion of ontologically nested cognizers may be unintuitive stems from the fact that our folk-psychology notion of what a cognizer is has been deeply influenced by our folk-biological manner of understanding biological individuals as units of reproduction.

Keywords Free energy principle · Active inference · Symbiosis · Nested Markov blankets · Multiscale integration · Cognitive evolution · Emergence · Physiological individuals

✉ Matthew Sims
s1670526@ed.ac.uk

¹ School of Philosophy, Psychology, and Language Sciences, University of Edinburgh, Edinburgh EH8 9AD, Scotland, UK

1 Introduction

Symbioses¹ are cooperative heterospecific associations in which each symbiont partner mutually benefits from (e.g., gaining nourishment, shelter, etc.) the presence of the other partner(s). These associations have presented an interesting problem case for the notion of *biological individuality*. What it is to be a biological individual refers to what it is that makes a living system a well delineated whole. Symbiosis raises the ontological question of whether beyond classifying each symbiont as a biological individual, there is a non-arbitrary manner of classifying the symbiotic assemblage itself as a well delineated biological unit. And if so, how can this non-arbitrary manner be spelled out? Answering this question is particularly important for at least two reasons; the first being that symbiosis is ubiquitous in nature; getting clear on the conditions under which a symbiotic association qualifies as a countable biological individual allows biological explanations to parse the world up in ways that ontologically cohere with the kinds of entities that populate the world. The second reason is that conceptualizing symbiotic associations as well delineated biological units plays a significant role in coming to understanding (and offer explanations of) the move from multi-organismality to multicellularity, an example of what has been called a major transition in evolution (Maynard-Smith and Szathmáry 1995). Under the assumption that multicellular organisms evolved from multi-organismal ensembles and that some of these ensembles were likely symbiotic associations, understanding this major transition in evolution requires a way to conceptualize biological individuals that is not restricted to units of reproduction or what are otherwise known as “Darwinian individuals” (see Godfrey-Smith 2013).

One way that philosophers have recently addressed this question is by conceptually developing and applying the notion of a *physiological individual*. A biological individual of this kind is a highly integrated functional unit, the heterogeneous parts of which cohere together through regulatory processes (e.g., metabolic, immuno-tolerance, etc.) so as to maintain the system’s integrity and resist environmental forces of decay (Pradeu 2011). Applying the notion of a physiological individual to symbioses suggests that at least some symbiotic ensembles qualify as biological individuals in their own right (Queller and Strassman 2009; Pradeu 2011; Godfrey-Smith 2013; Bouchard 2018). What qualifies some of these ensembles as physiological individuals is the high degree to which the symbionts are functionally integrated with one another. A fascinating and somewhat counterintuitive result of parsing biological world up in this manner is a kind of nesting of biological individuals in biological individuals; not only are the symbionts understood to be physiological individuals but also the symbiotic association comprised of those symbionts qualifies as a physiological individual.

This paper examines a related interesting question for any biologically informed cognitive science that takes two ideas seriously. The first idea is that cognitive processes are an enrichment of the organizational principles and properties definitive of living processes or what is called “the strong life-mind continuity thesis” (Godfrey-Smith 1996). The continuity thesis suggests that mind, because it is an enrichment of life,

¹ Throughout this paper, I will use the term symbiosis to refer specifically to mutualistic associations and not commensalism or parasitism.

might be exhibited in biological processes that prefigure those which cognitive science typically studies. It opens up the possibility that simple organisms (e.g., some bacteria and plants) may engage in cognitive processes that are not different in kind but different in degree to those processes that are paradigmatic of human cognition. The second idea is that some symbiotic associations actually do qualify as physiological individuals. Putting these ideas together the following question arises: if mind is an enriched version of life and some symbiotic associations indeed qualify as physiological individuals, is there any reason to think that some of these living associations might themselves be enriched enough so as to qualify as bio-cognizers?² This is not merely a question regarding a possible application of the extended mind thesis (i.e., the idea that cognitive processes can loop outside of a cognizing agent to envelop the use of external artefacts as scaffolding) (Clark and Chalmers 1998). Rather, this is a question about the kinds of physiological individuals that we can reasonably ascribe the term “cognizer” to. It is a question about how to count biological minds. The significance of this question lies in the fact that its answer may be used to shed light upon a possible *major transition in cognition*. Perhaps there are many such transitions: the move from reflex behaviour to sensorimotor coordination; the move from non-sentience to sentience; the move from individual intentionality to group intentionality. However the question which this paper will investigate concerns the transition away from short-term, slightly integrated cooperative interaction between simple cognizing physiological individuals towards the long-term, highly integrated cooperative interaction of the component parts of a complex cognizing symbiotic whole; each of the parts becoming more specialized in the function it plays in driving the behaviour of the whole integrated system.

Given the major transition in evolution from multi-organismality to multicellularity, and assuming that at least some of those component organisms that eventually constituted multicellular organisms were cognizers, symbiosis provides an opportunity to think about how an ensemble of minded physiological individuals became a single physiological individual with one mind. Like the nested-ness of physiological individuals, the possibility of a symbiotic mind suggests that a symbiotic cognizing unit may be comprised of nested symbiont cognizers. And much like the nested-ness of physiological individuals runs counter to our folk-biological conception of biological individuality, this idea that cognizers can have other cognizers as constituent parts seems to go against the grain of our folk-psychological conception of what cognizers are. One aim of this paper is to argue that the possibility of symbiotic minds provides a reason to believe that our folk-psychological conception of “cognizer” may be in need of serious revamping in this regard.

To support the notion of symbiotic minds and the notion of nested biological cognizers that falls out of it, this paper will develop and deploy the notion of *reciprocal multiscale integration* from within the Free Energy Principle. Reciprocal multiscale integration describes the case where each of two (or more) cognizing systems uses the other to provide evidence for its own model of the world (and itself acting in that world). It is when each cognizer reciprocally contributes to the cognitive processes of the other that an emergent action capacity arises. This is a capacity to act upon the envi-

² This paper’s scope will be strictly limited to investigating *biological* cognizers or what I will often call “bio-cognizers”.

ronment which outstrips the individual action capacities of each partner; the capacity of the whole being more than that possessed by the sum of its parts. It is this emergent action capacity that allows a symbiotic unit to provide evidence for its own model of the world and hence for its own continued existence. Reciprocal multiscale integration builds upon the multiscale integrationist view of the Free Energy Principle (Ramstead et al. 2019a, b; Hesp et al. 2019). This view uses the notion of nested Markov blankets to demonstrate that cognitive systems have a plurality of ontological boundaries, each relevant to the study of cognition. Reciprocal multiscale integration, on the other hand, is a manner of accounting for the nested boundaries of *cognizers* as opposed to the integrated boundaries of the nested non-cognizing constituents of cognitive systems. Like the multiscale integration view, reciprocal multiscale integration deploys the Markov blanket formalism ontologically. As it will be discussed in Sect. 2, the Free Energy Principle (FEP) (Friston 2010, 2013) is a framework that sees both life and cognition as coming about in virtue of minimizing the quantity variational free energy (henceforth free energy). The reason for using FEP to investigate and argue for a symbiotic mind is not only because the quantity “free energy” provides a measure of *cognition* across spatio-temporal scales (Ramstead et al. 2019a, b; Friston et al. 2015; Kirchhoff and Kiverstein 2018), but because FEP and its various corollaries suggest a plausible criterion for identifying biological cognizers across various spatio-temporal scales. In other words, the theoretical apparatus that falls out of FEP may be used to arrive at an account of the kind of enrichment of living processes that are required for those processes to qualify as cognitive processes ascribable to physiological individuals; thus, making good on the project of strong life-mind continuity.

As a case study in support of the notion of symbiotic minds I will look at the symbiosis of *Vibrio fischeri* bacteria and the bobtail squid. I will argue that the *Vibrio*-squid assemblage constitutes a functionally integrated cognitive whole, the Markov blanket of which constrains those of the *Vibrio* and the squid. In showing how the notion of a symbiotic mind can be supported by reciprocal multiscale integration within FEP, this paper contributes to the philosophy of cognitive science, demonstrating that our folk-psychological conception of what cognizers are require rethinking. Moreover, if the account of symbiotic cognizers presented is tenable, by bringing into relief and accounting for at least one possible transition in cognition, this paper contributes to a more complete understanding of the evolution of cognition.

This paper shall proceed as follows: In Sect. 2, after briefly reviewing FEP and exposing its criterion for being a cognizer, adaptive active inference, I will explicate the multiscale integrationist approach as a manner of individuating the boundaries of cognition. In Sect. 3 I will look at the symbiosis of *Vibrio fischeri* bacteria and the bobtail squid before turning to some philosophical implications of this symbiosis that are significant to thinking about symbiotic minds. In Sect. 4, building on the multiscale integration view, I shall provide an account of *reciprocal multiscale integration* and argue that some symbiotic associations and their symbiont constituents may be simultaneously construed as biological cognizers when reciprocally integrated. This paper concludes when some brief remarks addressing the problem of nested-cognizers.

2 The free energy principle—a short overview

The Free Energy Principle (FEP) (Friston 2010, 2013, 2019a, b), an ambitious unifying cognitive framework that may be applied to any biological system, starts from a particular view of life that is grounded in statistical physics. In nature, most systems self-organize to thermodynamic equilibrium. This means that all energy gradients for such systems have been consumed; something that entails death for living systems. Staying alive on the other hand requires that the states of a system behave in ways that counter the dissipative effects of random fluctuations. In other words, a biological system is one that self-organizes to a limited set of attracting states that is far from thermodynamic equilibrium (Nicolis and Prigogine 1977; Friston 2013). This set is referred to as a system's *nonequilibrium steady state density* (NESS) (Seifert 2012; Friston 2019a, b; Friston et al. 2020; Palacios et al. 2020). Importantly, the NESS density, towards which an organism's dynamics flow, corresponds to its phenotype (i.e., its regular patterns of behaviour, morphology, and physiology), to find itself in its characteristic or phenotypic states provides the organism is evidence that its behaviour is countering dispersive effects of random fluctuations. Hence, for a biological system to remain alive is for it to revisit the states defined by its NESS and to remain statistically separate from the environment in which it is embedded. Statistical separation is cast in terms of a conditional independence exhibited between the system's internal states and the external states of its environment. This brings us to the second primary feature of FEP: *the Markov blanket formalism*.³

The Markov blanket formalism as used in the context of FEP and biological systems describes a particular kind of statistical organization of a system relative to that which the system is not. Namely, it describes the statistical partitioning of internal states (r) and external states (φ) by sensory (s) and active states (a), these latter states being the states of the Markov blanket. The behaviour of the internal states may be predicted fully from knowing the states of the Markov blanket and as such the external states are rendered uninformative. Internal and external states are thus conditionally independent. This conditional independence falls out of the partitioning rule: internal states influence active states, which influence but are not influenced by external states; external states influence sensory states, which influence but are not influenced by internal states (Friston 2010). These dependencies/independencies do however allow for like states to influence like (see Fig. 1).

Active and sensory states create a statistical partition between internal states of system and external states of the environment, forming “a surface or boundary that defines the thing that exists (e.g., a cell membrane) (Friston 2019a, b, p. 176). Although this conditional independency implies that the internal states remain statistically distinct from external states it should not however be taken to imply that such a system

³ The notion of a Markov blanket (Pearl 1988) was originally introduced in the context of graph theory, where it describes a set of topological properties of a Bayesian nets. Specifically, a Markov blanket is the *smallest set of nodes* in a larger set of nodes, which if given information regarding their states would allow for the prediction of the unfolding states of a target node. This set of nodes includes *parents*, *children* and *co-parents*. The conditional independence between the target node and the nodes outside its Markov blanket mean that the latter are rendered uninformative or informationally redundant by the Markov blanket. The description of this set of nodes in terms of active, and sensory states is specific to FEP and its corollary active inference.




Fig. 1 Markov blanket partitioning rule: sensory states (s) influence but cannot be influenced by internal states (r). Active states (a) influence but cannot be influenced by external states (φ). This relation results in a conditional independence of internal and external states, where if one knows s and a at time, $t-1$, the knowledge of φ at time, $t-1$, provides no additional information about the internal state at t

is causally isolated from the environment. Biological systems are self-organizing thermodynamically open systems that operate far-from equilibrium. (Turner 1982). The manner which biological systems resist dispersive effects of fluctuations in their external milieu, and hence remain far from thermodynamic equilibrium, is by exchanging matter and energy with their environment (Friston 2013; Demirel 2014). These exchanges are mediated by a system's Markov blanket via circular causality: external states causally influence internal states by way of sensory state changes and internal states causally influence external states by way of active state changes. As such, although the system individuated by its Markov blanket is statistically independent from its environment, the formalism as used within FEP suggests that such systems remain tightly coupled to the dynamics of the environment (Friston 2013, 2019a, b). This brings us to a third essential feature of FEP: the notion of a *generative model*.

A generative model is a probabilistic model that describes how the evolution of sensory states of a Markov blanket could be caused by external states; it captures prior 'beliefs'—in the form of probability distributions—about unobserved external states, and a likelihood mapping external states to the evolution of sensory states. Such models are "implicit in the dynamics of internal states" (Palacios et al. 2020, p. 2). Importantly, a system's generative model may be cast in terms of its NESS. To see how this is the case requires understanding the notion of *dual information geometry* of self-organizing systems (Friston 2019a, b; Friston et al. 2020; Parr et al. 2019).

FEP describes the long-term dynamics of a self-organizing system in one of two mathematically equivalent ways that depend upon such a system's having a Markov blanket and, hence, its existing at nonequilibrium steady state. The evolution of system's states (i.e., its flow) may be described dually in terms of (1) its *phase space dynamics* or (2) its *statistical manifold dynamics*. With respect to its phase space description, every instantaneous state of the system is represented by a point in phase space (i.e., a space of all possible states that a system could occupy), every variable in the system corresponds to a dimension of the phase space, and the evolution of states over time corresponds to particular trajectory in phase space. Under this description, the probabilistic flow of a system's states over time is known as its *intrinsic* information geometry (Friston et al. 2020). A system's *extrinsic* information geometry

describes the probabilistic flow of a system's states over a statistical manifold, upon which each manifold coordinate corresponds to the sufficient statistics (i.e., the mean and the variance of a Gaussian) defining a probability density (i.e., a probabilistic 'belief' about external states). Adjacent points on a statistical manifold represent densities that are alike, and the distance between points (i.e., "information length") along a manifold path scores distances between probability densities encoded as points on that path respectively (Friston 2019a, b; Friston et al. 2020). In short, while intrinsic information geometry is a description of the probabilistic evolution of states system itself, extrinsic geometry is a description of 'belief' distributions over external states that internal states of a Markov-blanketed system encode.

FEP suggest that as long as a system possess a Markov blanket, rendering the system (i.e., its internal and active states) conditionally independent from its embedding environment, the flow in phase space towards the NESS density can be expressed as a probabilistic flow towards 'belief' distributions about external states. As a result, FEP allows for the remarkable observation that the NESS density towards which any self-organizing system tends to flow is that system's generative model. Given the presence of a system's Markov blanket, its intrinsic and extrinsic geometries coincide, meaning that a biological system's phenotype dually corresponds to its NESS density and to its generative model.⁴ How does a system's generative model contribute to the continued maintenance of its Markov blanket? This brings us to the fourth primary feature of FEP: *active inference*.

Remaining alive (i.e., statically separate from the environment) is a continuous feat that is accomplished despite the tendency for all systems to approach thermodynamic equilibrium. As such, living systems appear to resist the second law of thermodynamics. FEP proposes that living systems are able to do this in virtue of avoiding sensory states which are deleterious and actively bringing about those sensory states which allow them to maintain their structural and functional integrity. These latter kinds of phenotypically determined sensory states describe those which an organism expects itself to be in—given its generative model—and, as such, act as evidence for its existence (Kirchhoff and Kiverstein 2019b). This is just to say that an organism must thus pursue those actions which bring about evidence for itself over the long run (this includes exploration) and avoid those situations which could possibly result in its systemic dissipation. How does an organism accomplish this?

The free-energy principle that states that for any biological system to remain alive, it must minimize a quantity known as variational free energy. Variational free energy⁵ (henceforth, "free-energy") is an upper bound on the information theoretic surprise⁶ (i.e., the negative log probability of sensory states) and is a function of probabilistic

⁴ I would like to thank an anonymous reviewer for calling my attention to the importance of the phenotype's dual aspect.

⁵ Free energy is under some simplifying assumptions—the details of which I cannot treat in this paper—is mathematically equivalent to the quantity "prediction error" (Friston 2010) that is typically referred to in predictive processing. In what follows I shall use both free energy and prediction error to refer to the same quantity unless noted otherwise.

⁶ Surprise (i.e., negative Bayesian model evidence) is itself a computationally intractable quantity because it requires that biological systems have access to something that they generally do not have access to. Namely, it requires access to all the possible ways in which sensory states could have been caused (Friston 2010). Since free-energy is an upper bound on surprise (i.e., free-energy is always equal to or greater than surprise)

‘beliefs’ encoded by internal states of a Markov-blanketed system. Free-energy may be thought of as the difference between the sensory states an organism expects to observe given its generative model (or equivalently, its NESS density) and those that it actually observes at its sensorium. According to FEP, free-energy can be minimized by (1) inferentially optimizing the internal states of the Markov blanket so as to accommodate sensory states (i.e., perception), and (2) inferentially optimizing active states of the blanket (i.e., action) to sample sensory states that are expected given its phenotype. It is by engaging in this coupled process of *active inference*⁷ that an organism minimizes not only current free energy encountered but expected free energy (i.e., the free energy that would arise were a particular action policy selected and followed). It should be impressed that the inference of active inference does not presuppose personal level inferential processes; rather, it describes the dynamics which underwrite the optimization of internal states and prediction driven sensory feedback control.

Minimizing free energy over time is the same as maximizing *Bayesian model evidence* (i.e., self-evidence) (Friston 2013; Hohwy 2016). This is because when an organism’s generative model is well-tuned to its environment the discrepancy between the kinds of states that the organism expects to observe (i.e., states that confirm its own existence) and those that it does observe is at a minimum; its internal states may be thus understood as a model of the external states of its niche (Friston et al. 2015). It is thus by minimizing the free energy bound on surprise via active inference that biological systems “maximize the lower bound on the evidence for an implicit model of how their sensory samples were generated” (Friston 2013, p. 2). An interesting result of casting the NESS density as the generative model is that the gradient ascent of a system towards its NESS density may be described “as if” that system is minimizing variational free-energy and hence harvesting evidence for its generative model. To say that a system will behave “as if” it is performing a gradient descent on free-energy is a description that falls out of the fact that there is a dual information geometry (i.e., one can interpret a system’s gradient ascent on an attracting set through phase space in terms of its flow towards a generative model on a statistical manifold). In sum, to minimize free-energy is to maximize evidence for a generative model which is just maximising evidence that the flow of the system’s internal and blanket states is approaching its NESS density (i.e., flowing towards the system’s attracting states that counter fluctuation effects).

2.1 Enriching life with autonomy—adaptive active inference

There are two important points about FEP that are worth noting. Firstly, FEP is not only a theory of cognition but a theory of life; mind and life are both underwritten by the same fundamental principle of free energy minimization and come about via the kind of Markov blanket-preserving adaptive behaviour that is expressed by active inference. FEP as such is ideally placed as a framework within which to investigate the notion of

Footnote 6 continued

and because it is something that biological systems do have access to (i.e., it is a function of a generative model), it may be used as a proxy quantity for surprise. That is to say that by minimizing free-energy an organism can minimize surprise and maximize model evidence.

⁷ It should be impressed that “inference” as it is used within FEP is understood as a subpersonal process.

enrichment that is central to the strong life-mind continuity thesis. Moreover, because FEP does not specify the manner in which active inference is physically implemented, this framework may be used to theorize about cognition in a range of living systems of varying complexity (see Auletta 2013; Calvo and Friston 2017; Kirchhoff 2018; Sims 2019).

This brings us to the second point of importance. FEP falls short of making the claim that all Markov blankets draw a line around cognizers. It has been suggested that what determines whether a living Markov blanketed system is cognizer is whether or not it is the kind thing which engages in “adaptive active inference” rather than “mere active inference” (Kirchhoff et al. 2018). Adaptive active inference requires that a system *autonomously* engage in active inference, maximizing sensory evidence for its own existence (Kirchhoff and Kiverstein 2019a, b). Autonomy may be understood as the capacity for systemic self-evidencing through “actively monitor[ing] and react[ing] to perturbations that challenge homeostatic variables” (Kirchhoff et al. 2018, p. 5). Cells, tissues and organs, although they are self-organizing, free energy-minimizing systems, fail to rise to the level of bio-cognitive individuals because the manner in which they provide evidence for themselves (i.e. for their generative models) is highly dependent upon the machinery of and their interaction with other components of a global system (Kirchhoff and Kiverstein 2019a). As such, FEP (or more specifically its corollary adaptive active inference) boldly suggests that one manner in which life may be enriched so as to rise to the status of mind is tied to the increased degree of autonomy of the living Markov blanketed system.

The distinction between adaptive active inference and mere active inference offers a manner of addressing a *prima facie* worry that active inference (and FEP more generally) might be thought to face: ascribing Markov blankets to all living systems is intuitively unattractive if it brings with it the claim that all living systems engage in active inference. If something like a red blood cell actually engages in active inference, then despite the fact that inferences are understood as subpersonal descriptions of internal dynamics this result seems to make active inferential a trivial notion. This trivialization worry however may be avoided when recognizing that although every living thing is enshrouded by Markov blanket, this does not entail that every living thing engaged in the kind of adaptive active inference that is required to be a cognizer. In other words, one can understand the Markov blanket formalism as a genuine manner of individuating all living systems (it is not merely way of modelling them)⁸ while distinguishing the highly autonomous manner in which bio-cognitive systems maintain and preserve themselves and the highly interdependent manner that non-cognizing living systems maintain and preserve one another. The distinction between adaptive active inference and mere active inference demonstrates that active inference is far from a trivial notion—the devils in the detail.

To sum up: FEP (or more specifically its corollary adaptive active inference) boldly suggests that one manner in which life may be enriched so as to rise to the status of mind is tied to the increased degree of autonomy of the living Markov blanketed

⁸ That FEP holds Markov blankets to be ontologically robust may be supported by Friston when he writes that a Markov blanket “is not some statistical device by which we come to observe or model the world—it is a necessary attribute of a universe that can be carved into things (that are distinct from other things)” (Friston 2019a, b, p. 176).

system. With this understanding of FEP and adaptive active inference in hand, let us now briefly look at the multiscale integration view, which lays much of the theoretical ground upon which the notion of a symbiotic mind (and nested bio-cognizers) can be developed.

2.2 The multiscale integration

The multiscale integrationalist view (Ramstead et al. 2019a, b; Hesp et al. 2019) is a pluralist theory about cognitive boundaries; it uses the notion of nested Markov blankets to demonstrate that cognitive systems have a plurality of ontological boundaries, each relevant to the study of cognition.⁹ Cognition is seen as falling out of active inference that spans across and integrates various nested Markov blankets at different spatio-temporal scales (Ramstead et al. 2019a, b). To support this position, the multiscale integrationalist view offers a nuanced manner of spelling out just how the boundaries of cognitive systems are both nested and multiple via the process of multiscale integration. Take for example a human cognitive system which is composed of a brain; which is itself composed of cortical layers; which is composed of a collection of neurons; which are composed of individual neurons, etc. According to the multiscale integrationalist view each of the spatio-temporally nested components of this cognitive system may be ontologically picked out by deploying the Markov blanket formalism at different scales. Any one Markov blanket located at one spatio-temporal scale is simultaneously composed of other nested Markov blanket constituents at the scale below *and* is itself a constituent of a larger Markov blanket at the scale above it. The temporal dynamics between scales differ such that the smaller the scale is, the faster the dynamics are. The slower dynamics of macroscale systems *constrain* the faster dynamics of the microscale Markov blanketed systems, in effect acting as an order parameter¹⁰ (Kirchhoff et al. 2018; Ramstead et al. 2019a, b). The structure across scales is governed by the same conditional (in)dependencies described by the partitioning rule (Allen and Friston 2018; Clark 2017; Kirchhoff et al. 2018; Ramstead et al. 2019a, b; Friston 2019a, b). Since free energy is an additive quantity, the free energy which arises at the macroscale is the only free energy that there is for the ensemble and it may be quantified by summing the free energy associated with each constituent microscale Markov blanketed system. Minimizing free energy at the level of the composite system thus means minimizing free energy at the level of the constituent systems.

The process of multiscale integration may be generally understood in terms of inferences made on the part of individual Markovian partitioned microscale components (e.g. cells) of a macroscale system (e.g., a multicellular organism) that share the same generative model (see Friston et al. 2015; Kirchhoff et al. 2018; Palacios

⁹ If the multiscale integrationalist view is correct, then cognition can be a spread-out endeavour, the kind which often involves not only the brain but the dynamic and every-changing involvement of the body along with aspects of the environment—a view that is consonant with 4E theories of cognition (Clark 2008; Chemero 2009; Gallagher 2005).

¹⁰ An order parameter is a notion taken from synergetics (Haken 1985) and used dynamics systems theory. It denotes a measure of a global system's macroscale unstable (slow) dynamics that enslaves the fast dynamics of microscale component systems and results in a globally emergent pattern.

et al. 2020; Veissière et al. 2019). The idea of sharing a generative model may be construed in two manners given what has been said above (Sect. 2) regarding the dual information geometry of self-organizing Markov blanketed systems. When an organism's phenotype is cast in terms of its generative model, all component particles of the composite system will have the same expectations about the free-energy minimizing sensory states that should be encountered in addition to the kind of active states which are most likely to give rise to such sensory states. This means having the same of priors ('beliefs') acquired via evolution (Palacios et al. 2020). Integration, from the perspective of extrinsic information geometry, thus boils down to active inference on the part of those component particles, each instantiating a generative model with the same expectations as other integrating particles. The expected kinds of interaction (sensory feedback) are constrained by the partitioning rule that governs the organization of Markov blankets at any scale (e.g., If I infer that I am in an internal state of a larger blanketed system, then I expect to be influenced by all other states with the exception of external states). To minimize its free energy, and thus provide sensory evidence for itself as a constituent of the larger system, a microscale system allows itself to be enslaved to the dynamics of that macroscale system. The whole Markov blanket ensemble is integrated in virtue of the free-energy minimizing, self-evidencing dynamics occurring across/between each spatiotemporal scale (Ramstead et al. 2019a, b), resulting in the emergence of a composite system that possesses a *new* generative model at the superordinate level.

When taking the intrinsic perspective, where an organism's phenotype is cast as its NESS density, sharing a generative model may be seen as a case in which all of the subordinate level Markov blanketed particles of a superordinate system self-organize to a shared attracting set in phase space (i.e., a shared NESS density). From this perspective, the mechanics of the multiscale integration view may be understood as the idea that any living system may have multiple NESS densities, nested in successively larger spatio-temporal scales.¹¹ Because a Markov blanket in the case of living systems statistically maps onto a system as defined by its NESS density, multiscale integration can be understood as the idea that the coupled dynamics of multiple Markov blanket particles over time result in a new attracting state in their shared phase space. The emergence of this new attracting set is just the emergence of a system with a new NESS (i.e., a new generative model) at the superordinate level—a system that enslaves the dynamics of Markov blanketed systems at respective subordinate levels.

In providing a model of the integrated dynamics of Markov blankets across different spatio-temporal scales, the multiscale integration view offers an account of how the boundaries that compose *cognitive systems* are mutable, enveloping different component parts that contribute to the process of cognition as needed by the cognitive system. Multiscale integration of living systems that fail to possess the high degree of autonomy required of adaptive active inference (and hence fail to be cognizers) continue to engage in mere active inference at fast time scales (e.g., a mitochondrion in a cell).¹² On the other hand, multiscale integration of the nested systems that sub-

¹¹ I am grateful to an anonymous reviewer for clarifying this point.

¹² This is not to suggest that mitochondria do not exhibit some degree of autonomy and adaptivity. However, a mitochondrion's ability to seek out energy gradients required for its continued existence has become severely limited due to the evolution of the highly dependent relationship between it and its embedding cell.

tend the scale of a cognizing organism occurs in virtue of hierarchically “downward reaching” adaptive active inference; the behaviour of constituent nested Markov blanket systems are enslaved to the slower dynamics the autonomous cognizing organism; the subtending Markov blanketed systems become resources for the cognizing system to minimize the summated free energy across all nested blankets that compose it.¹³ This integration in such a case is *unidirectional*; it is the adaptive active inference at the slower scale of the cognizing system that drives the continued integration of the nested Markov blankets at scales below it and not vice versa. Similarly, multiscale integration between cognizers and non-living artefacts (e.g., spiders and their webs) may be construed in a similar fashion. Such integration is unidirectionally driven; it is the active inference dynamics of the biological cognizer that drives integration with the non-living Markov blanketed system and not vice versa.

This sets up the question which will occupy us for the remainder of this paper: what happens when the component parts of a biological system are themselves highly autonomous biological systems that engage in adaptive active inference? In other words, what happens in the case where integration involves physiological individuals that are cognizers in their own right? Whatever answering this question asks for, it most certainly requires something more than mere unidirectional integration. It requires a kind of reciprocal integration that I will argue undergirds the transition from a collection of cognizing physiological individuals to a symbiotic mind.

However, before presenting this account of reciprocal integration, it is important to understand just what kind of natural case in the world that the notion of a symbiotic mind might apply to. So, let us briefly look at one particular striking example of a symbiotic association: *V. fischeri* and bobtail squid symbiosis.

3 *Vibrio*-squid symbiosis: a case study

The Hawaiian bobtail squid *Eurprymna scolopes* is a nocturnal predator of the shallow reef of the Hawaiian archipelago that rests, buried in the sands, during the day and hunts at night (McFall-Ngai 2008). It is also prey to various marine predators. Whether or not it is preyed upon, however, is influenced by its association with bioluminescent bacteria *Vibrio fischeri*. The juvenile squid recruits *V. fischeri* into the epithelial-lined crypt spaces of its light organ by entraining seawater into itself with ciliated appendages (Nyholm and McFall-Ngai 2004). Subsequently the host promotes its colonization by *V. fischeri* and only *V. fischeri* via the production of mucus (i.e., bacteria food), the elimination of competing bacteria through *hemocyte* defences of its innate immune system (McAnulty and Nyholm 2017) and the eventual shedding of the ciliated appendages and the swelling of the crypt membranes preventing further entry into the light organ.¹⁴ Following gradients of chitin which they feed upon, *V. fischeri* migrate deeper into the crypts of the light organ, colonizing it, and causing

¹³ Inferring from large scale to small scale thus recapitulates the kind of adaptive active inference which would drive organogenesis (or morphogenesis) (see Friston et al. 2015).

¹⁴ The squid is able to control the amount of light emitted by its light organ by both contraction and expansion of the ink sac diverticula and by withholding oxygen to the *Vibrio* colonies which limits the amount of light the bacteria produce (McFall-Ngai 2008).

(after reaching a sufficiently high density) a biochemical reaction resulting in their emission of bioluminescent light. This occurs just in time for the bobtail's nightly hunt for prey. Because of the counterillumination against the light of the moon and stars that the bioluminescence allows, the squid casts no shadow and is thus camouflaged from potential predators on the seafloor below it whilst it hunts (McFall-Ngai and Montgomery 1990; Jones and Nishiguchi 2004). The bacteria also benefit from the bioluminescent reaction. The process consumes oxygen and induces hypoxic conditions in the squid's crypts that bring about the exocytosis of cytoplasmic substances out of the epithelial membrane (Visick and McFall-Ngai 2000). It is hypothesized that bioluminescence results in the production of nutrients for *V. fischeri* (Visick and McFall-Ngai 2000).

The particular population of *V. fischeri* that inhabit the light organ of the squid is, however, temporary. For the bobtail, cued by the light of dawn, vents up to 95% of the bacteria every morning (Nyholm and McFall-Ngai 2004). During the day, as the squid rests in the sand, *V. fischeri* that have not been expelled reproduce, repopulating the crypts so that by evening their density is high enough for the biochemical reaction that results in bioluminescence to occur again. After the initial colonization, this bacterial venting-repopulating cycle continues for the extent of the squid's life.

3.1 Philosophical implications of *Vibrio*-squid symbiosis

There are three philosophical implications that are exposed by this striking association that are significant to the account of symbiotic minds which follows. Firstly, the *Vibrio*-squid association is a clear example of a symbiotic physiological individual. What makes this association a symbiotic physiological individual opposed to a collection of physiological individuals? When deploying either an immunological criterion of physiological individuals (Pradeu 2011) and/or a metabolic criterion of physiological individuals (Godfrey-Smith 2013), given the high degree of immunological and/or metabolic functional integration between each symbiont, the *Vibrio*-squid association qualifies as a symbiotic physiological individual.¹⁵ But does the fact that the assemblage is temporary (i.e., that some of the *Vibrio* are expelled everyday) present a challenge to the status of the *Vibrio*-squid association as a symbiotic physiological individual? Does this daily expelling suggest that a proliferation of symbiotic individuals? This brings into relief a second crucial philosophical implication: physiological individuals are matters of degree; over time the same living system can be more or less of a (symbiotic) physiological individual (Bouchard 2018).

The last important philosophical implication may be thrown into relief when asking the critical question what is it that produces the bioluminescence—the squid, the bacte-

¹⁵ The immune criterion emphasizes (in)tolerance to elements with both exogenous and endogenous origins—relative to the system in question—as necessary for the kind of functional integration that a physiological individual requires. The metabolic criterion holds that a physiological individual must minimally possess its own core metabolic machinery (i.e., the machinery that result in processes of energy harvesting and energy breakdown). This machinery may also include external machinery (other environmental realizers) which regularly and closely influence those metabolic core processes. For considerations in favour of the metabolic criterion see Dupre and O'Malley (2009) and Godfrey-Smith (2013, 2014). For a detailed account of the physiological criterion see Pradeu (2011, 2016).

ria, or the temporary assemblage of these organisms? Fredric Bouchard (2018) argues that the bioluminescence is an emergent property of the *Vibrio*-squid assemblage. His reasoning is based upon the following facts about this case of bioluminescence. *V. fischeri* only begin to emit light when they sense a high concentration of autoinducers (i.e., a bacterial metabolic product) in the surrounding medium of the crypt (Miller and Bassler 2001). This process of quorum sensing induces the transcription of the *lux* gene that results in the bacterial bioluminescent light emission (Nyholm and McFall-Ngai 2004). Although *V. fischeri* populations in the seawaters could grow to such density as to quorum sense and generate light independently of the squid, *V. fischeri* do not actually do this in the wild (Bouchard 2018). The light organ provides just the right conditions (i.e., shelter and food) for exponential population growth (McFall-Ngai 2008) required of quorum sensing. The bobtail, similarly, cannot generate light itself without being colonized by *V. fischeri*. If this is correct, then it is not either the squid or *V. fischeri* that glow alone but rather “what glows is a temporary assemblage of species interacting in the right way” (Bouchard 2018, p. 190).

With this case study and these philosophical implications in mind, let us now return to the notion of reciprocal multiscale integration and to ultimately answering the question of how such a symbiotic physiological individual could be a symbiotic mind.

4 Unidirectional versus reciprocal multiscale integration

To understand what is meant by reciprocal integration, I will now contrast it with unidirectional integration. To make this contrast apparent I will deploy the terminology of users, U_i , and resources, R_i . Users are those living systems, the internal states of which inferentially generate subpersonal predictions that allow them to use other Markov blanketed systems that are external to them. A necessary characteristic of users is that they have the capacity to engage in adaptive active inference and use Markov blanketed systems that are external to them in order to minimize *long-term* free energy. Where being external is relative to the Markov blankets that play an essential role to the user’s autonomous self-evidencing. Resources, on the other hand, are those external Markov blanketed systems that given their structure and behaviour provide a means for users to minimize their long-term free energy but are themselves not required for autonomous self-evidencing. Importantly, resources can, but need not, have the capacity to engage in adaptive active inference; they can but need not be enriched enough so as to qualify as cognizers according to FEP. As such, a Markov blanketed system can be both a user and a resource.

In unidirectional multiscale integration, the internal states of U_1 generate predictions about the role that R_1 plays (e.g., sensory, internal or active state) in a superordinate level Markov blanket in which it also plays a role. In order to minimize short-term free energy that results from such predictions, U_1 acts in ways to bring about the kind of feedback that would result were it the case that both it and R_1 played those inferred roles. On the other hand, R_1 does not (and perhaps cannot) generate predictions about U_1 . In contrast, reciprocal multiscale integration describes the process whereby a user, U_1 , is simultaneously a resource, R_2 , for a distinct user, U_2 , that is itself a cognitive resource R_1 for U_1 (see Fig. 2).




Fig. 2 Unidirectional multiscale integration vs reciprocal multiscale integration: Plate **a** illustrates small/fast scale to large/slow scale unidirectional integration. It is unidirectional in terms of being driven by an inferential procedure, the source of which is the user and not the resource. Here R_1 is used by (represented by the arrow) a prediction generating cognizer U_1 but is not itself a cognizer (i.e., although R_1 has a Markov blanket, it does not engage in adaptive active inference). Plate **b** illustrates reciprocal integration from small/fast scale to large/slow scale and vice versa. It occurs when users, U_1 and U_2 , are each bio-cognizers and inferentially generate predictions about the sensory states that would be encountered were they playing different roles in the same larger Markov blanketed system. Resources, R_1 and R_2 , are used by U_2 and U_1 respectively as the users engage in adaptive active inference, adapting their generative models to one another. Because U_1 is R_2 and U_2 is R_1 , the integration between the users is reciprocal

U_1 generates inferences about the role that R_2 plays in a larger Markov blanket that it also plays a role in, and U_2 does the same with respect to R_1 . U_1 acts as a resource contributing to the adaptive active inference of U_2 and vice versa such that from the increasingly coupled state space of U_1 and U_2 a new shared ($U_1/R_2 + U_2/R_1$) NESS density emerges at the superordinate scale, enslaving both U_1 and U_2 . A new generative model emerges that not only encodes priors about the ‘preferred’ states of U_1 and U_2 , but also encodes priors concerning the kinds of environmental conditions in which the superordinate level unit thrives in addition to the kinds of actions which bring these conditions about. The result of reciprocal integration is an extended symbiotic phenotype that authors the conditions for its own existence; by behaving in ways that optimize action, sampling the world that it expects, the integrated symbiotic unit tends to provide evidence for its own generative model.¹⁶

To be clear, each constituent bio-cognizer continues to engage in adaptive active inference, acting in ways that ensure it remains statistically separate from its environment. That two individuals are enslaved by a state attractor at the superordinate scale does not imply that they must renounce their autonomy to the degree of forfeiting their status as bio-cognizers. Symbiosis is an effective means of cooperatively sustaining the Markov blankets of both U_1/R_2 and U_2/R_1 separately and as a unit. Importantly,

¹⁶ This kind of reciprocal multiscale integration between two (or more) cognizing systems may be interestingly interpreted as a form of niche construction (Constant et al. 2018; Veissière et al. 2019), whereby U_1 and U_2 , co-evolve and come to play the role of exo-genetically inherited (i.e., non-genetically specified) resources for each other. Inheritance, spanning long enough timescales, becomes consolidated in the emergence of a new ($U_1/R_2 + U_2/R_1$) NESS density at the superordinate level. I would like to thank an anonymous reviewer for bringing this insightful interpretation of reciprocal multiscale integration to my attention.

the integration of U_1/R_2 and U_2/R_1 into a unit determines a new shared Markov blanket where external states are defined by everything that is not $(U_1/R_2 + U_2/R_1)$. It is the unit which infers the states external to its blanket, differentially interacting with them in order to persist as a unit. Crucially, the manner in which $(U_1/R_2 + U_2/R_1)$ interacts with its environment cannot be reduced to the differential operations at the level U_1/R_2 and U_2/R_1 . It is when they are reciprocally integrated that an *emergent action capacity* arises. This is a capacity possessed by a unit to behave in a manner which outstrips the combined individual behavioural capacities of each user. I would like to suggest that such an emergent action capacity is evidence that adaptive active inference is occurring at the level of the symbiotic physiological individual.¹⁷ In other words, when the reciprocal integration of symbiotic partners that have been evolutionarily (i.e., adaptively) coerced into following a long-term reciprocal free energy minimizing strategy do result in an emergent action capacity, this capacity is possessed by a symbiotic cognizer.

To be sure, an emergent action capacity is a property of the reciprocally integrated $(U_1/R_2 + U_2/R_1)$ unit that constrains and affects the behaviour of the symbiotic components U_1/R_2 and U_2/R_1 . As such, how symbiotic cognizers and their action capacities arise is best characterized as a form of *ontological emergence*. Ontological emergence refers to two joint ideas: (1) some properties possessed by a composite macro-level system cannot be reduced to the structural (i.e., intrinsic) properties of its component parts and their governing micro-dynamics and (2) because of these new irreducible properties, a macro-level system has ontological status (e.g., is an entity) in its own right. (Santos 2015). Such emergent properties at the macro-level are more than the properties of their component parts in the sense that macro-level properties have causal powers that their component parts fail to possess, enabling a macro-level system to both facilitate and constrain the very micro-level parts and dynamics that give rise to it—what is referred to as ‘downward causation’. Taking into account however that both new environmental influences are specific to each hierarchical level (Gilbert and Sarkar 2000) and that causation involved in ontological emergence is diachronic (i.e., macro-levels of a system affect micro-levels and vice versa at different sequential timesteps) (see O’Connor and Wong 2005) emergence may be more accurately—and less mysteriously—analysed in terms of circular causation.^{18,19} In fact, this circular causation should be expected in FEP given an understanding of the kind coupled dynamic exchanges between self-organizing open systems and their external milieu that are required for such systems to resist dispersive effects. That is, because system $(U_1/R_2 + U_2/R_1)$ is causally coupled to its environment in a way that neither U_1/R_2 or U_2/R_1 are individually, $(U_1/R_2 + U_2/R_1)$ has causal powers that neither U_1/R_2 or

¹⁷ The *emergence* of an action capacity is distinct from the *realization* of an emergent action capacity. For example, although I may not be able to do 20 push-ups now after a long workout at the gym, I certainly still have the capacity to do so when I am rested.

¹⁸ I am indebted to an anonymous reviewer for pushing me to clarify more about the notion of emergence being used in this paper.

¹⁹ The contemporary notions of emergence are both various and nuanced (see for example Winning and Bechtel 2019). Although providing a more detailed analysis of the kind of emergence involved in reciprocal multiscale integration is an interesting and valuable philosophical endeavour, for reasons of limited space, the more general description of ontological emergence will have to do for the purposes of this paper.

U_2/R_1 possess individually. The emergent system has action capacities that, via facilitating free-energy minimizing exchanges with its environment at the superordinate scale, constrain the behaviour of components U_1/R_2 and U_2/R_1 at subordinate scales.

Let us now concretize this somewhat abstract description of how reciprocal integration might occur at the level of the symbiotic unit and return to the case study of the *Vibrio*-bobtail squid association.²⁰

4.1 *Vibrio*-squid reciprocal integration

The squid-*Vibrio* reciprocal integration might unfold as follows: the internal states of *V. fischeri*, U_1 and those of the bobtail system, U_2 , come to generate the same inferences about the sensory effects that should be observed were both of the to play the role of sensory and active states in a superordinate blanketed system. For example, both U_1 and U_2 might expect sensory feedback in the form of a reduced ambient oxygen gradient in the bobtail's light organ (i.e., an increased hypoxic state). The squid system, in order to reduce the free energy that arises as result of its inferences, acts to bring about that very pattern of protracted feedback that it expects. To do this, the squid harvests *V. fischeri*, R_2 , and only *V. fischeri* and thereafter provides conditions for their thriving from which these protracted sensory patterns may arise. When all goes well, reducing free-energy associated with the expected hypoxia allows the squid system to further influence the evolving states of its light organ; these evolving light organ states reflect the evolution of two convergent flows on a shared squid-*Vibrio* state space.

Similarly, in order to reduce the free-energy that arises from the *Vibrio* system's inferences about expected sensory states, the bacteria migrate further into deep crypts of the light organ, shed their flagella (reducing a degree of their own autonomy), and bring about a decreasing ambient oxygen gradient over time as both their density and quorum sensing increases. In this reciprocal manner, the generative models of the squid and *V. fischeri* adapt to one another, each user acting as a free-energy minimizing resource for the other. Crucially, as a result of the continued and increasing squid-*Vibrio* coupled activity, a new attractor emerges at a superordinate scale on their shared state space. This emergent NESS density corresponds to a new generative model that is distinct from the squid and *Vibrio* models that gave rise to it. Thus, by reciprocally behaving in ways to bring about the kinds of feedback that would accompany their inferred roles in a superordinate Markov blanketed system, U_1 and U_2 integrate across spatio-temporal scales, harvesting evidence for an emergent phenotype ($U_1/R_2 + U_2/R_1$). (See Fig. 3). Via its influence upon the environment, this new symbiotic bio-cognitive individual both constrains and facilitates the continued behaviour of its component users and realizers.

The notion of environmental influence may be understood in terms of the effects of an *emergent action capacity* that is a property of the symbiotic bio-cognizer and

²⁰ The description of symbiotic cognitive integration that follows is admittedly simplified in that I shall treat the *V. fischeri* colony as a single individual. *Vibrio*-squid symbiosis, although only involving two species, involves more than two organisms (e.g., a colony of *Vibrio* may reach 10^9 inside the light organ of a bobtail). Moreover, given that scale refers to spatio-temporal scale, on intuitive grounds I will stipulate that the *V. fischeri* system and the bobtail system live at different scales; the former, being smaller and faster, is nested in the latter.




Fig. 3 Reciprocal multiscale integration: the squid system, U_1/R_2 , infers (indicated by the wavy arrows) that *V. fischeri* system, U_2/R_1 , plays the role of sensory states (indicated by yellow nodes) and that it plays the role of active states (indicated by orange nodes) in a superordinate level Markov blanket system ($U_1/R_2 + U_2/R_1$). U_1/R_2 , becomes further coupled to the dynamics of U_2/R_1 , actively bringing about the kind of sensory feedback that would accompany each of the subordinate level system's playing its respective inferred role in the larger blanket. Similarly, the *V. fischeri* system, U_2/R_1 , integrates with the squid system by inferring that U_1/R_2 , plays the role of the active states and that it plays the role of sensory states in ($U_1/R_2 + U_2/R_1$). U_2/R_1 , then acts in ways to bring about the kinds of sensory feedback that would be observed were the two systems to play these inferred roles. The bacteria (user) system does this by engaging in tightly coupled behaviour with the squid(resource) system. This reciprocal integration over time results in the emergence of a new NESS density, the dynamics of which enslave those of U_1/R_2 and U_2/R_1 . The resulting global emergent action capacity possessed by the symbiotic unit ($U_1/R_2 + U_2/R_1$) allows for a downwards causal influence upon those very systems which give rise to it. Modified from Kirchhoff et al. (2018). (Color figure online)

not of its component parts. More precisely, as U_1 and U_2 become further integrated, flowing towards the new squid-*Vibrio* NESS density, the new system acquires a new action capacity, controlled bioluminescence, that only reaching maturation, is actualized nightly at the *peak of integration*.

This reveals a crucial point: Markov blanket integration (reciprocal and unidirectional) as located in FEP is a matter of degree. That this is the case is implied by the fact that Markov blankets can be more or less integrated over time and that integration does not happen instantaneously but is something that can increase (and decrease). The spider system that rebuilds its web after being forced to relocate to a different location is likely to integrate with its newly constructed web to a high degree only after interacting with it; this a process where internal states of the spider system calibrate to the sensory feedback that arises as external states (e.g., wind, prey, its own actions) affect the web (i.e., the resource). Integration between the bobtail and *V. fischeri* systems is a rolling diel process (i.e., a cyclical process that occurs over the course of a 24-h period). For

subsequent to the initial colonization, reciprocal integration peaks nightly with the emergence of controlled bioluminescence. This high degree of integration lasts until the symbiotic bio-cognitive individual that is evidenced by the emergent action capacity expels 95% of the *Vibrio* back into the water when cued by the first light of dawn. At this point the integration drops in degree to its lowest value of the diel cycle and the process of reciprocal integration begins anew from this lower degree until the unit, with the emergent bioluminescence, actively generates evidence for its own existence again.

Does this drop in the degree of integration that accompanies the diel cycle result in a problem of proliferating bio-cognitive individuals; short-lived individuals that spring into existence nightly and are the cause of their own dissipation at each dawn venting? I would like to argue that this problem may be avoided when it is recognized that being a bio-cognitive individual is a matter of degree.²¹ Even when the squid-*Vibrio* unit is not integrated enough to realize its capacity for controlled bioluminescence (by its glowing and using it), this does not mean that the squid-*Vibrio* unit is not integrated enough to *possess* the capacity. The long-lived bio-cognitive squid-*Vibrio* unit, through a continuous process of adaptive active inference retains this capacity and when enough free energy is minimized in the reciprocal integration process, the symbiotic unit may deploy the capacity. This suggests that when the capacity for controlled bioluminescence is not realized, the symbiotic unit is less of a bio-cognitive individual rather than losing its status full stop (See Fig. 4).

This kind of ebb and flow of the degree of cognitive individuality is familiar in our own case. For although when sleeping we fail to use our capacities to fully engage in the kind of self-evidencing active inference, we remain bio-cognizers. What I am suggesting is that in such cases we may be bio-cognizers to a lesser degree. Moreover, the daily expelling of the *Vibrio* may be seen itself as an instance of self-evidencing of the integrated symbiotic cognizer! For if the *Vibrio* were not vented, the existence of the unit would be at risk. Venting is an active way of fostering a condition in which *Vibrio* quorum sensing will continue effectively in the future by keeping *Vibrio* in an active growth phase (Visick and McFall-Ngai 2000). As such venting on the part of the symbiotic unit may be construed in terms of actively reducing *expected* free energy. Moreover, venting prevents the symbiotic association from becoming harmful to the squid (too high of a *Vibrio* density over a long period may result in *Vibrio* bringing about pathogenic responses).

If this picture of reciprocal integration of the squid and *Vibrio* is correct, then the *Vibrio*-squid physiological individual via reciprocal multiscale integration becomes a model of its niche; a model that is different from that of its individual symbionts. Both the existence of the symbiotic cognizer and the pattern of controlled bioluminescent masking that occurs over the course of night is evidence for the specific range of environmental conditions that have evolutionarily come to have *Vibrio*-squid-association-relative value. Expressing this point deploying the ecological notion of

²¹ This solution is consistent with the idea that a being a physiological individual is a matter of degree (Queller and Strassman 2009; Godfrey-Smith 2014; Bouchard 2016); the degree of physiological individuality is proportional to the degree of functional integration. As such, then although after the venting of the *Vibrio* the squid-*Vibrio* association is less of physiological individual it is still functionally integrated (metabolically and/or immunologically) enough to retain its status as a physiological individual.




Fig. 4 The diel cycle of *Vibrio*-squid integration and bio-cognitive individuality

affordances (Gibson 1979), while the moonlit waters of the Hawaiian reef may afford safety for the integrated *Vibrio*-squid association, it fails to offer an affordance of this sort to either of the lone symbionts.

To sum up: I have argued that a symbiotic unit (e.g., *Vibrio*-squid association) that is constituted of cognizing symbionts at different scales may itself be a biological cognizing system when the symbiont partners that compose it engage in a high degree of reciprocal integration, a form of Markov blanket integration that falls out FEP. If what I have argued is correct, it is possible that in some restricted cases symbionts can maintain their status as bio-cognizers whilst simultaneously giving rise to a new NESS density that corresponds to a symbiotic cognizing individual *in its own right*.

5 Conclusion

I began this paper by flagging a particular concern that arises when considering the notion of a symbiotic mind. The idea that cognizers may have other cognizers as nested constituent parts is an idea that seems to run counter to our folk-psychological intuitions about what it is to be a cognizer. This concern I noted reiterates a similar worry for the physiological account of biological individuality: the hierarchically nested ontology of physiological individuals in physiological individuals flouts out folk-biological intuition that an organism cannot be a constituent of another organism.²² This worry for the physiological account can be defanged however by pointing

²² See S. Chauvier's (2017) "the formal indivisibility principle" for an expression of this concern which is also called "the exclusion principle" (Godfrey-Smith 2014).

out that physiological individuals are distinct from Darwinian individuals (i.e., units of reproduction) and that it is the nesting of the latter that is intuitively problematic.²³ As such, folk-biological intuitions cannot be used when adjudicating the claim that physiological individuals can be nested in other physiological individuals in symbiotic associations (Godfrey-Smith 2014). With this in mind, I would like to suggest that the worry that cognizers cannot be nested in other cognizers may stem from a similar failure on behalf of folk-psychology to distinguish physiological individuals from Darwinian individuals, overlooking the former. For if our folk-psychology implicitly models the notion of cognizers upon that of Darwinian individuals—a much more scientifically old and ingrained conception of biological individuals—then it would seem that the notion of nested cognizers would remain as unintuitive as the notion of nested Darwinian individuals. However, by reconceptualizing our folk-psychological notion of “cognizer” to include physiological individuals the notion of nested biological cognizers is no more unintuitive as that of nested physiological individuals. Moreover, given that the conception of cognizers as physiological individuals is already implicit in FEP, this framework and its corollary active inference (mere and adaptive), offer an ideal programme for bringing our folk-psychological intuitions in line with the way nature has carved its own joints over evolutionary timescales.

Acknowledgements I would like to thank Dave Ward, Julian Kiverstein, Andy Clark, Mark Miller, Beren Millidge, Nina Poth, and Ben Little for their many helpful comments and fruitful conversations. Many thanks to Thomas Pradeu and all of the speakers and participants at the IDEM 2019 Symbiosis summer school for providing a lively and challenging platform for discussing many of the ideas that gave rise to this paper. I am also grateful to two anonymous reviewers for their generosity and extremely valuable suggestions.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Allen, M., & Friston, K. (2016). From cognitivism to autopoiesis: Towards a computational framework for the embodied mind. *Synthese*. <https://doi.org/10.1007/s11229-016-1288-5>.
- Auletta, G. (2013). Information and metabolism in bacterial chemotaxis. *Entropy*, 15(1), 311–326. <https://doi.org/10.3390/e15010311>.
- Bouchard, F. (2018). Symbiosis, transient biological individuality, and evolutionary processes. In D. J. Nicholson & J. Dupré (Eds.), *Everything flows: Towards a processual philosophy of biology*. Oxford: Oxford University Press.
- Calvo, P., & Friston, K. (2017). Predicting green: Really radical (plant) predictive processing. *Journal of the Royal Society Interface*, 14, 20170096.

²³ Aphid-Buchnera bacteria symbiosis presents an interesting case in which the symbiotic association may indeed qualify as both a Darwinian individual and a physiological individual (for more on this see Godfrey-Smith 2013). This suggests that our intuitions should not have the final say in biological ontologising.

- Chauvier, S. (2017). Individuality and aggregativity. *Philosophy Theory Practice Biology*. <https://doi.org/10.3998/ptb.6959004.0009.011>.
- Chemero, A. (2009). *Radical embodied cognitive science*. MA: MIT Press.
- Clark, A. (2008). *Supersizing the Mind: Embodiment, action and cognitive extension*. Oxford: Oxford University Press.
- Clark, A. (2017). How to knit your own Markov Blanket: resisting the second law with metamorphic minds. In T. Metzinger & W. Wiese (Eds.), *Philosophy and predictive processing: 3. Frankfurt am Main: MIND Group*. <https://doi.org/10.15502/9783958573031>.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 50, 7–19.
- Constant, A., Ramstead, M. J., Veissière, S. P., Campbell, J., & Friston, K. J. (2018). A variational approach to niche construction. *Journal of the Royal Society Interface*, 15, 20170685.
- Demirel, Y. (2014). *Nonequilibrium thermodynamics: Transport and rate processes in physical* (3rd ed.). Amsterdam: Elsevier B.V. Chemical and Biological Systems. <https://doi.org/10.1016/C2012-0-00459-0>.
- Dupre, J., & O'Malley, M. A. (2009). Varieties of living things: Life at the intersection of lineage and metabolism. *Philosophy Theory Biology*, 1(201306), 1–25. <https://doi.org/10.3998/ptb.6959004.0001.003>.
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11, 127–138.
- Friston, K. (2013). Life as we know it. *Journal of the Royal Society, Interface*, 10, 20130475. <https://doi.org/10.1098/rsif.2013.0475>.
- Friston, K. J. (2019). Beyond the desert landscape. In M. Colobmo, E. Irvine, & M. Stapleton (Eds.), *Andy Clark and his critics*. Oxford: Oxford University Press.
- Friston, K. J. (2019). A free energy principle for a particular physics. *arXiv*1906.10184. pp. 1–148.
- Friston, K., Levin, M., Sengupta, B., & Pezzulo, G. (2015). Knowing one's place: A free-energy approach to pattern regulation. *Journal of the Royal Society, Interface*, 12, 20141383. <https://doi.org/10.1098/rsif.2014.1383>.
- Friston, K. J., Wiese, W., & Hobson, J. A. (2020). Sentience and the origins of consciousness: From Cartesian duality to Markovian monism. *Entropy*, 22(5), 516.
- Gallagher, S. (2005). *How the body shapes the mind*. Oxford: Oxford University Press.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston, MA: Houghton Mifflin.
- Gilbert, S. F., & Sarkar, S. (2000). Embracing complexity: Organicism for the 21st century. *Developmental Dynamics*, 219(1), 1–9.
- Godfrey-Smith, P. (1996). *Complexity and the function of mind in nature*. Cambridge: Cambridge University Press.
- Godfrey-Smith, P. (2013). Darwinian individuals. In F. Bouchard & P. Huneman (Eds.), *From groups to individuals: Evolution and emerging individuality*. Cambridge: MIT Press.
- Godfrey-Smith, P. (2014). *Philosophy of biology*. Princeton and Oxford: Princeton University Press.
- Haken, H. (1985). Synergetics: An interdisciplinary approach to self-organization. *Geoform*, 16(2), 205–211.
- Hesp, C., Ramstead M., Constant A., Badcock P., Kirchoff M., & Friston K. (2019). A multi-scale view of the emergent complexity of life: A free-energy proposal. In G. Georgiev, J. Smart, C. Flores Martinez, & M. Price (Eds.), *Evolution, development and complexity*. Springer Proceedings in Complexity.
- Hohwy, J. (2016). The self-evidencing brain. *Noûs*, 50(2), 259–285. <https://doi.org/10.1111/nous.12062>.
- Jones, B. W., & Nishiguchi, M. K. (2004). Counterillumination in the Hawaiian bobtail squid, *Euprymna scolopes* Berry (Mollusca: Cephalopoda). *Marine Biology*, 144, 1151–1155. <https://doi.org/10.1007/s00227-003-1285-3>.
- Kirchoff, M. D. (2018). Autopoiesis, free energy, and the life-mind continuity thesis. *Synthese*. <https://doi.org/10.1007/s11229-016-1100-6>.
- Kirchoff, M. D., Parr, T., Placios, E., Friston, K. J., & Kiverstein, J. (2018). The Markov blanks of life: Autonomy, active inference and the free energy principle. *Journal of the Royal Society, Interface*, 15, 20170792. <https://doi.org/10.1098/rsif.2017.0792>.
- Kirchoff, M. D., & Kiverstein, J. (2019a). *Extended consciousness and predictive processing: A third wave view*. Abingdon: Routledge.
- Kirchoff, M. D., & Kiverstein, J. (2019). How to determine the boundaries of the mind: A Markov blanket. *Synthese*. <https://doi.org/10.1007/s11229-019-02370-y>.

- Maynard-Smith, J., & Szathmáry, E. (1995). *The major transitions in evolution*. New York: Oxford University Press.
- McAnulty, S., & Nyholm, S. (2017). The Role of Hemocytes in the Hawaiian Bobtail Squid, *Euprymna scolopes*: A model organism for studying beneficial host-microbe interactions. *Front. Microbiol.*, 7, 2013. <https://doi.org/10.3389/fmicb.2016.02013>.
- McFall-Ngai, M., & Montgomery, M. (1990). The anatomy and morphology of the adult bacterial light organ of *Euprymna scolopes* Berry (Cephalopoda: Sepioidae). *Biological Bulletin*, 179(3), 332–339.
- McFall-Ngai, M. (2008). Hawaiian bobtail squid. *Current Biology*, 18(22), R1043–R1044.
- Miller, M., & Bassler, B. (2001). Quorum sensing in bacteria. *Annual Review of Microbiology*, 55, 165–199.
- Nicolis, G., & Prigogine, I. (1977). *Self-organization in nonequilibrium systems*. New York, NY: Wiley.
- Nyholm, S. V., & McFall-Ngai, M. J. (2004). The winnowing: Establishing the squid–*Vibrio* symbiosis. *Nature Review Microbiology*, 2, 632–642.
- O'Connor, T., & Wong, H. Y. (2005). The metaphysics of emergence. *Noûs*, 39(4), 658–678.
- Parr, T., Da Costa, L., & Friston, K. J. (2019). Markov blankets, information geometry and stochastic thermodynamics. *Philosophy Transactions Royal Society A*, 378, 20190159. <https://doi.org/10.1098/rsta.2019.0159>.
- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. San Francisco, CA: Morgan Kaufmann Publishers.
- Palacios, E. R., Razi, A., Parr, T., Kirchhoff, M., & Friston, K. (2020). On Markov blankets and hierarchical self-organisation. *Journal of Theoretical Biology*, 486, 110089.
- Pradeu, T. (2011). A mixed self: The role of symbiosis in development. *Biology Theory*, 6, 80–88. <https://doi.org/10.1007/s13752-011-0011-5>.
- Pradeu, T. (2016). Organisms or biological individuals? Combining physiological and evolutionary individuality. *Biology and Philosophy*, 31, 797. <https://doi.org/10.1007/s10539-016-9551-1>.
- Queller, D. C., & Strassman, J. E. (2009). Beyond society: the evolution of organismality. *Philosophy Transactions Royal Society B*, 2009(364), 3143–3155. <https://doi.org/10.1098/rstb.2009.0095>.
- Ramstead, M. J. D., Kirchhoff, M. D., Constant, A., & Friston, K. (2019). Multiscale integration: beyond internalism and externalism. *Synthese*. <https://doi.org/10.1007/s11229-019-02115-x>.
- Ramstead, M. J. D., Kirchhoff, M. D., & Friston, K. (2019). A tale of two densities: Active inference is enactive inference. *Adaptive Behavior*. <https://doi.org/10.1177/1059712319862774>.
- Santos, G. C. (2015). Ontological emergence: How is that possible? Towards a new relational ontology. *Foundations of Science*, 20, 429–446.
- Seifert, U. (2012). Stochastic thermodynamics, fluctuation theorems and molecular machines. *Report Program Physics Physics Society*, 75, 126001.
- Sims, M. (2019). Minimal-perception: Responding to the challenges of perceptual constancy and veridicality with plants. *Philosophical Psychology*. <https://doi.org/10.1080/09515089.2019.1646898>.
- Turner, J. (1982). Self-organization in nonequilibrium chemistry and in biology. In W. C. Schieve & P. M. Allen (Eds.), *Self-organization and dissipative structures, applications in the physical and social sciences*. Austin: University of Texas Press.
- Veissière, S. P., Constant, A., Ramstead, M. J., Friston, K. J., & Kirmayer, L. J. (2019). Thinking through other minds: A variational approach to cognition and culture. *Behavioral and Brain Sciences*, 43, e90. <https://doi.org/10.1017/S0140525X19001213>.
- Visick, K. L., & McFall-Ngai, M. J. (2000). An exclusive contract: Specificity in the *Vibrio fischeri*-*Euprymna scolopes* partnership. *Journal of Bacteriology*, 182(7), 1779–1787.
- Winning, J., & Bechtel, W. (2019). Being emergence VS. Pattern emergence: Complexity, control and goal-directedness in biological systems. In S. Gibb, R. F. Hendry, & T. Lancaster (Eds.), *The Routledge handbook of emergence*. London New York: Routledge.