

## SPECIAL ISSUE ARTICLE OPEN ACCESS

The Relevance of a Philosophical Toolkit to Advance Neuroscience

# The Possibility Space Concept in Neuroscience: Possibilities, Constraints, and Explanations

Lauren N. Ross<sup>1</sup>  | Viktor Jirsa<sup>2</sup> | Anthony R. McIntosh<sup>3</sup><sup>1</sup>Logic and Philosophy of Science, University of California, Irvine, California, USA | <sup>2</sup>Institut de Neurosciences des Systèmes, Aix-Marseille Université, Marseille, France | <sup>3</sup>Biomedical Physiology and Kinesiology, Simon Fraser University, Burnaby, Canada**Correspondence:** Lauren N. Ross ([rossl@uci.edu](mailto:rossl@uci.edu))**Received:** 6 July 2024 | **Revised:** 2 December 2024 | **Accepted:** 16 February 2025**Associate Editor:** Markus Kunze**Funding:** This work was supported by the National Science Foundation (NSF) Career, 1945647; John Templeton Foundation, 63021; and NSERC Discovery Grant, RGPIN-2018-04457.

## ABSTRACT

Although the brain is often characterized as a complex system, theoretical and philosophical frameworks often struggle to capture this. For example, mainstream mechanistic accounts model neural systems as fixed and static in ways that fail to capture their dynamic nature and large set of possible behaviors. In this paper, we provide a framework for capturing a common type of complex system in neuroscience, which involves two main aspects: (i) constraints on the system and (ii) the system's possibility space of available outcomes. Our analysis merges neuroscience examples with recent work in the philosophy of science to suggest that the possibility space concept involves two essential types of constraints, which we call hard and soft constraints. Our analysis focuses on a domain-general notion of possibility space that is present in manifold frameworks and representations, phase space diagrams in dynamical systems theory, and paradigmatic cases, such as Waddington's epigenetic landscape model. After building the framework with such cases, we apply it to three main examples in neuroscience: adaptability, resilience, and phenomenology. We explore how this framework supports a philosophical toolkit for neuroscience and how it helps advance recent work in the philosophy of science on constraints, scientific explanations, and impossibility explanations. We show how fruitful connections between neuroscience and philosophy can support conceptual clarity, theoretical advances, and the identification of similar systems across different domains in neuroscience.

## 1 | Introduction


Neuroscience research is engaged in the study of complex systems, exemplified by its focus on the brain. In this research, there is interest in identifying various states and capacities of the brain, capturing what produces these states, and understanding how they change over time. Although there is significant diversity across types of research with this focus, there are recurring and foundational concepts that are shared across this work and are important

for understanding the behaviors of complex systems. One of these concepts is the notion of a system's possibility space, which refers to some total set of available states given some system of interest. Our analysis focuses on a domain-general notion of possibility space (Figure 1) that is present in manifold frameworks and representations, phase space diagrams in dynamical systems theory, and classic cases such as Waddington's epigenetic landscape model. This paper examines the possibility space concept in complex systems, focusing on the brain and current methodology in

**Abbreviations:** FCD, functional connectivity dynamics; FC, functional connectivity; TVB, The Virtual Brain.

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial](https://creativecommons.org/licenses/by-nc/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2025 The Author(s). *European Journal of Neuroscience* published by Federation of European Neuroscience Societies and John Wiley & Sons Ltd.



**FIGURE 1** | Manifold of possibilities: (A) Representation of Waddington's epigenetic landscape, in which a ball rolling along a landscape is analogized to an undifferentiated cell taking on different states. The ball's trajectory is guided by a landscape of valleys and hills, which constrain its movement through space. Similarly, an undifferentiated cell can travel through different states, which are also guided and constrained by various factors. The landscape captures possibilities or possible states of the ball, states that are more (and less) likely, and also constraints that guide these sequences of states. (B) A network diagram captures similar features, as it involves a set of nodes that can be differentially engaged in a network. There is a possible space dictated by connections among these nodes that determine which network states can be occupied and in what order.

neuroscience. Our goal is to provide a philosophical toolkit that supports conceptual clarity around discussions of possibilities, impossibilities, and relevant constraints as they present in complex systems. We explore how the possibility space framework captures different explanations of complex systems, different constraints on complex systems, and unique features these systems have that are more challenging to capture in alternative frameworks, such as mechanistic perspectives.

Our interest in this topic grew from the realization that many descriptions and explanations of brain capacities rely on characterizations of the brain's possibility space and that the possibility space concept relates to other important notions in neuroscience and philosophy. For example, in the neuroscience literature, recent accounts of *hidden repertoires* rely on characterizations of unrealized states, as they capture inaccessible pockets of possibility space in the context of brain dynamics (McIntosh and Jirsa 2019). Such proposals are typically couched in the dynamical systems framework leveraging multistability and structured flows on manifolds. Additionally, notions of adaptability, resilience, and phenomenology often refer to brain states that are impossible or possible for a system, and more or less accessible to the system as a whole. Further motivation for this analysis is the

emergence of related work in the philosophy of science that examines different types of constraints, their role in scientific explanation, and their relation to the possibility space concept.<sup>1</sup> A growing literature distinguishes types of constraints in science, the role of these constraints in providing scientific explanations, and their unique ability to provide "impossibility explanations," which are a rare form of scientific explanation (Anderson 2016; Lange 2016; Ross 2023). Other work examines explanatory constraints that are top-down, structural, and causal, with a focus on capturing how they differ from traditional mechanistic models of complex systems and accounts of explanation (Bolt et al. 2018; Raja and Anderson 2021; Silberstein and Chemero 2013).

In this work, we first review mechanistic accounts of explanation in neuroscience and highlight some of their challenges in accommodating complex systems. We then examine two distinct types of constraints—hard and soft constraints—which capture different limitations on complex systems. We use these constraint types to develop a possibility space framework that is applied to three main neuroscience examples—adaptability, resilience, and phenomenology. Although this framework is applied to neuroscience cases in particular, we also demonstrate its ability to generalize to other examples, including gene regulatory networks, immune

cell development, and Waddington's classic epigenetic landscape example.

## 2 | Moving Beyond the Mechanism Tradition

In neuroscience and philosophy, many research programs have characterized neural and brain systems through scientific explanations, revealing the inner workings of complex systems through mechanisms. (Craver 2007; Bechtel and Richardson 2010; Bechtel and Levy 2013). Largely inspired by advances in molecular and cellular work, the mechanism concept in neuroscience originated with causal descriptions of systems in terms of their lower level parts. Modern neuroscience now admits explanations at higher scales, as seen in appeals to network, dynamical, and other macroscale explanatory models. Acceptance of such higher level explanations is also seen in claims that higher level explanatory regularities can be multiply realized by distinct neural details in ways that capture more general, universal explanatory models. Despite this acceptance of explanation across scales, reductive language and assumptions remain. This is seen in frequent explanatory efforts that aim to get “under the hood” and suggestions that deeper explanations require lower level neural details. Even when explanatory factors operate at higher scales, it can be challenging to resist some mention of the importance of lower scale information or future promises that such lower levels may someday give a deeper understanding.

An interesting feature of the mechanism paradigm—and its focus on capturing explanatorily relevant causal structure—is that the structure it captures is typically quite fixed and static. Mainstream philosophical accounts describe mechanisms as discrete systems with lower level causal parts, where these parts interact to produce some higher level outcome of the system (Craver 2007; Bechtel and Richardson 2010). Such a picture is strictly hierarchical in the sense that relevant causes are always at lower scales than the effects they produce. There is often interest in how these causes are organized and most accounts allow for more abstract mechanisms, where causal parts need not be spatially proximal but where they remain at lower scales than the outcomes they explain. Although such a framework may seem unobjectionable and fairly accurate of most causal systems and explanations, its limitations have been increasingly exposed in recent work. One of these limitations is that this model views mechanism as highly static and fixed—in most of these cases there is a clear causal story in terms of particular fixed parts, which interact to produce a single outcome (Dupré 2013). In particular, most accounts view these systems as mechanisms “for” some particular outcome, where this outcome defines the mechanism, its causal boundaries, and sets up conditions. This is supported by mechanists’ claims that explanations often start with the specification of some explanatory target of interest, after which one “drills down” and “decomposes” the system into its relevant parts. This framework “fixes” the behavior that a system produces, making it difficult to capture the plethora of diverse outcomes that many systems can be driven to exhibit.

Part of what this framework misses is a natural way to capture systems that are dynamic, with many different available states, causal trajectories, and endpoints. Such dynamic aspects involve some set of possible states for a system, the actual states it manifests, and

other factors that enable or constrain the behavior of the systems. Where standard causal mechanism models work best in capturing repeating, regular causal parts, interactions, and outcomes, such models are limited in capturing various features of various complex and dynamic systems (Dupré 2013). These complex systems are less amenable to approaches that first fix the final behavior of the system while then searching for lower level causes that are “under the hood” and repeatedly produce system-level behavior. We do not deny the utility of the mechanistic framework for some systems but rather emphasize its limitations in capturing all types of systems in neuroscience. This point has been made in the context of the pathway and mechanism concepts in neuroscience and biology. Representations of roadmaps of pathways in the life sciences often capture many different routes that a system can take, which is not something that mechanism models are able to capture (Ross 2021a).

Our analysis in this paper expands on traditional mechanistic frameworks by providing an account that captures systems with constraints, dynamical behaviors, and wide landscapes of possible outcomes and states. As various philosophers have indicated, mainstream new mechanistic accounts outline a “fixed” notion of causal mechanisms, which struggles to capture causal systems in neuroscience that are more dynamical, changing, and open to large sets of possibilities (Dupré 2013; Silberstein and Chemero 2013). Other limitations of this mechanism framework include its strict hierarchical character, which is challenged in capturing systems with causes that are higher scale, top-down, and that have a constraining influence. This hierarchical notion accommodates causal systems with lower level parts that produce higher level outcomes—a feature also captured with the decomposition and localization strategies emphasized in new mechanist literature (Bechtel and Richardson 2010). For these reasons, recent philosophical projects aim to expand on new mechanist accounts and articulate frameworks that capture other types of causal systems in science. Various projects in philosophy of science have contributed to this project, by capturing distinct forms of explanation, understanding, and modeling consistent with dynamical systems and network approaches. This is seen in work on dynamical explanations, explanatory constraints and enabling conditions, and network explanations (Anderson 2016; Lange 2016; Silberstein and Chemero 2013; Ross 2023).

The emphasis of our framework here is on the interplay between two features of neural systems: *constraints* and the system's *possibility space*. We provide an account of how these features show up in many complex systems in neuroscience and how they help capture the diverse range of states that these systems can exhibit, the trajectories that depict how these systems change, sets of states that are available but uncommon, and also states that are strictly off-limits for the system. This outlines a more nuanced model of neural systems and a framework that is intended to generalize to complex systems in other domains as well.

## 3 | Hard Constraints

One implication of the fact that systems have properties that can present in a range of available states is that they also have states that are impossible as they cannot be realized physically or are off-limits. Many of these impossible states are not just

uncommon or rarely realized—they are strictly impossible and unavailable to the system. For example, it is impossible for a household oven to heat itself to 2000°C, it is impossible for the human auditory system to detect sound below 10 Hz or above 25 kHz, and it is impossible for human neurons to transmit signals at 500 m/s (or faster). It might be possible for other systems to exhibit these states, but for the system in question, these states are unavailable. This limit on a system's states, particularly the border between states that are impossible versus possible for a system is captured by what we call *hard constraints*. In this sense, *hard constraints* capture limits on the system's behaviors that are hard in that they specify what is strictly unavailable. In the context of brain behaviors, clear examples of hard constraints are differences in anatomical connectivity (or macroscale wiring), limits when it comes to signal speed and processing, and limits on the environmental conditions under which brains can perform (such as temperature and energy resources).<sup>2</sup>

An ordinary life example that helps capture hard constraints and impossible states—and one that is often used to capture complex systems—is the game of chess (Holland 2014). At any point in a game of chess, there are available moves that the player can make, but there are also strict limits to these moves.<sup>3</sup> The rules of chess are similar to what we call *hard constraints*—these rules constrain local moves and how the game can evolve (Holland 2014; McIntosh and Jirsa 2019). Examples of hard constraints from scientific contexts include various laws of nature and mathematical properties that constrain living systems and, in so doing, explain outcomes that they can and cannot produce (Lange 2016; Ross 2023; Silberstein and Chemero 2013).<sup>4,5</sup> The role of hard constraints in these cases always depends on specified properties of interest, including the system, environment, timescale, and traits in question. This notion of hard constraints is similar to Maynard Smith et al.'s (1985) discussion of the “inescapability” and “bindingness” of constraints—essentially, how fixed, strict, and unbreakable constraints are for a system.<sup>6</sup> Although the rules of chess and the law of gravity are “inescapable” for the systems above, the following section considers constraints that are less “binding” and that operate as suggesting, guiding, and channeling—we call these soft constraints, as discussed in the next section.

Constraints have received significant attention in recent philosophical work. One reason is that constraints appear to provide important types of scientific explanations, but they differ from standard explanatory factors (Raja and Anderson 2021; Ross 2023; Silberstein and Chemero 2013). Many accounts of scientific explanation focus on explanations of why a system exhibits one possible outcome over another. This is seen in explanatory why-questions that ask why an organism shows a particular eye color (from some available range), a particular height (from some set of possibilities), or some diseased state versus a nondiseased state. In these cases, the explanatory target consists of a range of possible outcomes, whereas the explanatory factors (often causes or mechanisms) explain why one possible outcome was produced instead of another. In contrast to this picture, one explanatory role of constraints—particularly the hard constraints discussed in this section—is to explain why some outcomes are strictly unavailable or off-limits to a system. As we will see soon, this is related to claims that constraints provide “impossibility

explanations,” which differ from explanations of why one possible outcome (over other possibilities) is produced (Lange 2016).

To examine this further, consider a framework suggested by Ross (2024) for capturing constraints and their role in scientific explanation.<sup>7</sup> In this account, constraints, in general, are factors that are (i) often viewed as more external to a system or parts of interest, (ii) relatively fixed compared to other explanatory factors, and (iii) structure or guide the system as opposed to triggering or determining its behavior (Ross 2024). Notice how these three features are present in our chess example—the rules of chess are external to the moves of the game, these rules are fixed relative to players choices about which next move to make, and these rules limit or guide plays in the game, without determining which exact play will occur next.<sup>8</sup>

However, although many generic constraints have these three features, *hard* constraints have the additional feature of (iv) specifying states that are strictly off-limits and impossible for the system. To see that some constraints lack (iv), consider that some constraints (such as a lack of resources, types of peer pressure, and imposed speeding laws) can make various behaviors less likely, without strictly outlawing them. A unique feature of hard constraints is that they capture and explain why some behaviors are strictly off-limits. This helps capture how hard constraints provide explanations and understanding of complex systems. First, (3.1) hard constraints can be used to provide “impossibility explanations,” which explain why particular outcomes are impossible or off-limits for a system (Lange 2016; Ross 2021b). In order to see this, consider the seven bridges of Königsberg example, in which a graphical or network model captures “topological constraints” that explain. In this case, the question of interest whether it was possible to walk a single path across each of the seven bridges of Königsberg exactly and only once. After much interest in the question, Euler provided proof demonstrating that bridge systems with such a path (what we now call an Eulerian path) need to meet two criteria. When the bridge system is represented graphically, all nodes need to have at least one connection, and there need to be either zero or two nodes of odd degree (Euler 1956; Ross 2021b). As the Königsberg bridges failed to meet these criteria, such a path was strictly impossible for this system. In this case, the bridge topology is a constraint explaining why something is impossible for a system. These cases are viewed as “explanations by constraint,” which have the feature of explaining impossibilities and being stronger or “inevitable to a stronger extent” than other scientific explanations (Lange 2016). Similarly, the hard constraints on human audition and nerve signaling physiology explain why some outcomes are impossible for these systems. This explanatory pattern differs from standard causal or mechanistic explanations, as these explain why some actual state is produced among a set of possible outcomes (as opposed to explaining why some states are strictly off-limits).

In addition to providing impossibility explanations, hard constraints can also explain (3.2) changes in a system over time and (3.3) comparative explanations across distinct systems. With respect to (3.2), some systems acquire changes in their hard constraints, which explain changes in what is possible (and impossible) over time. This is seen in various insults to the brain, such as strokes and trauma-induced lesions, where functional



areas are irreversibly damaged in ways that prevent forms of cognitive functioning. In these cases, reduced cognitive functioning is explained by changes in hard constraints, which further limit what is possible for the system. This is also seen in developmental examples, where differentiation of neural crest/stem cells is genetically determined, forming the basis for sensory, motor, and autonomic functions. This genetic program is a hard constraint that explains the narrowing of cell type functionality through development. The retina connects with visual thalamic and midbrain regions, whereas the cochlea connects with auditory brainstem regions. The strengths of these connections within pathways vary with development and experience, but the initial connections are hard constraints for the system architecture.

Third and finally, (3.3) hard constraints also provide comparative explanations that capture variations across distinct systems. Suppose we want to explain why dogs can hear higher pitched sounds (at 30–47kHz), whereas humans cannot. This is explained by the fact that dogs and humans have different hard constraints on hearing possibilities. In other words, the difference in hard constraints across distinct systems explains the difference in impossible outcomes across them. As Raja and Anderson state, “... if we have two similar systems but only one of them is constrained, the unconstrained one will exhibit more degrees of freedom in its behavior precisely because of the lack of constraints” (Raja and Anderson 2021). This comparative explanation also has implications for phenomenology, wherein hard constraints bind the experience of the world. This may seem trivial, but consider the difference in perception of walking in a forest for a dog versus a human. The olfactory and auditory systems dominate the perception of the forest for the dog. This percept is impossible for humans to imagine because it lies outside the boundaries of hard constraints.

This section highlights the main features of hard constraints, the ways they provide explanations, and how they capture important features of complex systems. Although hard constraints capture the border between impossible and possible outcomes, other constraints operate within a system's possibility space. Next, we discuss soft constraints and possibility space considerations and how such a framework captures further features of complex systems and the behaviors they produce.

## 4 | Soft Constraints and Possibility Space

Soft constraints operate as predisposing causes that guide, encourage, and influence a system to exhibit some states over others.<sup>9</sup> For example, consider the different fates of a stem cell, the different routes blood flows through dense vasculature, and whether an individual moves into a state of experiencing a mood disorder. In the stem cell case, examples of soft constraints include local concentrations of molecules (coenzymes, growth factors, etc.) and environmental conditions (pH, temperature, etc.) that selectively guide the developmental trajectory of a stem cell, encouraging it to mature along one developmental pathway over another. Similarly, when blood flows through dense vasculature, differences in molecular, chemical, and physical factors (tissue damage factors, pH, vessel size, etc.) encourage more flow to areas recovering from exercise, injury, or other types of

stress. Finally, whether an individual experiences a mood disorder (such as depression) is influenced by their genetic profile and environmental factors (such as stressful life events), which can each encourage or discourage such conditions (without necessarily strongly determining them). In all these cases, soft constraints made various states of the system more or less likely without strongly determining which states will present (and without dictating which states are impossible).

Other examples of soft constraints are various cognitive limitations that restrict the capacity for optimal human decision-making in various contexts. This relates to Herbert Simon's notion of “bounded rationality,” in which humans engage in decision-making with “universal cognitive limitations” that prevent the selection of perfectly optimized outcomes or decisions (Bendor 2001; Simon 1996). For example, when selecting moves in chess, our cognitive abilities are a “binding constraint” that can explain why many of our decisions are suboptimal (otherwise, both players would choose the most optimal moves every time, and the game would be rote). Interestingly, because of these limitations, humans rely on various heuristics and strategies to enhance performance in these cases. Other examples of soft constraints are various cognitive limitations and “information processing constraints”—these channel, guide, and limit behavioral outcomes without strongly dictating what is impossible or off-limits, as seen in hard constraints (Bendor 2001).

*Possibility space* representations—complete with information about all soft constraints and trajectories driven by the non-linear flow in state space—capture a more nuanced, dynamic picture of complex systems compared to models of static, fixed systems with a limited number of variables or causal factors.<sup>10</sup> Furthermore, the possibility space framework captures states that are available but unrealized, together with different biases on which states are more likely and how such likelihoods change as early decisions are made and as the system evolves. These possibility space models capture an additional layer of complexity in cases where the possibility space itself—in particular, its soft constraints and trajectories—changes over time.<sup>11</sup> These changing possibility space models capture systems that adapt, evolve, and vary in ways that produce highly diverse and seemingly disrupted landscapes with residual features. Another advantage of this possibility space framework is that it captures systems that contain a massive number of possibilities but where the system only implements some of these potential outcomes for purposes of adaptation, evolvability, or response to insult. An example of this from immunology is clonal selection, in which a huge diversity of pre-existing lymphocytes exists in the human body (Rajewsky 1996). This diversity covers the range of possible antigen types, such that the presence of a given antigen binds an existing lymphocyte to create specific antibodies. The possibility space framework can help capture the extensive range of potential antibody responses that the immune system can have and, on the other hand, the particular response that the system manifests.

Other examples of such soft constraints are attractors, detractors, and basins in dynamical system theory frameworks. Dynamical systems evolve in time and are described by a set of variables that unambiguously define the system's state. In classical mechanics, a pendulum has the state variables'


position and velocity; in thermodynamics, a gas is fully described by the totality of the kinetic variables (position and velocity in 3D) of its molecules; and in neuroscience, a neuron's electric activity is captured by its membrane potential and the kinetic variables associated with ion channel opening probabilities (Hodgkin Huxley equations). The more components a system has, the more state variables it has, and its dimension increases. As time evolves, the system traces out a trajectory in the state space spanned by its state variables. In general, complex systems will be nonlinear, which refers to the relationship between the state variables and their change of rate. Linear approximations work well only for certain conditions (e.g., the linear pendulum for small angles). Outside of these ranges, complex systems—from artifacts to living organisms—can express their nonlinear behavior, of which the coexistence and multistability of behaviors is one fundamental property. Importantly, they display this repertoire of states for the same system configuration, that is the setting of the system's parameters (such as the pendulum's length in the previous example). As the system's configuration changes, the system's repertoire of states may also change. For example, a dimmer light bulb can be in many different states, including “off” or in the states of emitting light at “low,” “medium,” or “high” levels, depending on the system parameter's setting (the dimmer's position). Human auditory systems can detect the presence and absence of sound, in addition to varying pitches within a given range. A neuron's firing frequency and response speed can take on value within a range of states specific to the neuron. The possibility space is the set of available states for a system, defined by the nonlinear properties of the system itself (establishing its dynamic repertoire) and its configuration (its parameter settings). It depends on many factors, including the particular property or behavior in question (such as light emission, sound detection, and firing frequency), how this property is defined or characterized, and features in the environment, background, or context of the system of interest, to name a few. Strictly speaking, the states do not have to be static in terms of “having no change of rate” in terms of time evolution, thus representing characteristic behaviors captured by different domains in state space. The system may occupy such domains for a finite time and then reorganize and evolve somewhere else. For instance, an oscillatory behavior may be shown until the system fatigues or runs out of energy, then changes its behavior (or state) and transits to rest. We use the term “state” in this more general sense of behavior.

Well-known and paradigmatic examples of the possibility space concept include Waddington's epigenetic landscape; phase space diagrams in dynamical systems theory; manifold frameworks and representations; network or pathway “maps” of developmental, metabolic, and other outcomes; and decision tree figures in various contexts (Huang 2012; Izhikevich 2007; Waddington 1957). These cases rely on various representations that capture a system's possible outcomes, how these possibilities have distinct features that can vary over time, and how a system flows through this space in exhibiting some outcomes over others. For example, in Waddington's epigenetic landscape, the changing state of a developing cell is represented by a ball rolling through a varied landscape, which contains valleys, hills, and varying slopes (Waddington 1957). Different locations of the ball's trajectory along the landscape

represent different states of the cell, whereas landscape features constrain, enable, and channel which states are realized. Similar concepts are captured in phase space diagrams with attractors, detractors, and basins, which reveal possible states that are more or less likely than others. In other cases, network models, pathway diagrams, and decision trees capture how a system's states change in sequence, represented by a trajectory through space (Ross 2021a). These illustrations reflect how some routes through space are more likely than others and how early “moves” change which downstream outcomes are available (or not), similar to the notion of path dependence. Additionally, these possibility space landscapes can capture available states that the system can manifest but also states that are difficult to reach, less likely to occur, and less accessible than others.

The possibility space examples typically contain three key elements: (4.1) a total possibility space, (4.2) constraints that capture guiding influences within the space, and (4.3) trajectories through the space that capture changes in an entity's state over time. The guiding constraints within the possibility space differ from the hard constraints discussed in Section 3. We suggest that these factors are well understood as soft constraints because they guide what states of the system are likely to manifest, instead of specifying what outcomes are completely off-limits (i.e., hard constraints). In Waddington's epigenetic landscape, examples of soft constraints are the valleys, hills, and slopes that channel the ball's trajectory as it flows through the possibility space. Although these constraints have some influence on the future state of the system, they do not capture impossible states, and they rarely determine which specific possible outcomes occur.

Before we discuss neuroscience cases, it will help to illustrate the applicability of these concepts with an expansion of the Waddington Epigenetic landscape metaphor (Figure 2). In a recent review paper, Huang (2012) connects the epigenetic landscape metaphor with dynamical systems theory to show how the evolution of gene regulatory networks can be linked to the epigenetic landscape. A link between manifolds and flows occurs when gene expression patterns are couched in network expressions. Here, the manifold contains all potential gene expressions (4.1). The network architecture per se is a hard constraint, whereas the specific network state determines the gene expression, that is, the soft constraints. As these networks develop, either across evolutionary time or in the organism's development, the broader manifold (the epigenetic landscape) is formed, containing all attractor basins consistent with gene expression capacity (4.2). The manifold also contains attractors consistent with the architecture—possible—but unoccupied, unused, and sometimes hidden. In Huang's consideration, these unused attractors are as follows: “inevitable mathematical by-products of the network that has evolved to produce the set of ‘useful’ attractors” (Huang 2012, supp text, 3). In other words, as an adaptive system, the gene regulatory networks can form attractors that are potential configurations. Huang (2012) makes two more assertions about the unused attractors. First, on the negative side, these are ordinarily difficult to access, but if occupied, they could lead to tumorigenesis when cells, facilitated by somatic mutations that make them accessible, are trapped in such attractors



**FIGURE 2** | Revised epigenetic landscape: Gene regulatory networks can be linked to the epigenetic landscape. Modified from Huang (2012).

and remain immature. Second, the unused attractors could also convey the potential for evolution, where adaptive pressures distort the manifold to make the attractor accessible and open sets of gene expression programs for exploration.

The relevance of this discourse is the link between hard and soft constraints that define impossible and possible spaces from the perspective of dynamical systems. Insofar as the brain is a dynamic system, the updated epigenetic landscape analogy would be equally applicable to the brain. The possibility framework makes the connection between attractors in the manifold that support resilience and adaptability. This can be realized within the lifetime of the brain as the pressure for adaptation will change with development. How the developmental pressure is addressed will impart resilience to the system.

Moreover, as the attractor space develops, unused attractors will also form, making the connection to the hidden repertoires noted in the introduction. As we elaborate below, the hidden repertoires have the same dual consequence: one that leads to pathology and another that supports a new adaptation. Because a brain's manifold(s) contain the entire space for possible function, we can link that architecture to phenomenology, where our personal experience of a situation is a consequence of what happened and what *did not happen* but was possible. Our current experience exists in the context of what is possible.

The dynamical systems framework, structured flows on manifolds, provides a mathematical expression that complements the updated epigenetic landscape. Pillai and Jirsa (2017) introduced this framework, which was then linked to system configuration changes in neuroscience displaying resilience and adaptability (Jirsa 2020). The structured flows on manifolds framework emphasizes spatiotemporal processes encapsulated in relatively low-dimensional manifolds. The manifolds set the boundary conditions for the available functional configuration (4.1). The functional configuration, or flow, that is enacted depends on the

specific demands of the situation. (4.2) Thus, the actual network used exists atop a space of possible networks that could have been enacted within a given manifold architecture (4.3).

The proposal of hidden repertoires came from the structured flows on manifolds framework (McIntosh and Jirsa 2019). An example, which is consistent with the negative impact of unused attractors, is epilepsy, where the expression of seizure activity is not commonly seen but is a capacity present in the system. The hidden repertoire was further elaborated to suggest that another pathological state, status epilepticus, was a potential that could be realized with the appropriate shift in a slow variable that governs seizure expression. Key to this is that the attractor space for epilepsy exists in ostensibly healthy manifolds but is not expressed. The second and positive aspect of hidden repertoires is less definitive but can be inferred from manifold dynamics that support healthy behavior but with different attractor dynamics. We explore this possibility in the next section. These considerations show how the possibility space concept emerges in various modeling contexts—from epigenetic landscape models to dynamical system frameworks and suggest a way to understand multiple shared components and features across these contexts. We examine these frameworks in more detail in the next section, focusing on concrete neuroscience cases.

## 5 | Neuroscience Examples

The possibility space framework—complete with soft constraints and trajectories—has several advantages in capturing complex brain behaviors, dispositions, and capacities. One clear advantage, articulated earlier, is the ability of this framework to capture a broader set of possibilities for a system, regardless of whether they have been observed by researchers or realized by the system. A second main advantage, discussed in more detail subsequently, is the ability of this framework to capture changes in the possibility space over time and, in

particular, unique pockets of possibility space that result from changes to the brain. Changes to the brain due to adaptation or evolution can create residual, hard-to-access zones and lead to remaining areas that are unused by enabling frequently occupied zones. Many of these less-used pockets of possibility space are “hidden repertoires,” which are a type of hidden state space and harder to access, and “invisible to other approaches,” many of which focus on actual, realized outcomes and more static representations of brain capacity (McIntosh and Jirsa 2019). A third main advantage is the ability of this framework to capture different explanatory patterns, such as explaining impossibilities, biases in states that are more (or less) likely to manifest, and ways in which systems evolve and change over time.

*A word about plasticity.* Many features that support the possibility space in the brain can be linked to the notion of “plasticity.” Given that the brain is a complex *adaptive* system, plasticity can be considered as the means to adapt and change. This may not be the same for all complex adaptive systems, such as socioeconomic systems, where adaptation comes through other mechanisms. However, as the behavior of complex adaptive systems can be characterized as reflecting hard and soft constraints, it is essential to note that the adaptive aspects of possibility space may not all come from plasticity in the case of the brain. To be clear, possibility space can be modified, changed, and driven by plasticity, but this is not the only factor influencing adaptation and change in neural systems and the brain (e.g., neurogenesis and pruning, neuromodulation, and epigenetics).

In modern network neuroscience, particularly neuroimaging, there is a focus on defining structural and functional networks and relating their architectures to cognition and behavior. Some contemporary work emphasizes particular spatiotemporal scales, such as that accessible by resting-state fMRI, which provides a partial picture of what is supported by brain networks as characterized by measures of statistical dependency between regions or functional connectivity (FC). For example, FC measured over long-time scales on the order of minute (static FC) may characterize a capacity that enables the differentiation between persons but misses the potential faster time scale dynamics (seconds or shorter) that represent what each person actually does (dynamic FC). Dynamic FC may capture the flow of mental processes. Static FC, by the calculation, will smooth over that flow and provide only a glimpse of the landscape that was traversed. For instance, Khambhati et al. (2018) proposed that linear models best describe resting state dynamics after evaluating their predictive power using second-order derived metrics (same as static FC), biasing the analyses toward such smoothened behaviors. Indeed, static FC could indicate hard constraints that define the boundaries of the functional configuration in a brain. Dynamic FC, on the other hand, exploits the dynamic repertoire within the soft constraints of the system. They may not relate to impossibility per se, which is better aligned to structural and biophysical constraints. Hard constraints in the functional space could define the collection of available network configurations but not how those configurations are related across time. Dynamic FC would be more sensitive to time dependence and give an estimate of space possibilities. The key distinction here is that the trajectories a person engages to flow through this space

embody the flow of cognition. This flow can be quantified, and the paths taken are related to cognitive status.

We discuss applications of this possibility space framework (and the hidden repertoire concept) in the context of three examples: *adaptability*, *resilience*, and *phenomenology*. Adaptability and resilience are related concepts, but we will consider them separately.

## 5.1 | Adaptability

Adaptability refers to the ability of a system to evolve and develop over time. In the context of the brain, this involves the capacity for various functional configurations to form, dissolve, and reform, supported by the brain's inherent plasticity. Throughout development, the brain can adopt multiple configurations, enabling it to adapt to changing circumstances.

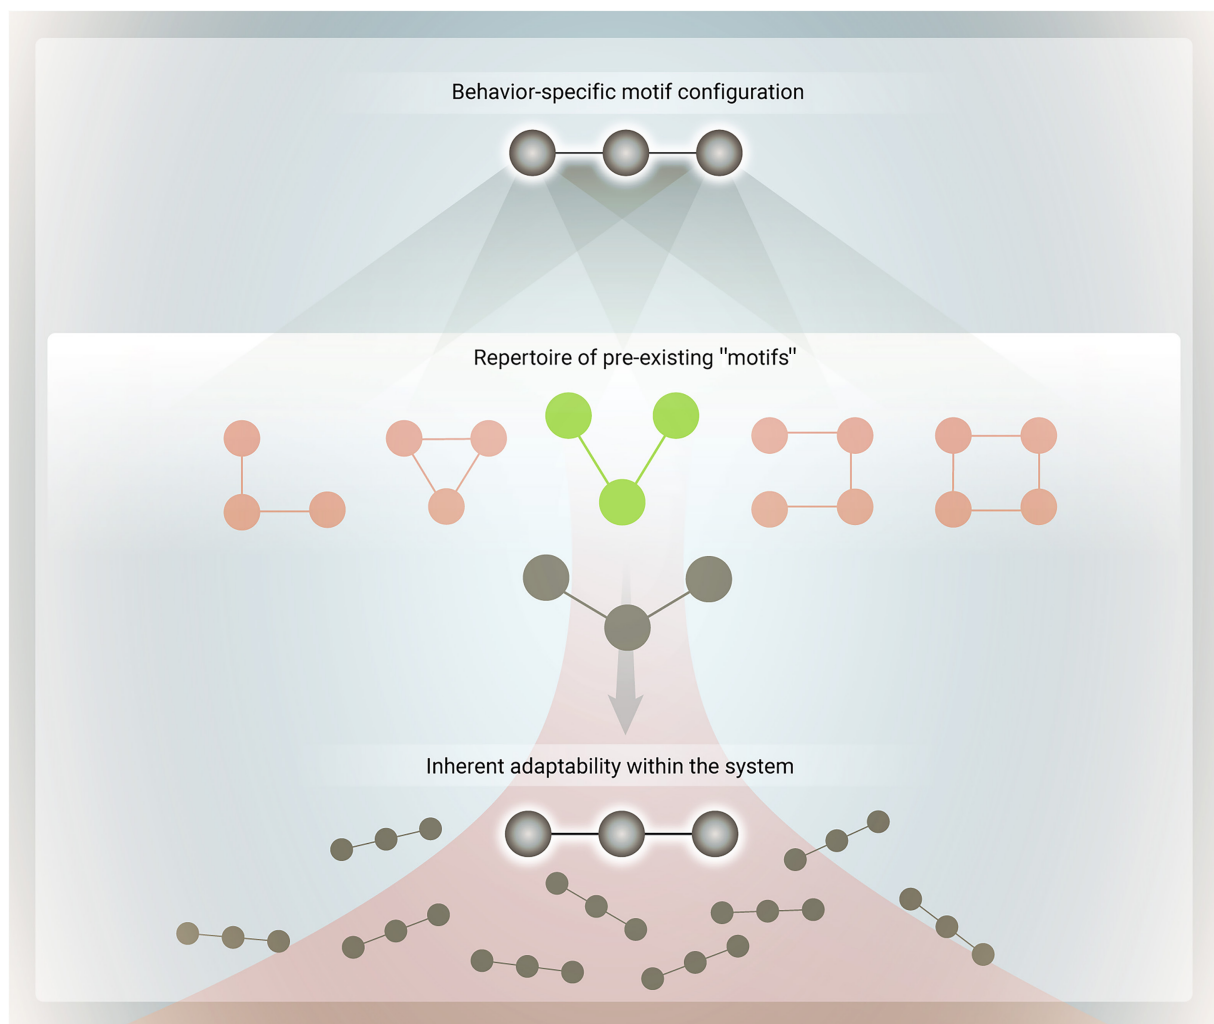
Learning relates to adaptability, representing the emergence of more effective behavior for a given context. As a reflection of possibility space, there may be several brain configurations that are appropriate for a given behavior (i.e., many-to-one mapping), but some configurations will be reinforced as optimal for that behavior. Figure 3 shows a schematic wherein a given behavioral challenge is addressed optimally by a particular motif of neural interactions. For the organism, that precise motif is not available, but pre-existing configurations that closely resemble it are. By small adaptations to existing motifs, the organism adapts its behavior and that motif becomes part of the dynamic repertoire.

Some recent work has suggested that potential configurations pre-exist in brain circuitry, which can be engaged for a given task and selected if the configuration is adaptive. For example, foundational research suggested that hippocampal ensembles “replayed” experiences, especially during sleep, which was thought to help reinforce memory traces (Wilson and McNaughton 1994). More recent work has now shown “pre-play” where ensemble activity related to learning seems to exist prior to the experience (Dragoi and Tonegawa 2011). A study by Mocle et al. (2024) provides an interpretation of the pre-play that is similar to adaptive immunity. Some ensemble activities prior to training showed similar patterns to the ensemble activity engendered by learning.<sup>12</sup> Thus, the brain establishes a possibility space that captures potentially useful configurations that can be reinforced if engaged in a given behavior.

This adaptation using pre-existing motifs is similar to “adaptive immunity” in the immune system (Rajewsky 1996; Rees 2020). As noted earlier, the immune system produces an astonishing diversity of antibodies that do not necessarily match any antigen the organism has encountered. However, when a new antigen is encountered, the immune system can mobilize the collection of pre-existing possible antibodies, adapting the configuration to best match the given antigen. Thus, the broad possibility space for the immune system enables rapid response (or adaptation) to a particular event (Rees 2020).

Numerous examples illustrate brain adaptation, often linked to neural plasticity. For instance, changes in neural firing





**FIGURE 3** | Symbolic depiction of an adaptive system (e.g., the brain and immune system), which contains several network motifs that support possible function capacity (pre-existing motifs). A specific external challenge requires a specific motif configuration. Although the configuration does not exactly map to pre-existing motifs, the system's adaptive capacity modifies an existing motif to match the challenge. The system may also replicate the optimal motif to impart additional resilience.

patterns in response to the varying significance of stimuli are some of the earliest examples of adaptability (Galambos et al. 1956; McGann 2015; Recanzone et al. 1992; Weinberger and Diamond 1987). This adaptability extends to neural circuits and networks, where within-session learning can lead to shifts in network responses (Buchel et al. 1999; Buchel and Friston 1997; McIntosh and Gonzalez-Lima 1994, 1998; McIntosh et al. 2001). These shifts can be modeled as changes in the possibility space landscape (Roy et al. 2014). These changes can introduce new repertoires within the possibility space or alter the accessibility of existing possibilities, sometimes making previously accessible possibilities less reachable. Both theoretical and empirical works suggest that this dynamic is represented by changing soft constraints and the available trajectories within the possibility space landscape. Moreover, work on the robustness and variability of neural circuits shows how neural systems can maintain function and adapt despite changes in their components (Marder and Goaillard 2006).

Given that neural plasticity is a fundamental property, adaptability plays a crucial role in maintaining desirable behavioral

outcomes across the lifespan, even though the capacity for change diminishes with age. Neuroimaging studies have shown that older adults often use different neural networks to perform behaviors comparable to those of younger adults (Grady 2012; Park and Reuter-Lorenz 2009). For instance, research indicates that older adults may engage different brain regions than younger adults to support similar behaviors (McIntosh et al. 1999). Older adults who do not shift to new networks often exhibit memory deficits compared to their peers (Cabeza 2002). Additionally, studies on network dynamics have shown that network topology and network flows change in healthy aging (Cabral et al. 2017). These changes can be linked to exploring and mobilizing possibility space to support optimal behavior. A lack of exploration of possibility space is often associated with poorer cognitive function (McIntosh 2019).

Adaptability, aimed at achieving functional outcomes, is a multidimensional process well-suited to the possibility space framework. This framework captures potential but unrealized outcomes, changes in capacity over time, and alterations in the landscape that are not evident in simpler models that assume a one-to-one mapping between brain configurations and function.

Possibility space frameworks are particularly useful for modeling how neural systems adapt and change over time. This can be achieved by modeling changes in a system's soft constraints over time, which capture states that are more or less likely to manifest, such as newly uncovered states and hidden states. This approach provides a comprehensive picture of a system's new (adapted) capacity and possible manifestations. Further research can identify various factors that act as soft constraints and explore the influences that hard and soft constraints have on adaptability.

## 5.2 | Resilience

Although brain resilience is defined in various ways, many of these definitions characterize it as the ability of the brain to produce positive, constructive outcomes in response to damaging or negative insults. This can involve a response to adversity that involves “bending and not breaking” or simply direct resistance through coping processes (McEwen et al. 2015). The possibility space framework helps model resilience for a few reasons. First, this framework helpfully represents potential backup responses or capacities when an insult strikes. Second, this framework can represent eliminated possibilities and situations where an insult triggers a new, unique response. Third, the possibility space model can capture how insults result in reconfigurations of the space that open hidden repertoires or previously isolated pockets of space. These opened spaces can include pathological states (cf. epilepsy) or functional repertoires. As Raja and Anderson state, “changing constraints in different ways can result in positive new capacities ... [even]...dysfunctional ones” (p. 212).


Resilience per se is hard to assess empirically prior to an event. For example, there is little to no data on the premorbid state in persons with a stroke, making it hard to link characteristics of

the hard and soft constraints that relate to better resilience. Most work has compared equivalent lesions or disease states (e.g., neurodegeneration) across an extensive demographic range to identify factors that relate to better recovery or reduce disease burden.

Jirsa (2020) has proposed a mathematical formalism for resilience in dynamical systems wherein the concept of “degeneracy,” or more formally “neural degeneracy,” is captured in the mapping between state variables that drive flows on manifolds (e.g., excitatory and inhibitory currents), versus parameter space that captures the distribution of parameter values for the state variables. Figure 4 illustrates two degeneracy manifolds in parameter space, one supporting state variables for singing and the other supporting state variables for walking. The collection of parameters within the degenerate manifolds leads qualitatively to the same flows in state space.

Any complex system with multiple scales and nonlinear properties will ultimately show degeneracy as a fundamental characteristic of its system behavior. The system can exploit the degeneracy manifolds to respond to altered conditions (such as injury) but within the hard constraints by moving along the degeneracy manifold and realizing different parameterizations of the system to accomplish the same behavior. The possibility space can thus be preserved. If a perturbation eliminates some aspects of the parameter distribution, the behavior in state space is relatively preserved.

The parameter distribution defines the possibility space, but here, in the case of resilience, the possibility space supports resilience in a somewhat paradoxical way. For adaptability, a broader possibility space would facilitate adaptations in an enhanced capacity to create new configurations or modify existing ones to enhance adaptation. One could envision the opposite for



**FIGURE 4** | Schematic of degenerate manifolds in parameter space that support different behaviors in state space. The parameter space is spanned by P1, P2, and P3, wherein manifolds of parameter values that support the same state behaviors are contained: in (A) the singing and in (B) walking. The distribution of parameters in the degenerate manifold imparts a degree of resilience in that a loss of a given parameter combination, indicated by points in the manifold, will have minimal effect on the state space behavior.

resilience, where too broad a possibility space would result in less resilience because there are too many options. This apparent paradox can be resolved by considering the configurations of systems like the brain, which occupy an intermediate position between strong resilience and strong adaptability. Others have used graph theory to capture the features of the brain from the perspective of a small world. Some have linked the brain feature to complexity, where the brain maximizes complexity through the intermediate position, which maximizes segregation and integration (Tononi et al. 1994).

The hard and soft constraints factor here in that hard constraints establish the boundaries of possibilities. The broader the possibility space, the more opportunities there are for adaptation. The hard constraints also provide the boundaries for redundancy, defining the possibility space that imparts resilience. This balance between adaptability and resilience is central to the brain as a complex adaptive system.

This balance represents flexibility captured by the possibility space containing alternative repertoires or trajectories that serve similar functions, which could support a positive outcome following stress or insult. This not only allows us to explain why one individual cannot maintain resilience (and overcome the insult), but it also helps capture comparative differences across individuals because of having different capacities in their possibility spaces. Suppose one individual has some or many backup options, whereas another has fewer or none. In that case, this helps explain why the first individual will be more resilient to ensuing threats, stresses, or insults. This is important because much of this explanation cites capacity, possibility, and potential spaces—this is a main advantage of this framework that many models cannot capture.

A notable exception is the conceptual model of “cognitive reserve” (Stern 2003). Especially in neurodegeneration, cognitive reserve has been used to convey an enhanced capacity for resilience in aging, where a person maintains cognitive function in the face of age-related decline in brain network integrity. Cognitive reserve imparts more resilience and is related to biological and lifestyle factors. The link to possibility space is clear, wherein a broader and more navigable possibility space would engender greater resilience and hence cognitive reserve. The two constructs can provide complementary views of resilience, but possibility space extends beyond cognition to cover any brain operations.

The intersection of adaptability and resilience relates to hidden repertoires. Although the notion of degeneracy mapping from Jirsa covers resilience for one possible state space configuration, configurations may produce behavioral equivalents from different state variables (e.g., different regions that interact). With adaptation and experience, some configurations will likely be used more frequently, but as with the epigenetic landscape, the configurations that are less used still exist. If the system is faced with a challenge from disease or trauma, the preferred configurations may be compromised, which would provide an avenue to access the less used “hidden” repertoires to maintain function. It may very well be that the decline in performance seen in aging may relate to the mobilization of configurations that produce behavioral equivalents qualitatively but not quantitatively.

### 5.3 | Phenomenology

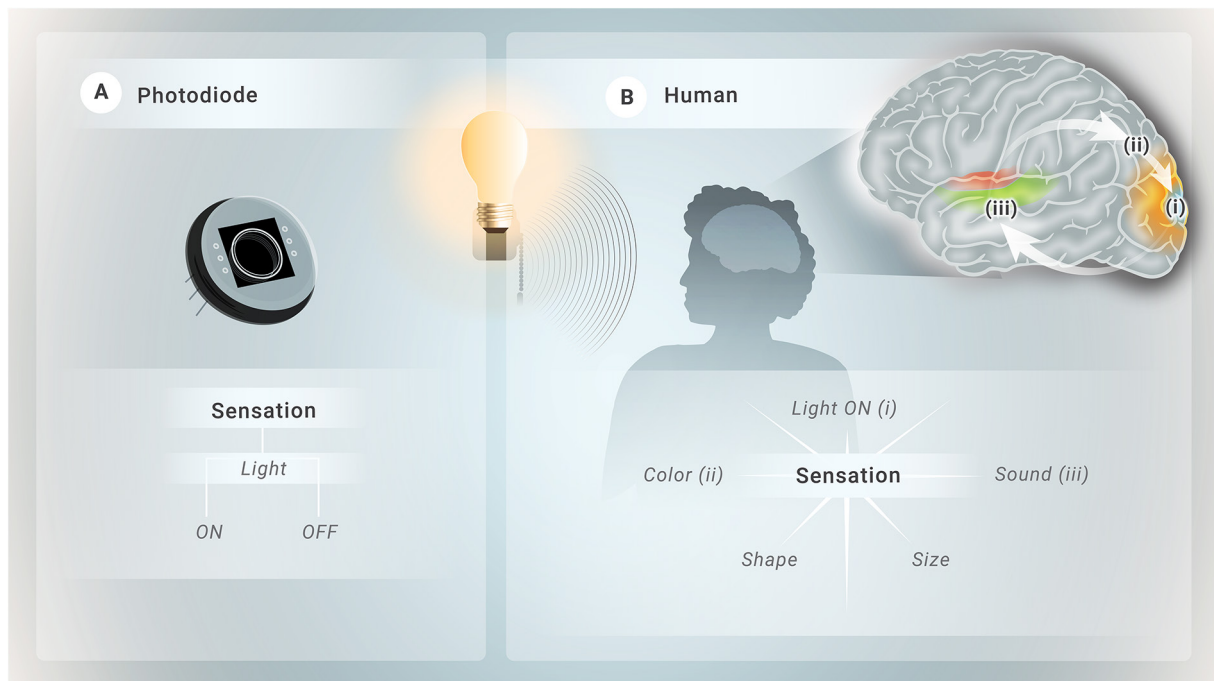
A final implication we address is that possibility space and hidden repertoires may be helpful in capturing notions of phenomenology. The possibility space defined by adaptive evolution relates to immediate experience, which is colored by what has happened and what could have happened—that is, what is possible.

Here, we can use the analogy Tononi (2004) developed about a photodiode and a human in a dark room. A small light comes on in the room. The photodiode registers the change, as does the human. If we express this detection as a landscape with an attractor configuration for “light on,” one can place this attractor in the broader landscape of possibility space. For the photodiode, its possibility space is only light-on versus light-off (this is even more impoverished as there would be no concept of “light” either). For the human, the possibility space has a light on versus off but also has myriad other possibilities: orange-yellow light, green light, and so on. The photodiode would only register “light-on” if the light were orange-yellow. If there was a “click” when the light came on, corresponding to the pull cord on the light, the human would register both the light and the sound of a click and note the orange-yellow hue of the light. The photodiode would only register light on. The click and “orange-yellow” do not exist for the photodiode—these are impossibility explanations (Figure 5).

Tononi further elaborates on this analogy by suggesting that our possibility provides both the experience of what did happen and what did not. The light was orange-yellow and accompanied by a click—it was not green, it was not multicolor, it was a light, not a smell, and so on. This richness of what did and did not happen influences our phenomenology. The poor photodiode will never know anything but “light on” versus “light off.” It would not even know it does not know that—such experiences are impossible for this system due to its various hard constraints. For the photodiode, it would be unfathomable.

We can further concretize this by connecting to clinical conditions that follow brain damage, which impairs access to possibility space and changes the person's phenomenology. One of the most profound examples is anosognosia from hemispatial neglect that follows damage to the right parietal lobe, thus affecting the left part of the world (Parton et al. 2004). For that patient, the left hemispace no longer exists as dictated by hard constraints resulting from brain damage. It is not a matter of a sensory deficit but rather that half of their world is not there. The symptoms can be so extreme that some patients will not acknowledge the possibility of a world to the left. For them, such spatial experiences are impossible—it would be ludicrous to suggest otherwise.

This example has further connections to the underlying brain dynamics. A study by Schirner et al. (2023) constructed large-scale brain network models integrating a “decision” module. The models were built from Human Connectome Project data, covering healthy young adults with individual variations in cognitive function as measured by fluid intelligence. The models showed differences in the exploration of attractor space/landscape in relation to both the task and fluid



**FIGURE 5** | Representation of the experiences a photodiode (left) and human (right) would have to a orange-yellow light, turned on with a pull chain. The photodiode can only detect the presence of light, whereas the human, with a broader possibility space, has more capacity. The binary space for the photodiode registers only on versus off. The more elaborate space for the human, encompassing brain networks that link inputs, captures light, and that it is a orange-yellow light and registers the click from the chain. The broad possibility for the human also captures other features of the light, embedded in the space of other possibilities that did not happen.

intelligence. Specifically, those with higher fluid intelligence measures explored the attractor space longer when task difficulty increased, leading to more accurate behavior. Although a complete characterization of the manifolds across variations in cognition was not done, one could infer that the manifold architecture changes with cognitive capacity, rendering more possibilities to be considered.

Like most complex adaptive systems, we consider the idea of an expansive manifold that encapsulates possibility space as an optimization target. To narrow a possibility space would engender rapid behavior for a range of tasks (e.g., the photodiode responds very rapidly to luminance change) but would falter as demands changed (red vs. green light). Broadening a space could lead to decision paralysis or engagement of attractors that may not be optimal for a behavioral demand. In this manner, possibility space frameworks have the potential to capture the richness, depth, and degree of phenomenology and experience. This is seen in comparing distinct systems (photodiode and human) and how particular systems change over time.

## 6 | What Is Next?

Accounts of constraints, possibility space, and impossibility explanation have received significant attention in philosophy but have received less formal attention in the neuroscience literature. These notions can support discussions around brain adaptation, resilience, and phenomenology by providing a framework to connect complex systems with similar properties. Such a framework provides conceptual clarity, a toolbox of

principled distinctions, and supports the identification of different processes that support function.

For neuroscience, in particular, the notion of impossibility explanations and possibility space open new avenues of question. If we consider brain dynamics as an expression of possibility space, how do we study it? Methods already exist that give us glimpses. Indeed, some of the work on network synchrony measured with electrophysiology can be considered a means to characterize possibility space. Buzsáki (2006), for example, considers oscillatory activity as supporting the coordination and integration of neural processes, which can be linked to exploring possibility space, as these dynamics reflect the brain's capacity to transition between different states. The perspective of Varela et al. (2001) underscores the importance of temporal coordination in brain function, which is a crucial aspect of navigating possibility space. In functional MRI, functional connectivity dynamics (FCD), for example, would express possibility space. Quantification of FCD is vital here, as one would expect the nature of possibility space to vary. As we mentioned earlier, the features of FCD change with age, and the nature of the change relates to cognitive measures.

Spatial determinants of possibility space need also be considered. Some recent intriguing work suggests that rudimentary features in the brain's spatial architecture constrain how information is represented at different spatial scales (Munn et al. 2024). Local scales carry more unique information with minimal redundancy between elements at the same scale, whereas larger scales show more redundancy, which enhances resilience (cf. Section 5.2). This would imply that possibility space would differ by spatial



scale. Moreover, this has some profound implications for how possibility space transcends spatial scales and interacts with temporal scales.

One can also consider the topology of possibility space as another feature to target. For example, the dimensions identified by principal components analysis would be a straightforward expression of the topology, although mainly in the sense of the number of dimensions that capture the functional space for a brain. Graph theory applications are ideal for defining hard and soft constraints that arise from anatomical and functional connections (Bullmore and Sporns 2009). More sophisticated approaches that derive more geometric characterizations, such as topological data analysis (Saggar et al. 2018), have provided more detail relating metrics like modularity to cognitive status. The application of control theory to brain networks further illuminates how control principles can be used to understand and potentially guide brain dynamics, making it a compelling addition to the study of possibility space (Bassett and Khambhati 2017).

A challenge with studying possibility space is that most empirical work will capture only part of a person's possibility space. To really characterize possibility space (over an individual's life) would involve measuring the brain from birth to death, 24 h a day, 7 days a week. This is obviously impractical. There are recent trends in "precision neuroimaging" that aspire to better characterize individual variation in functional brain networks by performing several sessions of data acquisition at longer durations. These data have demonstrated reliable stationary FC patterns across sessions, suggesting some proof of the concept of increased precision. More recently, FCD has been explored across data sets with precision scanning, suggesting that a core set of network hubs appears common across individuals, but the exploration around those hubs, where different regions are engaged and disengaged, appears unique to each person. How this relates to cognitive function remains to be explored (Saggar et al. 2022).

Precision neuroimaging gets us closer to characterizing a possibility space. However, we are still left with the challenge that all we can see is what happened during those scans—not the broader realm of possibilities. This is where computational modeling can help. The connectome-based brain simulation platform, The Virtual Brain (TVB), integrates structural and functional neuroimaging data into a generative network model, which can be personalized with an individual's brain imaging data. The virtual brain model for that person then becomes the vehicle to derive explanations for observed behavior. In the Schirner et al. paper, the dynamics were characterized in relation to manifold exploration. Explaining how each person's manifold is traversed is an expression of possibility space. Within the virtual brain model, parameter adjustments expand the characterization to express the potential better, or possibility, of using different trajectories. The multitude of options in this range follows from the soft constraints of the system.

Identifying hidden repertoires would also benefit from combining empirical and modeling work. The idea of a hidden repertoire first came from the modeling work of El Houssaini et al. (2015) on the Epileptor model. A detailed bifurcation analysis showed that the model could move from seizure activity to refractory

status epilepticus. This latter capacity of shift between seizure states in the repertoire was not evident during the initial creation of the model but was, in fact, a consequence of the modeling of the system's seizure activity. The existence of the state of refractory status epilepticus necessarily followed from properties of the seizure state. In other words, it was a "hidden repertoire," akin to the behaviors that are a mathematical consequence of dynamical systems, as phrased by Huang (2012). It stands to reason that parameter explorations done in other modeled systems may also reveal hidden repertoires. The modeled dynamics that arise from attractor navigation in these repertoires may provide clues as to whether they represent pathological configurations or potential for beneficial adaptation that reflects resilience.

Theoretical approaches to further characterizing (im)possibility constraints from a brain perspective could be fruitful. The structured flows on manifolds framework provides a formidable foundation in dynamical systems theory, emphasizing the emergence of networks that guide the evolution of behavior (Huys et al. 2014; Jirsa 2020; Pillai and Jirsa 2017). Manifolds move from simple forms to more complicated landscapes as connections become more heterogeneous/asymmetric. The attractors forming on these manifolds begin to define possibility space, further constrained by how sequential flows and manifolds are organized. What needs to be added to this formulation is why such architectures arise. A provocative complement to structured flows on manifolds is the free energy principle (Friston 2010), which provides optimization criteria to direct certain architectures over others. The principle reflects the assertion that the brain operates as a predictive machine, constantly updating its internal models to minimize free energy or surprise. The brain continuously adapts and optimizes its internal generative models (cf. Figure 3) by reducing prediction errors and guiding perception and action.

Integrating structured flows on manifolds and the free energy principle can illustrate how the brain's possibility space evolves as it learns and adapts. Structured flows on manifolds provide the paths and attractors, whereas free energy minimization drives the navigation and optimization of these paths. Over time,<sup>13</sup> this dynamic interplay can show how new paths and attractors emerge (representing adaptation and resilience) and how certain areas of the possibility space become more or less accessible depending on the brain's current state and experiences.

This paper has introduced a framework that provides conceptual clarity about possibility space in the context of complex systems in general and neuroscience in particular. We have paid particular attention to capturing a complex system's impossible states, possible states, and different types of possible states, including those that are more (or less) likely, more (or less) available, and hidden repertoires. This framework has further developed existing analytic philosophical work on constraints and scientific explanation in order to bridge philosophical and neuroscientific discussions of the possibility space concept.

An advantage of this work is that it captures challenging aspects of complex systems, such as different types of states that are "unrealized" for a system. For any system, there is a large range of possible yet unactualized states distinct from states that are strictly impossible for a system. These points can be hard

to appreciate because the range of possible states is extensive. Fewer concepts exist to clarify actual impossibilities for systems and the principled rationale for why they are impossible. Mere lack of realization is not an impossibility, and lack of realization can be difficult to model, represent, and conceptually appreciate. The importance of these points is only further magnified by the fact that when we consider systems that change, adapt, and evolve, these possible and impossible state landscapes change as well. Capturing these features is an important step in understanding, modeling, and advancing theoretical frameworks for complex systems.

## Author Contributions

**Lauren N. Ross:** writing – original draft, writing – review and editing. **Viktor Jirsa:** writing – original draft, writing – review and editing. **Anthony R. McIntosh:** writing – original draft, writing – review and editing.

## Acknowledgements

Art work for all figures by Alana McPherson (iamsoci.com).

## Conflicts of Interest

The authors declare no conflicts of interest.

## Data Availability Statement

The authors have nothing to report.

## Peer Review

The peer review history for this article is available at <https://www.webofscience.com/api/gateway/wos/peer-review/10.1111/ejn.70038>.

## Endnotes

<sup>1</sup> This is seen in work by Silberstein and Chemero (2013) on “global constraints” and “multiscale contextual constraints”; Lange (2016) on “constraint-based explanations”; Anderson (2016) on “enabling constraints”; and Ross (2023) on causal, mathematical, and physical law constraints.

<sup>2</sup> In this manner, structural connectivity “establishes a deterministic architecture” that captures spatial and temporal constraints. We discuss these and other brain-related cases in more detail shortly.

<sup>3</sup> For example, in the game’s opening, the Queen cannot move diagonal six spaces, and at no point in the game can the knight move just a single space forward.

<sup>4</sup> In fact, the rules of chess are often analogized to laws of nature—both play a constraining role on systems of interest, what they can and cannot exhibit, and how they can evolve over time.

<sup>5</sup> For example, both (a) physical laws (such as gravity) and (b) mathematical relationships (such as the square cube law) together explain why there are limits on the sizes and types of bodies of living systems on planet Earth, a finding identified as early as the work of Galileo. Both (a) and (b) clarify why some large body frames are structurally impossible (they would collapse under their own weight) and why small body frames are metabolically impossible (as they could not retain sufficient heat given to support metabolism).

<sup>6</sup> Maynard Smith et al. (1985) consider constraints in the context of organism development and identify various distinctions across constraint types, including how binding they are and how universal or local their constraining influence is.

<sup>7</sup> We do not argue that this is the only (or best) account of constraints relevant to explanation in neuroscience. In fact, we connect this framework to other helpful analyses of constraints found in the works of Lange (2016), Silberstein and Chemero (2013), and Raja and Anderson (2021).

<sup>8</sup> We should note that there are some less common situations in which the rules of chess (and hard constraints) can determine the next play of a game. This occurs when the rules dictate that there is only one available option (as all other options are unallowed). In this manner, when hard constraints limit to such an extreme degree, they can actually explain why a particular outcome presents (or will present) (Ross 2024).

<sup>9</sup> Our notion of soft constraints bears some similarity to the notion of “enabling constraints” discussed by Anderson (2016), Bolt et al. (2018), and Raja and Anderson (2021). Main differences are that our notion of soft constraints does not rely on a notion of function, we adopt the three criteria in Section 2 as capturing genetic constraints, and we distinguish soft from hard constraints (which meet these criteria yet differ in other ways). Other work in this area categorizes constraints differently (such as in terms of constraints that are structural, functional, strong, and weak) and a fruitful areas of future work involves identifying further types of constraints.

<sup>10</sup> For an excellent account of various type of constraints in the life sciences and ways in which constraints provide nonmechanistic explanations, see Silberstein and Chemero (2013, 380).

<sup>11</sup> In dynamical systems theory, these are often referred to as fast-slow systems or systems with multiple time scales.

<sup>12</sup> Interestingly, this pre-existing pattern was only present in the waking period that encompassed the training, suggesting that new arrays of possible motifs may be generated daily.

<sup>13</sup> Time in this context can be either developmental or evolutionary. In the former, the initial architectures, established through genetic programming, are modified through environmental experience. In the latter, similar to epigenetic landscapes, certain configurations are selected as more adaptive and maintained but with allowable variation within. Evolutionary pressures force the emergence of new possibilities, perhaps through the mobilization of hidden repertoires, that support the emergence of new architectures.

## References

- Anderson, M. L. 2016. “Beyond Componential Constitution in the Brain: Starburst Amacrine Cells and Enabling Constraints.” In *Open MIND, 2-Vol. Set: Philosophy and the Mind Sciences in the 21st Century*, 0. MIT Press. <https://doi.org/10.7551/mitpress/10603.003.0004>.
- Bassett, D. S., and A. N. Khambhati. 2017. “A Network Engineering Perspective on Probing and Perturbing Cognition With Neurofeedback.” *Annals of the New York Academy of Sciences* 1396, no. 1: 126–143. <https://doi.org/10.1111/nyas.13338>.
- Bechtel, W., and A. Levy. 2013. “Abstraction and the Organization of Mechanisms.” *Philosophy of Science* 80: 241–261.
- Bechtel, W., and R. C. Richardson. 2010. *Discovering Complexity*. Cambridge, MA: MIT Press.
- Bendor, J. 2001. “Bounded Rationality.” In *International Encyclopedia of the Social and Behavioral Sciences*, edited by N. J. Smelser and P. B. Baltes, 1303–1307. Elsevier.
- Bolt, T., M. L. Anderson, and L. Q. Uddin. 2018. “Beyond the Evoked/Intrinsic Neural Process Dichotomy.” *Network Neuroscience* 2, no. 1: 1–22. [https://doi.org/10.1162/NETN\\_a\\_00028](https://doi.org/10.1162/NETN_a_00028).
- Buchel, C., J. T. Coull, and K. J. Friston. 1999. “The Predictive Value of Changes in Effective Connectivity for Human Learning.” *Science* 283, no. 5407: 1538–1541.
- Buchel, C., and K. Friston. 1997. “Modulation of Connectivity in Visual Pathways by Attention: Cortical Interactions Evaluated With Structural Equation Modeling and fMRI.” *Cerebral Cortex* 7, no. 8: 768–778.

- Bullmore, E., and O. Sporns. 2009. "Complex Brain Networks: Graph Theoretical Analysis of Structural and Functional Systems." *Nature Reviews Neuroscience* 10, no. 3: nrr2575. <https://doi.org/10.1038/nrn2575>.
- Buzsáki, G. 2006. *Rhythms of the Brain*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195301069.001.0001>.
- Cabeza, R. 2002. "Hemispheric Asymmetry Reduction in Older Adults: The HAROLD Model." *Psychology and Aging* 17, no. 1: 85–100. <https://doi.org/10.1037/0882-7974.17.1.85>.
- Cabral, J., D. Vidaurre, P. Marques, et al. 2017. "Cognitive Performance in Healthy Older Adults Relates to Spontaneous Switching Between States of Functional Connectivity During Rest." *Scientific Reports* 7, no. 1: 5135. <https://doi.org/10.1038/s41598-017-05425-7>.
- Craver, C. 2007. *Explaining the Brain*. Oxford University Press.
- Dragoi, G., and S. Tonegawa. 2011. "Preplay of Future Place Cell Sequences by Hippocampal Cellular Assemblies." *Nature* 469, no. 7330: 397–401. <https://doi.org/10.1038/nature09633>.
- Dupré, J. 2013. "Living Causes." *Proceedings of the Aristotelian Society* 87: 19–37.
- El Houssaini, K., A. I. Ivanov, C. Bernard, and V. K. Jirsa. 2015. "Seizures, Refractory Status Epilepticus, and Depolarization Block as Endogenous Brain Activities." *Physical Review E, Statistical, Nonlinear, and Soft Matter Physics* 91, no. 1: 010701.
- Euler, L. 1956. "The Seven Bridges of Königsberg." In *The World of Mathematics*, edited by J. R. Newman, vol. 1, 573–580. London: George Allen and Unwin Ltd.
- Friston, K. J. 2010. "The Free-Energy Principle: A Unified Brain Theory?" *Nature Reviews Neuroscience* 11, no. 2: 127–138. <https://doi.org/10.1038/nrn2787>.
- Galambos, R., G. Sheatz, and V. G. Vernier. 1956. "Electrophysiological Correlates of a Conditioned Response in Cats." *Science (New York, N.Y.)* 123, no. 3192: 376–377. <https://doi.org/10.1126/science.123.3192.376>.
- Grady, C. 2012. "The Cognitive Neuroscience of Ageing." *Nature Reviews Neuroscience* 13, no. 7: 491–505. <https://doi.org/10.1038/nrn3256>.
- Holland, J. H. 2014. *Complexity: A Very Short Introduction*. Oxford Univ Press.
- Huang, S. 2012. "The Molecular and Mathematical Basis of Waddington's Epigenetic Landscape: A Framework for Post-Darwinian Biology?" *BioEssays* 34, no. 2: 149–157. <https://doi.org/10.1002/bies.201100031>.
- Huys, R., D. Perdikis, and V. K. Jirsa. 2014. "Functional Architectures and Structured Flows on Manifolds: A Dynamical Framework for Motor Behavior." *Psychological Review* 121, no. 3: 302–336. <https://doi.org/10.1037/a0037014>.
- Izhikevich, E. M. 2007. *Dynamical Systems in Neuroscience*. MIT Press.
- Jirsa, V. 2020. "Structured Flows on Manifolds as Guiding Concepts in Brain Science." In *Selbstorganisation - ein Paradigma für die Humanwissenschaften*, 89–102. Wiesbaden: Springer Fachmedien. [https://doi.org/10.1007/978-3-658-29906-4\\_6](https://doi.org/10.1007/978-3-658-29906-4_6).
- Khambhati, A. N., A. E. Sizemore, R. F. Betzel, and D. S. Bassett. 2018. "Modeling and Interpreting Mesoscale Network Dynamics." *NeuroImage* 180: 337–349. <https://doi.org/10.1016/j.neuroimage.2017.06.029>.
- Lange, M. 2016. *Because Without Cause: Non-Causal Explanations in Science and Mathematics*. Oxford University Press.
- Marder, E., and J. M. Goaillard. 2006. "Variability, Compensation and Homeostasis in Neuron and Network Function." *Nature Reviews Neuroscience* 7, no. 7: 563–574. <https://doi.org/10.1038/nrn1949>.
- McEwen, B. S., J. D. Gray, and C. Nasca. 2015. "Recognizing Resilience: Learning From the Effects of Stress on the Brain." *Neurobiology of Stress* 1: 1–11.
- McGann, J. P. 2015. "Associative Learning and Sensory Neuroplasticity: How Does It Happen and What Is It Good for?" *Learning & Memory (Cold Spring Harbor, N.Y.)* 22, no. 11: 567–576. <https://doi.org/10.1101/lm.039636.115>.
- McIntosh, A. R. 2019. "Neurocognitive Aging and Brain Signal Complexity." In *Oxford Research Encyclopedia of Psychology*. Oxford University Press. <https://doi.org/10.1093/acrefore/9780190236557.013.386>.
- McIntosh, A. R., and F. Gonzalez-Lima. 1994. "Network Interactions Among Limbic Cortices, Basal Forebrain and Cerebellum Differentiate a Tone Conditioned as a Pavlovian Excitor or Inhibitor. Fluorodeoxyglucose and Covariance Structural Modeling." *Journal of Neurophysiology* 72, no. 4: 1717–1733.
- McIntosh, A. R., and F. Gonzalez-Lima. 1998. "Large-Scale Functional Connectivity in Associative Learning: Interrelations of the Rat Auditory, Visual and Limbic Systems." *Journal of Neurophysiology* 80: 3148–3162.
- McIntosh, A. R., and V. K. Jirsa. 2019. "The Hidden Repertoire of Brain Dynamics and Dysfunction." *Network Neuroscience* 3, no. 4: 1–34. [https://doi.org/10.1162/netn\\_a\\_00107](https://doi.org/10.1162/netn_a_00107).
- McIntosh, A. R., M. N. Rajah, and N. J. Lobaugh. 2001. "Large-Scale Functional Connectivity in Learning With and Without Awareness During Sensory Differential Trace Conditioning." Submitted.
- McIntosh, A. R., A. B. Sekuler, C. Penpeci, et al. 1999. "Recruitment of Unique Neural Systems to Support Visual Memory in Normal Aging." *Current Biology* 9, no. 21: 1275–1278.
- Mocle, A. J., A. I. Ramsaran, A. D. Jacob, et al. 2024. "Excitability Mediates Allocation of Pre-Configured Ensembles to a Hippocampal Engram Supporting Contextual Conditioned Threat in Mice." *Neuron* 112, no. 9: 1486. <https://doi.org/10.1016/j.neuron.2024.02.007>.
- Munn, B. R., E. Müller, I. Favre-Bulle, E. Scott, M. Breakspear, and J. M. Shine. 2024. "Phylogenetically-Preserved Multiscale Neuronal Activity: Iterative Coarse-Graining Reconciles Scale-Dependent Theories of Brain Function." *BioRxiv* 2024.06.22.600219. <https://doi.org/10.1101/2024.06.22.600219>.
- Park, D. C., and P. Reuter-Lorenz. 2009. "The Adaptive Brain: Aging and Neurocognitive Scaffolding." *Annual Review of Psychology* 60: 173–196. <https://doi.org/10.1146/annurev.psych.59.103006.093656>.
- Parton, A., P. Malhotra, and M. Husain. 2004. "Hemispatial Neglect." *Journal of Neurology, Neurosurgery, and Psychiatry* 75, no. 1: 13–21.
- Pillai, A. S., and V. K. Jirsa. 2017. "Symmetry Breaking in Space-Time Hierarchies Shapes Brain Dynamics and Behavior." *Neuron* 94, no. 5: 1010–1026. <https://doi.org/10.1016/j.neuron.2017.05.013>.
- Raja, V., and M. L. Anderson. 2021. "Behavior Considered as an Enabling Constraint." In *Neural Mechanisms, Studies in Brain and Mind*, edited by F. Calzavarini and M. Viola, vol. 17. Cham: Springer. [https://doi.org/10.1007/978-3-030-54092-0\\_10](https://doi.org/10.1007/978-3-030-54092-0_10).
- Rajewsky, K. 1996. "Clonal Selection and Learning in the Antibody System." *Nature* 381, no. 6585: 751–758.
- Recanzone, G. H., C. E. Schreiner, and M. M. Merzenich. 1992. "Plasticity in the Frequency Representation of Primary Auditory Cortex Following Discrimination Training in Adult Owl Monkeys." *Journal of Neuroscience* 13, no. 1: 87–103.
- Rees, A. R. 2020. "Understanding the Human Antibody Repertoire." *MAbs* 12, no. 1: 1729683. <https://doi.org/10.1080/19420862.2020.1729683>.
- Ross, L. N. 2021a. "Causal Concepts in Biology: How Pathways Differ From Mechanisms and Why It Matters." *British Journal for the Philosophy of Science* 72: 131–158.
- Ross, L. N. 2021b. "Distinguishing Topological and Causal Explanation." *Synthese* 198, no. 10: 9803–9820.
- Ross, L. N. 2023. "The Explanatory Nature of Constraints: Law-Based, Mathematical, and Causal." *Synthese* 202, no. 2: 56.

- Ross, L. N. 2024. "What Is Social Structural Explanation? A Causal Account." *Noûs* 58, no. 1: 163–179.
- Roy, D., R. Sigala, M. Breakspear, et al. 2014. "Using the Virtual Brain to Reveal the Role of Oscillations and Plasticity in Shaping Brain's Dynamical Landscape." *Brain Connectivity* 4, no. 10: 791–811. <https://doi.org/10.1089/brain.2014.0252>.
- Saggar, M., J. M. Shine, R. Liégeois, N. U. F. Dosenbach, and D. Fair. 2022. "Precision Dynamical Mapping Using Topological Data Analysis Reveals a Hub-Like Transition State at Rest." *Nature Communications* 13, no. 1: 4791. <https://doi.org/10.1038/s41467-022-32381-2>.
- Saggar, M., O. Sporns, J. Gonzalez-Castillo, et al. 2018. "Towards a New Approach to Reveal Dynamical Organization of the Brain Using Topological Data Analysis." *Nature Communications* 9: 1399. <https://doi.org/10.1038/s41467-018-03664-4>.
- Schirner, M., G. Deco, and P. Ritter. 2023. "Learning How Network Structure Shapes Decision-Making for Bio-Inspired Computing." *Nature Communications* 14, no. 1: 2963. <https://doi.org/10.1038/s41467-023-38626-y>.
- Silberstein, M., and A. Chemero. 2013. "Constraints on Localization and Decomposition as Explanatory Strategies in the Biological Sciences." *Philosophy of Science* 80, no. 5: 958–970. <https://doi.org/10.1086/674533>.
- Simon, H. A. 1996. *The Sciences of the Artificial*. 3rd ed. MIT Press.
- Smith, J. M., R. Burian, S. Kauffman, et al. 1985. "Developmental Constraints and Evolution: A Perspective From the Mountain Lake Conference on Development and Evolution." *Quarterly Review of Biology* 60, no. 3: 265–287.
- Stern, Y. 2003. "The Concept of Cognitive Reserve: A Catalyst for Research." *Journal of Clinical and Experimental Neuropsychology* 25, no. 5: 589–593. [http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list\\_uids=12815497](http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=12815497).
- Tononi, G. 2004. "An Information Integration Theory of Consciousness." *BMC Neuroscience* 5, no. 1: 42. [http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list\\_uids=15522121](http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=15522121).
- Tononi, G., O. Sporns, and G. M. Edelman. 1994. "A Measure of Brain Complexity: Relating Functional Segregation and Integration in the Nervous System." *Proceedings of the National Academy of Science USA* 91: 5033–5037.
- Varela, F., J. P. Lachaux, E. Rodriguez, and J. Martinerie. 2001. "The Brainweb: Phase Synchronization and Large-Scale Integration." *Nature Reviews Neuroscience* 2, no. 4: 229–239. <https://doi.org/10.1038/35067550>.
- Waddington, C. H. 1957. *The Strategy of the Genes*. Routledge.
- Weinberger, N. M., and D. M. Diamond. 1987. "Physiological Plasticity in Auditory Cortex: Rapid Induction by Learning." *Progress in Neurobiology* 29, no. 1: 1–55.
- Wilson, M., and B. McNaughton. 1994. "Reactivation of Hippocampal Ensemble Memories During Sleep." *Science* 256: 676–679.