



Broken detailed balance and entropy production in the human brain

Christopher W. Lynn^{a,b}, Eli J. Cornblath^{c,d}, Lia Papadopoulos^e, Maxwell A. Bertolero^d, and Danielle S. Bassett^{c,e,f,g,h,i,j,1}

^aInitiative for the Theoretical Sciences, Graduate Center, City University of New York, New York, NY 10016; ^bJoseph Henry Laboratories of Physics, Princeton University, Princeton, NJ 08544; ^cDepartment of Bioengineering, School of Engineering & Applied Science, University of Pennsylvania, Philadelphia, PA 19104; ^dDepartment of Neuroscience, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA 19104; ^eInstitute of Neuroscience, University of Oregon, Eugene, OR 97403; ^fDepartment of Electrical & Systems Engineering, School of Engineering & Applied Science, University of Pennsylvania, Philadelphia, PA 19104; ^gDepartment of Physics & Astronomy, College of Arts & Sciences, University of Pennsylvania, Philadelphia, PA 19104; ^hDepartment of Neurology, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA 19104; ⁱDepartment of Psychiatry, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA 19104; and ^jSanta Fe Institute, Santa Fe, NM 87501

Edited by Marcus E. Raichle, Washington University in St. Louis, St. Louis, MO, and approved October 11, 2021 (received for review May 28, 2021)

Living systems break detailed balance at small scales, consuming energy and producing entropy in the environment to perform molecular and cellular functions. However, it remains unclear how broken detailed balance manifests at macroscopic scales and how such dynamics support higher-order biological functions. Here we present a framework to quantify broken detailed balance by measuring entropy production in macroscopic systems. We apply our method to the human brain, an organ whose immense metabolic consumption drives a diverse range of cognitive functions. Using whole-brain imaging data, we demonstrate that the brain nearly obeys detailed balance when at rest, but strongly breaks detailed balance when performing physically and cognitively demanding tasks. Using a dynamic Ising model, we show that these large-scale violations of detailed balance can emerge from fine-scale asymmetries in the interactions between elements, a known feature of neural systems. Together, these results suggest that violations of detailed balance are vital for cognition and provide a general tool for quantifying entropy production in macroscopic systems.

broken detailed balance | entropy production | cognitive neuroscience

The functions that support life—from processing information to generating forces and maintaining order—require organisms to break detailed balance (1, 2). For a system that obeys detailed balance, the fluxes of transitions between different states vanish (Fig. 1A). The system ceases to produce entropy and its dynamics become reversible in time. By contrast, living systems exhibit net fluxes between states or configurations (Fig. 1B), thereby breaking detailed balance and establishing an arrow of time (2). Critically, such broken detailed balance leads to the production of entropy, a fact first recognized by Sadi Carnot (3) in his pioneering studies of irreversible processes. At the molecular scale, metabolic and enzymatic activity drives nonequilibrium processes that are crucial for intracellular transport (4), high-fidelity transcription (5), and biochemical patterning (6). At the level of cells and subcellular structures, broken detailed balance enables sensing (7), adaptation (8), force generation (9), and structural organization (10).

Despite the importance of nonequilibrium dynamics at the microscale, there remain basic questions about the role of broken detailed balance in macroscopic systems composed of many interacting components. Do violations of detailed balance emerge at large scales? And, if so, do such violations support higher-order biological functions, just as microscopic broken detailed balance drives molecular and cellular functions?

To answer these questions, we study large-scale patterns of activity in the brain. Notably, the human brain consumes up to 20% of the body's energy to perform an array of cognitive functions, from computations and attention to planning and motor execution (11, 12), making it a promising system in which to probe for macroscopic broken detailed balance. Indeed, metabolic and enzymatic activity in the brain drives a number of nonequilibrium

processes at the microscale, including neuronal firing (13), molecular cycles (14), and cellular housekeeping (15). One might therefore conclude that the brain—indeed any living system—must break detailed balance at large scales. However, by coarse graining a system, one may average over nonequilibrium degrees of freedom, yielding “effective” macroscopic dynamics that produce less entropy (16, 17) and regain detailed balance (18). Thus, even though nonequilibrium processes are vital at molecular and cellular scales, it remains independently important to examine the role of broken detailed balance in the brain—and in complex systems generally—at large scales.

Fluxes and Broken Detailed Balance in the Brain

Here we develop tools to probe for and quantify broken detailed balance in macroscopic living systems. We apply our methods to analyze whole-brain dynamics from 590 healthy adults both at rest and across a suite of seven cognitive tasks, recorded using functional magnetic resonance imaging (fMRI) as part of the Human Connectome Project (19). For each cognitive task (including rest), the time-series data consist of blood-oxygen-level-dependent (BOLD) fMRI signals from 100 cortical parcels

Significance

To perform biological functions, living systems must break detailed balance by consuming energy and producing entropy. At microscopic scales, broken detailed balance enables a suite of molecular and cellular functions, including computations, kinetic proofreading, sensing, adaptation, and transportation. But do macroscopic violations of detailed balance enable higher-order biological functions, such as cognition and movement? To answer this question, we adapt tools from nonequilibrium statistical mechanics to quantify broken detailed balance in complex living systems. Analyzing neural recordings from hundreds of human subjects, we find that the brain violates detailed balance at large scales and that these violations increase with physical and cognitive exertion. Generally, we provide a flexible framework for investigating broken detailed balance at large scales in complex systems.

Author contributions: C.W.L. and D.S.B. designed research; C.W.L. performed research; C.W.L., E.J.C., and L.P. contributed new reagents/analytic tools; C.W.L. and M.A.B. analyzed data; and C.W.L. and D.S.B. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹To whom correspondence may be addressed. Email: dsb@seas.upenn.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2109889118/-DCSupplemental>.

Published November 17, 2021.

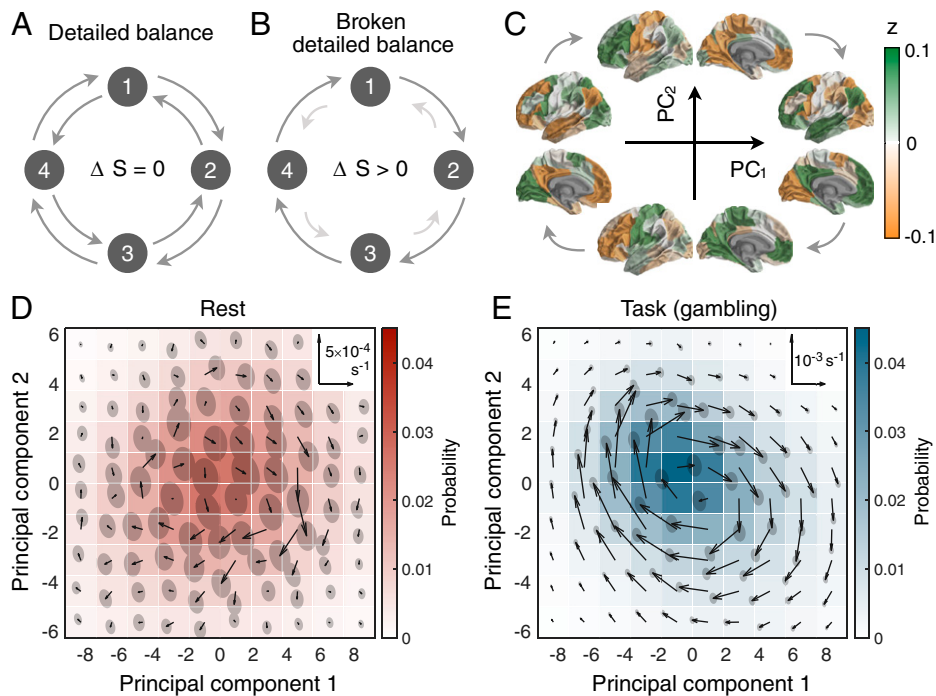


Fig. 1. Macroscopic broken detailed balance in the brain. (A and B) A simple four-state system, with states represented as circles and transition rates as arrows. (A) During detailed balance, there are no net fluxes of transitions between states, and the system does not produce entropy. (B) Systems that break detailed balance exhibit net fluxes of transitions between states, thereby producing entropy. (C) Brain states defined by the first two principal components of the neuroimaging time series of regional activity, computed across all time points and all subjects. Colors indicate the z-scored activation of different brain regions, ranging from high-amplitude activity (green) to low-amplitude activity (orange). Arrows represent hypothetical fluxes between states. (D and E) Probability distribution (color) and net fluxes between states (arrows) for neural dynamics at rest (D) and during a gambling task (E). To use the same axes in D and E, the dynamics are projected onto the first two principal components of the combined rest and gambling time-series data. The flux scale is indicated in the upper right, and the disks represent 2-SD confidence intervals that arise due to finite data (*Materials and Methods*).

(20), which we concatenate across all subjects (see *Materials and Methods* for an extended description of the neural data).

We begin by visually examining whether the neural dynamics break detailed balance. To visualize the dynamics, we must project the time series onto two dimensions. For example, here we project the neural dynamics onto the first two principal components of the time-series data, which we compute after combining all data points across all subjects (Fig. 1C). In fact, this projection defines a natural low-dimensional state space (21), capturing over 30% of the variance in the neural activity (*SI Appendix*). One can then probe for broken detailed balance by calculating the net fluxes of transitions between different regions of state space (22) (*Materials and Methods*). Moreover, we can repeat this analysis for different cognitive tasks to investigate whether the fluxes between neural states depend on the cognitive function being performed.

We first consider the brain's behavior during resting scans, wherein subjects are instructed to remain still without executing a specific task. At rest, we find that the brain exhibits net fluxes between states (Fig. 1D), thereby establishing that neural dynamics break detailed balance at large scales. But are violations of detailed balance determined solely by the structural connections in the brain, or does the nature of broken detailed balance depend on the specific function being performed?

To answer this question, we study task scans, wherein subjects respond to stimuli and commands that require attention, information processing, and physical and cognitive effort. For example, here we consider a gambling task in which subjects play a card guessing game for monetary reward. Interestingly, during the gambling task the fluxes between neural states are nearly an order of magnitude stronger than those present during rest (Fig. 1E). Moreover, these fluxes combine to form a distinct

loop in state space, a characteristic feature of broken detailed balance in steady-state systems (23), and we verify that the brain does indeed operate at a stochastic steady state (*SI Appendix*). To confirm that fluxes between neural states reflect broken detailed balance and are not simply artifacts of our data processing, we show that if the time series are shuffled—thereby destroying the temporal order of the system—then the fluxes between states vanish and detailed balance is restored (*SI Appendix*). Together, these results demonstrate that the brain fundamentally breaks detailed balance at large scales and that the strength of broken detailed balance depends critically on the cognitive function being performed.

Emergence of Macroscopic Broken Detailed Balance

We have established that the brain breaks detailed balance at large scales, exhibiting net fluxes between macroscopic neural states. This result builds upon recent measurements of broken detailed balance in a number of living systems (2, 4–10), including the brain itself (24). But can the large-scale violations of detailed balance that we observe in the brain emerge from fine-scale fluxes involving only one or two elements at a time?

To answer this question, we consider a canonical model of stochastic dynamics in complex systems. In the Ising model, the interactions between individual elements (or spins) are typically constrained to be symmetric, yielding simulated dynamics that obey detailed balance (25). However, connections in the brain—from synapses between neurons to white matter tracts between entire brain regions—are inherently asymmetric (12, 24, 26). Moreover, analytic techniques such as hidden Markov models have revealed that the effective connections between brain regions and neural states are also asymmetric (27–30). In the Ising model, if we allow asymmetric interactions, then the system

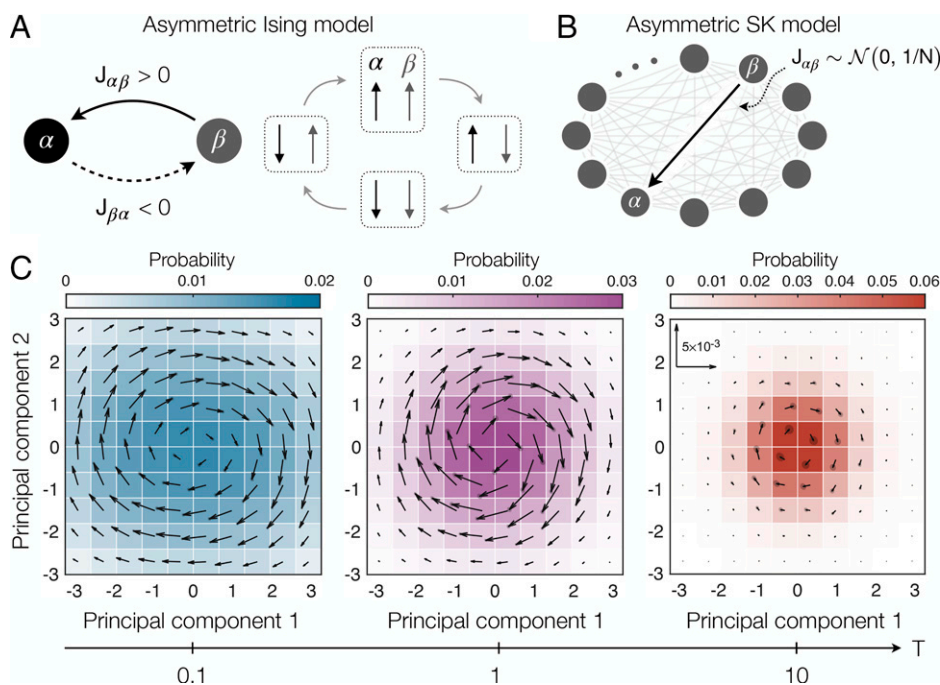


Fig. 2. Emergence of macroscopic broken detailed balance in a complex system. (A) Two-spin Ising model with interactions $J_{\alpha\beta}$ representing the strength of the influence of spin β on spin α (Left). If the interactions are asymmetric (such that $J_{\alpha\beta} \neq J_{\beta\alpha}$), then the system exhibits a loop of flux between spin states (Right). (B) Asymmetric SK model, wherein directed interactions are drawn independently from a zero-mean Gaussian with variance $1/N$, where N is the size of the system. (C) For an asymmetric SK model with $N = 100$ spins, we plot the probability distribution (color) and fluxes between states (arrows) for simulated time series at temperatures $T = 0.1$ (Left), $T = 1$ (Center), and $T = 10$ (Right). To visualize the dynamics, the time series are projected onto the first two principal components of the combined data across all three temperatures. The scale is indicated in flux per timestep, and the disks represent 2-SD confidence intervals that arise due to finite data (Materials and Methods).

diverges from broken detailed balance at small scales, displaying loops of flux involving pairs of spins (Fig. 24). But can such fine-scale fluxes combine to generate large-scale violations of detailed balance?

To understand whether (and how) microscopic asymmetries give rise to macroscopic broken detailed balance, we study a system of $N = 100$ spins (matching the 100 parcels in our neuroimaging data). Importantly, the system does not contain large-scale structure, with the interaction between each directed pair of spins drawn independently from a zero-mean Gaussian with variance $1/N$ (Fig. 2B). This model is the asymmetric generalization of the Sherrington–Kirkpatrick (SK) model of a spin glass (31). After simulating the system at three different temperatures, we perform the same procedure that we applied to the neuroimaging data (Fig. 1): projecting the time series onto the first two principal components of the combined data and measuring net fluxes in this low-dimensional state space.

At high temperature, stochastic fluctuations dominate the system, and we observe only weak fluxes between states (Fig. 2C, Right). By contrast, as the temperature decreases, the interactions between spins overcome the stochastic fluctuations, giving rise to clear loops of flux (Fig. 2C, Center and Left). These loops of flux demonstrate that asymmetries in the fine-scale interactions between elements alone can give rise to large-scale broken detailed balance. Moreover, by varying the strength of microscopic interactions, a single system can transition from exhibiting small violations of detailed balance to dramatic loops of flux, just as observed in the brain during different cognitive tasks (Fig. 1D and E).

Quantifying Broken Detailed Balance: Entropy Production

While fluxes in state space reveal violations of detailed balance, quantifying this behavior requires measuring the “distance” of a system from detailed balance. One such measure is entropy

production, the central concept of nonequilibrium statistical mechanics (2, 16, 32, 33). In microscopic systems, the rate at which entropy is produced—that is, the distance of the system from detailed balance—can often be directly related to the consumption of energy needed to drive cellular and subcellular functions (2, 7, 8). In macroscopic systems, this physical entropy production is lower bounded by an information-theoretic notion of entropy production, which can be estimated simply by observing a system’s coarse-grained dynamics (16, 33).

To begin, consider a system with joint transition probabilities $P_{ij} = \text{Prob}[x_{t-1} = i, x_t = j]$, where x_t is the state of the system at time t . We remark that P_{ij} differs from the conditional transition probabilities $\text{Prob}[x_t = j | x_{t-1} = i]$, which have been studied extensively in neural dynamics (28, 29). If the dynamics are Markovian (as, for instance, is true for the Ising system), then the information entropy production is given by (32)

$$\dot{S} = \sum_{ij} P_{ij} \log \frac{P_{ij}}{P_{ji}}, \quad [1]$$

where the sum runs over all states i and j . For simplicity, we refer to the information entropy production above simply as entropy production, not to be confused with the physical production of entropy at the microscale.

Inspecting Eq. 1, it becomes clear why entropy production is a natural measure of broken detailed balance: It is the Kullback–Leibler divergence between the forward transition probabilities P_{ij} and the reverse transition probabilities P_{ji} (34). If the system obeys detailed balance (that is, if $P_{ij} = P_{ji}$ for all pairs of states i and j), then the entropy production vanishes. Conversely, any violation of detailed balance (that is, any flux of transitions such that $P_{ij} \neq P_{ji}$) leads to an increase in entropy production. Moreover, we note that a system can still break detailed balance (thereby producing entropy $\dot{S} > 0$) even if the state probabilities

remain constant in time, as is the case in the Ising system (Fig. 2) and as we find to be the case in the neural data (SI Appendix). Such systems are said to operate at a nonequilibrium steady state (2, 32).

Calculating the entropy production requires estimating the joint transition probabilities P_{ij} . However, for complex systems the number of states grows exponentially with the size of the system, making a direct estimate of the entropy production infeasible. To overcome this hurdle, we employ a hierarchical clustering algorithm that groups similar states in a time series into a single cluster, yielding a reduced number of coarse-grained states (Fig. 3A and Materials and Methods). By choosing these clusters hierarchically (35), we prove that the estimated entropy production can only increase with the number of coarse-grained states—that is, as our description of the system becomes more fine grained (ignoring finite data effects; SI Appendix). Indeed, across all temperatures in the Ising system, we verify that the estimated entropy production increases with the number of clusters k (Fig. 3B). Furthermore, as the temperature decreases, the entropy production increases (Fig. 3B), thereby capturing the growing violations of detailed balance at low versus high temperatures (Fig. 2C).

Entropy Production in the Human Brain

We are now prepared to investigate the extent to which the brain breaks detailed balance when performing different functions. We study seven tasks, each of which engages a specific cognitive process and associated anatomical system: emotional processing, working memory, social inference, language processing, relational matching, gambling, and motor execution (36). To estimate the entropy production of the neural dynamics, we cluster the neuroimaging data (combined across all subjects and task settings, including rest) into $k = 8$ coarse-grained states, the largest number for which all transitions were observed at least once in each task (SI Appendix). Across all tasks and at rest, we find that the neural dynamics produce entropy, confirming that the brain breaks detailed balance at large scales (Fig. 4A). Specifically, for all settings the entropy production is significantly greater than the noise floor that arises due to finite data (one-sided t test with $P < 0.001$).

Interestingly, the neural dynamics produce more entropy during all of the cognitive tasks than at rest (Fig. 4A). In the motor task, for example—wherein subjects are prompted to perform specific physical movements—the entropy production is 20 times larger than for resting-state dynamics (Fig. 4A). In fact, while each cognitive task induces a unique pattern of fluxes between neural states, these fluxes nearly vanish during resting scans

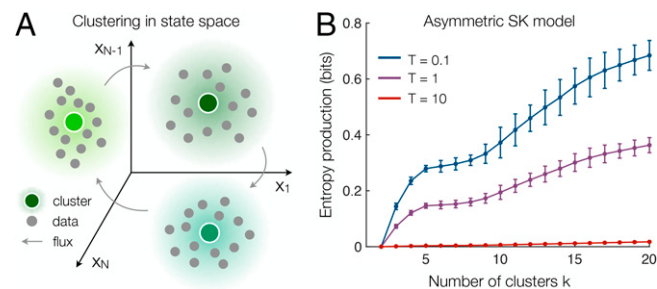


Fig. 3. Estimating entropy production using hierarchical clustering. (A) Schematic of clustering procedure, where axes represent the activities of individual components (e.g., brain regions in the neuroimaging data or spins in the Ising model), points reflect individual states observed in the time series, shaded regions define clusters (or coarse-grained states), and arrows illustrate possible fluxes between clusters. (B) Entropy production in the asymmetric SK model as a function of the number of clusters k for the same time series studied in Fig. 2C, with error bars reflecting 2-SD confidence intervals that arise due to finite data (Materials and Methods).




Fig. 4. Entropy production in the brain varies with physical and cognitive demands. (A) Entropy production at rest and during seven cognitive tasks, estimated using hierarchical clustering with $k = 8$ clusters. (B) Entropy production as a function of response rate (i.e., the frequency with which subjects are asked to physically respond) for the tasks listed in A. Each response induces an average 0.07 ± 0.03 bits of produced entropy (Pearson correlation $r = 0.774$, $p = 0.024$). (C) Entropy production for low cognitive load and high cognitive load conditions in the working memory task, where low and high loads represent 0-back and 2-back conditions, respectively, in an n -back task. The brain produces significantly more entropy during high-load than low-load conditions (one-sided t test, $P < 0.001$, $t > 10$, $df = 198$). In A–C, raw entropy productions (Eq. 1) are divided by the fMRI repetition time $\Delta t = 0.72$ s to compute an entropy production rate, and error bars reflect 2-SD confidence intervals that arise due to finite data (Materials and Methods).

(SI Appendix). Thus, we find that the extent to which the brain breaks detailed balance and the manner in which it does so depend critically on the specific task being performed.

The above results demonstrate that the brain breaks detailed balance at large scales as it executes physical movements, processes information, and performs cognitive functions. Indeed, just as energy is expended at the microscale to break detailed balance (2), one might expect violations of detailed balance in neural dynamics to increase with physical and cognitive exertion. To test the first hypothesis—that broken detailed balance in the brain is associated with physical effort—we compare the brain’s entropy production in each task with the frequency of physical movements (Fig. 4B). Across tasks, we find that entropy production does in fact increase with the frequency of motor responses, with each response yielding an additional 0.07 ± 0.03 bits of information entropy. Additionally, we confirm that this relationship between entropy production and physical effort also holds at the level of individual humans (SI Appendix).

Second, to study the impact of cognitive effort and information processing on broken detailed balance, we focus on the working memory task, which splits naturally into two conditions with high and low cognitive loads. Importantly, the frequency of physical responses is identical across the two conditions, thereby controlling for the effect of physical effort studied previously. We find that the brain operates farther from detailed balance when

exerting more cognitive effort (Fig. 4C), with the high-load condition inducing a twofold increase in entropy production over the low-load condition. Moreover, at the level of individuals, we find that entropy production increases with task errors (SI Appendix), once again indicating that violations of detailed balance intensify with cognitive demand.

Finally, we verify that our results do not depend on the Markov assumption in Eq. 1, are robust to reasonable variation in the number of clusters k , and cannot be explained by head motion in the scanner (a common confound in fMRI studies) (37), variance in the activity time series, or the block lengths of different tasks (SI Appendix). Moreover, across all tasks we confirm that the brain operates at a nonequilibrium steady state (SI Appendix). Together, these findings demonstrate that large-scale violations of detailed balance in the brain robustly increase with measures of both physical effort and cognitive demand. These conclusions, in turn, suggest that broken detailed balance in macroscopic systems may support higher-order biological functions.

Discussion

In this study, we describe a method for investigating macroscopic broken detailed balance by quantifying entropy production in living systems. While microscopic nonequilibrium processes are known to be vital for molecular and cellular operations (4–10), here we show that broken detailed balance also arises at large scales in complex living systems. Analyzing whole-brain imaging data, we demonstrate that the human brain breaks detailed balance at large scales. Moreover, we find that the brain’s entropy production (that is, its distance from detailed balance) varies critically with the specific function being performed, increasing with both physical and cognitive demands.

These results open the door for a number of important future directions. For example, while entropy production in the brain appears to increase with physical and cognitive exertion, these results do not preclude the possibility that other task- and stimulus-related factors may contribute to broken detailed balance. Specifically, one might suspect that by imposing external rhythms, such as repeated task blocks or oscillatory stimuli, one may be able to shift the brain farther from detailed balance. Additionally, given that large-scale violations of detailed balance can emerge from fine-scale asymmetries in a system (Fig. 2), future research should examine the relationship between broken detailed balance in the brain and asymmetries in the structural and functional connectivity between brain regions. Finally, recent work suggests that turbulent flow in the brain may facilitate the transfer of energy and information between regions (38). Given the intimate relationship between broken detailed balance and energy consumption at the molecular and cellular scales, one might consider whether entropy production in the brain is associated with increases in neural metabolism.

More generally, we remark that the presented framework is noninvasive, applying to any system with time-series data. Thus, the methods not only apply to the brain, but also can be used broadly to investigate broken detailed balance in other complex living systems, including emergent behavior in human and animal populations (39), correlated patterns of neuronal firing (40), and collective activity in molecular and cellular networks (41, 42). In fact, the framework is not even limited to living systems, which internally violate detailed balance, but can also be applied to nonbiological active systems, which are driven out of equilibrium by external forces (43).

Materials and Methods

Calculating Fluxes. Consider time-series data gathered in a time window t_{tot} and let n_{ij} denote the number of observed transitions from state i to state j . The flux rate from state i to state j is given by $\omega_{ij} = (n_{ij} - n_{ji})/t_{\text{tot}}$. For the flux currents in Figs. 1D and E and 2C, the states of the system are points (x, y) in two-dimensional space, and the state probabilities are

estimated by $p(x, y) = t_{(x,y)}/t_{\text{tot}}$, where $t_{(x,y)}$ is the time spent in state (x, y) . The magnitude and direction of the flux through a given state (x, y) are defined by the flux vector (22)

$$\mathbf{u}(x, y) = \frac{1}{2} \begin{pmatrix} \omega_{(x-1,y),(x,y)} + \omega_{(x,y),(x+1,y)} \\ \omega_{(x,y-1),(x,y)} + \omega_{(x,y),(x,y+1)} \end{pmatrix}. \quad [2]$$

In a small number of cases, two consecutive states in the observed time series $\mathbf{x}(t) = (x(t), y(t))$ and $\mathbf{x}(t+1) = (x(t+1), y(t+1))$ are not adjacent in state space. In these cases, we perform a linear interpolation between $\mathbf{x}(t)$ and $\mathbf{x}(t+1)$ to calculate the fluxes between adjacent states.

Estimating Finite-Data Errors Using Trajectory Bootstrapping. The finite length of time-series data limits the accuracy with which quantities—such as entropy production and the fluxes between states—can be estimated. To calculate error bars on all estimated quantities, we apply trajectory bootstrapping (22, 44). We first record the list of transitions

$$I = \begin{pmatrix} i_1 & i_2 \\ i_2 & i_3 \\ \vdots & \vdots \\ i_{L-1} & i_L \end{pmatrix}, \quad [3]$$

where i_ℓ is the ℓ th state in the time series, and L is the length of the time series. From the transition list I , one can calculate all of the desired quantities; for instance, the fluxes are estimated by

$$\omega_{ij} = \frac{1}{t_{\text{tot}}} \sum_{\ell} \delta_{i,i_{\ell,1}} \delta_{j,i_{\ell,2}} - \delta_{j,i_{\ell,1}} \delta_{i,i_{\ell,2}}. \quad [4]$$

We remark that when analyzing the neural data, although we concatenate the time series across subjects, we include only transitions in I that occur within the same subject. That is, we do not include the transitions between adjacent subjects in the concatenated time series.

To calculate errors, we construct bootstrap trajectories (of the same length L as the original data) by sampling the rows in I with replacement. For example, by calculating the entropy production in each of the bootstrap trajectories, we are able to estimate the size of finite-data errors in Figs. 3B and 4. Similarly, to compute errors for the flux vectors $\mathbf{u}(\mathbf{x})$ in Figs. 1D and E and 2C, we first estimate the covariance matrix $\text{Cov}(u_1(\mathbf{x}), u_2(\mathbf{x}))$ by averaging over bootstrapped trajectories. Then, for each flux vector, we visualize its error by plotting an ellipse with axes aligned with the eigenvectors of the covariance matrix and radii equal to twice the square root of the corresponding eigenvalues (SI Appendix). All errors throughout this article are calculated using 100 bootstrap trajectories.

The finite-data length also induces a noise floor for each quantity, which is present even if the temporal order of the time series is destroyed. To estimate the noise floor, we construct bootstrap trajectories by sampling individual data points from the time series. We contrast these bootstrap trajectories with those used to estimate errors above, which preserve transitions by sampling the rows in I . The noise floor, which is calculated for each quantity by averaging over the bootstrap trajectories, is then compared with the estimated quantities. For example, rather than demonstrating that the average entropy productions in Fig. 4A are greater than zero, we establish that the distribution over entropy productions is significantly greater than the noise floor using a one-sided t test with $P < 0.001$.

Simulating the Asymmetric Ising Model. The asymmetric Ising model is defined by a (possibly asymmetric) interaction matrix J , where $J_{\alpha\beta}$ represents the influence of spin β on spin α (Fig. 2A), and a temperature $T \geq 0$ that tunes the strength of stochastic fluctuations. Here, we study a system with $N = 100$ spins, where each directed interaction $J_{\alpha\beta}$ is drawn independently from a zero-mean Gaussian with variance $1/N = 0.01$ (Fig. 2B). One can additionally include external fields h_α , but for simplicity here we set them to zero. The state of the system is defined by a vector $\mathbf{x} = (x_1, \dots, x_N)$, where $x_\alpha = \pm 1$ is the state of spin α . To generate time series, we employ Glauber dynamics with synchronous updates, a common Monte Carlo method for simulating Ising systems (25). Specifically, given the state of the system $\mathbf{x}(t)$ at time t , the probability of spin α being “up” at time $t+1$ (that is, the probability that $x_\alpha(t+1) = 1$) is given by

$$\text{Prob}[x_\alpha(t+1) = 1 | \mathbf{x}(t)] = \left[1 + \exp \left(- \frac{2}{T} \sum_{\beta} J_{\alpha\beta} x_\beta(t) \right) \right]^{-1}. \quad [5]$$

Stochastically updating each spin α according to Eq. 5, one arrives at the new state $\mathbf{x}(t+1)$. For each temperature in the Ising calculations in Figs. 2C and 3B, we generate a different time series of length $L = 100,000$ with 10,000 trials of burn-in.

Table 1. Parameters of task and rest scans

Task	Duration, s	Response rate, s^{-1}	Block length, s
Rest	873	0	—
Emotion	136	0.260	18
Working memory	301	0.053	25
Social	207	0.025	23
Language	237	0.058	30
Relational	176	0.159	16
Gambling	192	0.172	28
Motor	212	0.477	12

Hierarchical Clustering. To estimate the entropy production of a system, one must first calculate the joint transition probabilities $P_{ij} = n_{ij}/(L - 1)$. For complex systems, the number of states i (and therefore the number of transitions $i \rightarrow j$) grows exponentially with the size of the system N . For example, in the Ising model each spin α can take one of two values ($x_\alpha = \pm 1$), leading to 2^N possible states and 2^{2N} possible transitions. To estimate the transition probabilities P_{ij} , one must observe each transition $i \rightarrow j$ at least once, which requires significantly reducing the number of states in the system. Rather than defining coarse-grained states a priori, complex systems (and the brain in particular) often admit natural coarse-grained descriptions that are uncovered through dimensionality-reduction techniques (21, 29, 45).

Although one can use any coarse-graining technique to implement our framework and estimate entropy production, here we employ hierarchical k -means clustering for two reasons: 1) k -means is perhaps the most common and simplest clustering algorithm, with demonstrated effectiveness fitting neural dynamics (29, 45); and 2) by defining the clusters hierarchically, we are able to prove that the estimated entropy production becomes more accurate as the number of clusters increases (ignoring finite-data effects; *SI Appendix*).

In k -means clustering, one begins with a set of states (for example, those observed in our time series) and a number of clusters k . Each observed state x is randomly assigned to a cluster i , and one computes the centroid of each cluster. On the following iteration, each state is reassigned to the cluster with the closest centroid (here we use cosine similarity to determine distance). This process is repeated until the cluster assignments no longer change. In a hierarchical implementation, one begins with two clusters; then one cluster is selected (typically the one with the largest spread in its constituent states) to be split into two new clusters, thereby defining a total of three clusters. This iterative splitting is continued until one reaches the desired number of clusters k . In *SI Appendix*, we show that hierarchical clustering provides an increasing lower bound on the entropy production, and we provide a principled method for choosing the number of clusters k .

Neural Data. The whole-brain dynamics used in this study are measured and recorded using BOLD fMRI collected from 590 healthy adults as part of the Human Connectome Project (19, 36). For each subject, recordings were taken during seven different cognitive tasks and also during rest (see Table 1 and

ref. 36 for details of task designs). BOLD fMRI estimates neural activity by calculating contrasts in blood oxygen levels, without relying on invasive injections and radiation (46). Specifically, blood oxygen levels (reflecting neural activity) are measured within three-dimensional nonoverlapping voxels, spatially contiguous collections of which each represent a distinct brain region (or parcel). Here, we consider a parcellation that divides the cortex into 100 brain regions that are chosen to optimally capture the functional organization of the brain (20). After processing the signal to correct for sources of systematic noise such as head motion (*SI Appendix*), the activity of each brain region is discretized in time, yielding a time series of neural activity. For each subject, the shortest scan (corresponding to the emotional processing task) consists of 176 discrete measurements in time. To control for variability in data size across tasks, for each subject we study only the first 176 measurements in each task.

Citation Diversity Statement. Recent work in several fields of science has identified a bias in citation practices such that papers from women and other minorities are undercited relative to the number of such papers in the field (47–52). Here we sought to proactively consider choosing references that reflect the diversity of the field in thought, form of contribution, gender, and other factors. We obtained predicted gender of the first and last author of each reference by using databases that store the probability of a name being carried by a woman (51, 53). By this measure (and excluding self-citations to the first and last authors of our current paper), our references contain 7% woman(first)/woman(last), 14% man/woman, 21% woman/man, and 58% man/man. This method is limited in that 1) names, pronouns, and social media profiles used to construct the databases may not, in every case, be indicative of gender identity, and 2) it cannot account for intersex, nonbinary, or transgender people. We look forward to future work that could help us to better understand how to support equitable practices in science.

Data Availability. The data analyzed in this paper and the code used to perform the analyses are publicly available at GitHub, github.com/ChrisWlynn/Broken_detailed_balance (54). Previously published data were used for this work (19).

ACKNOWLEDGMENTS. We thank Erin Teich, Pragya Srivastava, Jason Kim, and Zhixin Lu for feedback on earlier versions of this manuscript. C.W.L. acknowledges support from the James S. McDonnell Foundation 21st Century Science Initiative Understanding Dynamic and Multi-Scale Systems postdoctoral fellowship award. We also acknowledge support from the John D. and Catherine T. MacArthur Foundation, the Institute for Scientific Interchange Foundation, the Paul G. Allen Family Foundation, the Army Research Laboratory grant (W911NF-10-2-0022), the Army Research Office grants (W911NF-14-1-0679, W911NF-18-1-0244, W911NF-16-1-0474, and DCIST-W911NF-17-2-0181), the Office of Naval Research, the National Institute of Mental Health grants (2-R01-DC-009209-11, R01-MH112847, R01-MH107235, and R21-MH-106799), the National Institute of Child Health and Human Development grant (1R01HD086888-01), the National Institute of Neurological Disorders and Stroke grant (R01 NS099348), and the NSF grants (NSF PHY-1554488, BCS-1631550, and NCS-FO-1926829).

1. E. Schrödinger, *What is Life? The Physical Aspect of the Living Cell and Mind* (Cambridge University Press, Cambridge, UK, 1944).
2. F. S. Gnesotto, F. Mura, J. Gladrow, C. P. Broedersz, Broken detailed balance and non-equilibrium dynamics in living systems: A review. *Rep. Prog. Phys.* **81**, 066601 (2018).
3. S. Carnot, *Reflexions Sur la Puissance Motrice du Feu* (Bachelier, Paris, France, 1824).
4. C. P. Brangwynne, G. H. Koenderink, F. C. MacKintosh, D. A. Weitz, Cytoplasmic diffusion: Molecular motors mix it up. *J. Cell Biol.* **183**, 583–587 (2008).
5. H. Yin, I. Artsimovitch, R. Landick, J. Gelles, Nonequilibrium mechanism of transcription termination from observations of single RNA polymerase molecules. *Proc. Natl. Acad. Sci. U.S.A.* **96**, 13124–13129 (1999).
6. K. C. Huang, Y. Meir, N. S. Wingreen, Dynamic structures in *Escherichia coli*: Spontaneous formation of MinE rings and MinD polar zones. *Proc. Natl. Acad. Sci. U.S.A.* **100**, 12724–12728 (2003).
7. P. Mehta, D. J. Schwab, Energetic costs of cellular computation. *Proc. Natl. Acad. Sci. U.S.A.* **109**, 17978–17982 (2012).
8. G. Lan, P. Sartori, S. Neumann, V. Sourjik, Y. Tu, The energy-speed-accuracy tradeoff in sensory adaptation. *Nat. Phys.* **8**, 422–428 (2012).
9. M. Soares e Silva *et al.*, Active multistage coarsening of actin networks driven by myosin motors. *Proc. Natl. Acad. Sci. U.S.A.* **108**, 9408–9413 (2011).
10. B. Stuhmann, M. Soares e Silva, M. Depken, F. C. Mackintosh, G. H. Koenderink, Nonequilibrium fluctuations of a remodeling *in vitro* cytoskeleton. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* **86**, 020901 (2012).
11. J. J. Harris, R. Jolivet, D. Attwell, Synaptic energy use and supply. *Neuron* **75**, 762–777 (2012).
12. C. W. Lynn, D. S. Bassett, The physics of brain network structure, function and control. *Nat. Rev. Phys.* **1**, 318 (2019).
13. M. Erecińska, I. A. Silver, ATP and brain function. *J. Cereb. Blood Flow Metab.* **9**, 2–19 (1989).
14. K. Norberg, B. K. Siejö, Cerebral metabolism in hypoxic hypoxia. II. Citric acid cycle intermediates and associated amino acids. *Brain Res.* **86**, 45–54 (1975).
15. F. Du *et al.*, Tightly coupled brain activity and cerebral ATP metabolic rate. *Proc. Natl. Acad. Sci. U.S.A.* **105**, 6409–6414 (2008).
16. M. Esposito, Stochastic thermodynamics under coarse graining. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* **85**, 041125 (2012).
17. I. A. Martinez, G. Bisker, J. M. Horowitz, J. M. R. Parrondo, Inferring broken detailed balance in the absence of observable currents. *Nat. Commun.* **10**, 3542 (2019).
18. D. A. Egolf, Equilibrium regained: From nonequilibrium chaos to statistical mechanics. *Science* **287**, 101–104 (2000).
19. D. C. Van Essen *et al.*, The WU-Minn Human Connectome Project: An overview. *Neuroimage* **80**, 62–79 (2013).
20. B. T. Yeo *et al.*, The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J. Neurophysiol.* **106**, 1125–1165 (2011).
21. J. P. Cunningham, B. M. Yu, Dimensionality reduction for large-scale neural recordings. *Nat. Neurosci.* **17**, 1500–1509 (2014).
22. C. Battle *et al.*, Broken detailed balance at mesoscopic scales in active biological systems. *Science* **352**, 604–607 (2016).
23. R. K. P. Zia, B. Schmittmann, Probability currents as principal characteristics in the statistical mechanics of non-equilibrium steady states. *J. Stat. Mech.* **2007**, P07012 (2007).
24. K. J. Friston *et al.*, Parcels and particles: Markov blankets in the brain. *Netw. Neurosci.* **5**, 211–251 (2021).
25. M. Newman, G. Barkema, *Monte Carlo Methods in Statistical Physics* (Oxford University Press, New York, 1999).

26. P. Kale, A. Zalesky, L. L. Gollo, Estimating the impact of structural directionality: How reliable are undirected connectomes? *Netw. Neurosci.* **2**, 259–284 (2018).
27. A. P. Baker *et al.*, Fast transient networks in spontaneous human brain activity. *eLife* **3**, e01867 (2014).
28. D. Vidaurre, S. M. Smith, M. W. Woolrich, Brain network dynamics are hierarchically organized in time. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 12827–12832 (2017).
29. E. J. Cornblath *et al.*, Temporal sequences of brain activity at rest are constrained by white matter structure and modulated by cognitive demands. *Commun. Biol.* **3**, 261 (2020).
30. T. S. Zarghami, K. J. Friston, Dynamic effective connectivity. *Neuroimage* **207**, 116453 (2020).
31. D. Sherrington, S. Kirkpatrick, Solvable model of a spin-glass. *Phys. Rev. Lett.* **35**, 1792 (1975).
32. U. Seifert, Entropy production along a stochastic trajectory and an integral fluctuation theorem. *Phys. Rev. Lett.* **95**, 040602 (2005).
33. E. Roldán, J. M. Parrondo, Estimating dissipation from single stationary trajectories. *Phys. Rev. Lett.* **105**, 150607 (2010).
34. T. M. Cover, J. A. Thomas, *Elements of Information Theory* (John Wiley & Sons, 2012).
35. S. Lamrous, M. Taïleb, “Divisive hierarchical k-means” in *International Conference on Computational Intelligence for Modelling Control and Automation* (IEEE, 2006), p. 18.
36. D. M. Barch *et al.*, Function in the human connectome: Task-fMRI and individual differences in behavior. *Neuroimage* **80**, 169–189 (2013).
37. K. J. Friston, S. Williams, R. Howard, R. S. Frackowiak, R. Turner, Movement-related effects in fMRI time-series. *Magn. Reson. Med.* **35**, 346–355 (1996).
38. G. Deco, M. L. Kringelbach, Turbulent-like dynamics in the human brain. *Cell Rep.* **33**, 108471 (2020).
39. C. Castellano, S. Fortunato, V. Loreto, Statistical physics of social dynamics. *Rev. Mod. Phys.* **81**, 591 (2009).
40. J. M. Palva *et al.*, Neuronal long-range temporal correlations and avalanche dynamics are correlated with behavioral scaling laws. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 3585–3590 (2013).
41. G. H. Koenderink *et al.*, An active biopolymer network controlled by molecular motors. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 15192–15197 (2009).
42. L. Van Aelst, C. D'Souza-Schorey, Rho GTPases and signaling networks. *Genes Dev.* **11**, 2295–2322 (1997).
43. S. Ramaswamy, The mechanics and statistics of active matter. *Annu. Rev. Condens. Matter Phys.* **1**, 323–345 (2010).
44. C. E. Shannon, A mathematical theory of communication. *Bell Syst. Tech. J.* **27**, 379–423 (1948).
45. X. Liu, J. H. Duyn, Time-varying functional network information extracted from brief instances of spontaneous brain activity. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 4392–4397 (2013).
46. M. E. Raichle, Behind the scenes of functional brain imaging: A historical and physiological perspective. *Proc. Natl. Acad. Sci. U.S.A.* **95**, 765–772 (1998).
47. S. M. Mitchell, S. Lange, H. Brus, Gendered citation patterns in international relations journals. *Int. Stud. Perspect.* **14**, 485–492 (2013).
48. M. L. Dion, J. L. Sumner, S. M. Mitchell, Gendered citation patterns across political science and social science methodology fields. *Polit. Anal.* **26**, 312–327 (2018).
49. N. Caplar, S. Tacchella, S. Birrer, Quantitative evaluation of gender bias in astronomical publications from citation counts. *Nat. Astron.* **1**, 1–5 (2017).
50. D. Maliniak, R. Powers, B. F. Walter, The gender citation gap in international relations. *Int. Organ.* **67**, 889–922 (2013).
51. J. D. Dworkin *et al.*, The extent and drivers of gender imbalance in neuroscience reference lists. *Nat. Neurosci.* **23**, 918–926 (2020).
52. A. Maxwell *et al.*, Racial and ethnic imbalance in neuroscience reference lists and intersections with gender. *bioRxiv* [Preprint] (2020). <https://www.biorxiv.org/content/10.1101/2020.10.12.336230v1> (Accessed 30 April 2020).
53. D. Zhou *et al.*, Diversity statement and code notebook v1.1 (2020). <https://github.com/dalejn/cleanBib> (Accessed 30 April 2020).
54. C. W. Lynn, E. J. Cornblath, L. Papadopoulos, M. A. Bertolero, D. S. Bassett, Broken detailed balance and entropy production in the human brain. GitHub. https://github.com/ChrisWLYnn/Broken_detailed_balance. Deposited 16 September 2021.