Neuroscience and Biobehavioral Reviews 102 (2019) 1–12

Contents lists available at ScienceDirect

Neuroscience and Biobehavioral Reviews

journal homepage: www.elsevier.com/locate/neubiorev

Prediction and perception: Insights for (and from) tinnitus

Jeffrey Hullfish[a,*,1], William Sedley[b], Sven Vanneste[a,c]

[a] School of Behavioral and Brain Sciences, The University of Texas at Dallas, Richardson, TX 75080, USA
[b] Institute of Neuroscience, Newcastle University Medical School, Newcastle upon Tyne, UK
[c] Global Brain Health Institute, Institute of Neuroscience, Trinity College Dublin, Dublin, Ireland

ARTICLE INFO

Keywords:
Active inference
Bayes
Belief
Inference
Learning
Perception
Phantom perception
Prediction error
Predictive coding
Tinnitus
Unsupervised learning

ABSTRACT

More than 150 years have passed since Helmholtz first described perception as a process of unconscious inference about the causes of sensations. His ideas have since inspired a wealth of literature investigating the mechanisms underlying these inferences. In recent years, much of this work has converged on the notion that the brain is a hierarchical generative model of its environment that predicts sensations and updates itself based on prediction errors. Here, we build a case for modeling tinnitus from this perspective, i.e. predictive coding. We emphasize two key claims: (1) acute tinnitus reflects an increase in sensory precision in related frequency channels and (2) chronic tinnitus reflects a change in the brain's default prediction. We further discuss specific neural biomarkers that would constitute evidence for or against these claims. Finally, we explore the implications of our model for clinical intervention strategies. We conclude that predictive coding offers the basis for a unifying theory of cognitive neuroscience, which we demonstrate with several examples linking tinnitus to other lines of brain research.

1. Introduction

Tinnitus is a simple auditory percept that occurs in the absence of a corresponding acoustic stimulus. It is best known as a "ringing in the ears" though it can also manifest as more noise-like, e.g. hissing, buzzing, chirping, etc. The simplicity of tinnitus, however, belies the fact that it is a symptom of a complex—and, to date, incompletely understood—pathology. More to the point, a growing body of evidence suggests that there are several different mechanisms of action that can all lead to tinnitus, which has led to the idea that there are so-called "subtypes" of tinnitus. This complexity has hindered attempts by both academics and clinicians alike to understand and ultimately treat this condition.

In the search for a unifying mechanism, some tinnitus researchers have begun to turn to models of brain function predicated on predictive processing; see e.g. (De Ridder et al., 2014c, 2014a, Hullfish et al., 2018a, 2019; Kumar et al., 2014; Lee et al., 2017; Sedley et al., 2016a; Vanneste and De Ridder, 2016). This framework traces its roots back to Helmholtz (1867), who, in his seminal work on vision, famously described perception as a process of unconscious inference about the causes of sensory data. In other words, perception is the brain's inferred cause of sensory neural activity. The contemporary version of this framework relies on predictive coding (Friston, 2018; Rao and Ballard, 1999), which we discuss in detail throughout this paper but especially in Sections 2 and 3. The idea in brief is that the brain predicts sensations and the only aspect of sensory input encoded is prediction errors, which it uses to update its predictions via Bayesian inference (Fig. 1). While predictive coding and related theories (e.g. active inference; see Friston, 2010) have the potential to resolve several troubling paradoxes in tinnitus research (Sedley et al., 2016a), the tinnitus community at large has been slow to adopt them. Given that this framework has been applied elsewhere to studies of both normative and disordered brain function, this is somewhat surprising. The few tinnitus models based on it are nevertheless promising, with recent studies offering empirical evidence (Hullfish et al., 2019, 2018a, 2018b).

Here, we build a case for integrating predictive coding more fully into tinnitus research and vice versa. We therefore proceed from the perspective of tinnitus, exploring how predictive coding can improve our understanding of various aspects of this perceptual phenomenon. At the same time, we conversely explore the broader applicability of tinnitus research to the study of brain function (and dysfunction) by using predictive coding as a bridge. To this end, we incorporate selected research from outside the tinnitus field to discuss parallels that suggest fruitful avenues for future interdisciplinary research. Our review

* Corresponding author at: Lab for Clinical & Integrative Neuroscience, School of Behavioral & Brain Sciences, University of Texas at Dallas, 800 W. Campbell Rd., Richardson Dallas, Texas, 75080, USA.
E-mail address: Jeffrey.Hullfish@utdallas.edu (J. Hullfish).
[1] website: http://www.lab-clint.org.

https://doi.org/10.1016/j.neubiorev.2019.04.008
Received 28 January 2019; Accepted 14 April 2019
Available online 15 April 2019
0149-7634/

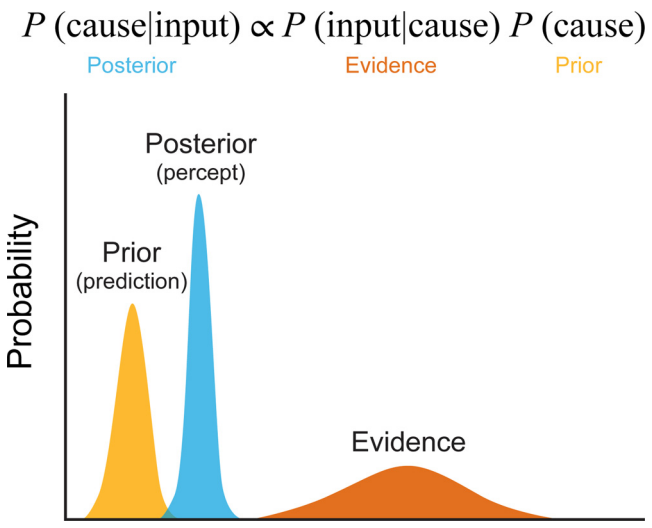

Fig. 1. Bayesian inference. (top) The posterior belief (probability of the cause, given the input) is directly proportional to the product of the evidence (probability of the input, given the cause) and the prior belief (probability of the cause). The constant of proportionality is the probability of the input itself. (bottom) Illustration of Bayesian inference using probability density functions. Content in this figure was adapted from (Sedley et al., 2016a), which was made freely available under the Creative Commons Attribution License (CC BY) (https://creativecommons.org/licenses/by/4.0/).

features a deeper and expanded discussion of a recently proposed predictive coding framework for tinnitus (Sedley et al., 2016a), considers its applicability to a broader range of tinnitus phenomena, introduces new roles for dopaminergic frontostriatal circuits, and lays out testable, refutable hypotheses for future mechanistic and therapeutic research studies.

2. A précis of predictive coding

Neural implementation of Bayesian inference is most often framed in terms of predictive coding (Clark, 2013; Friston, 2018; Rao and Ballard, 1999). It was originally conceived as a data compression strategy that enabled a system to efficiently process complex information by predicting its content and only encoding the values of prediction errors, as opposed to the information itself (Clark, 2013). Predictive coding thus becomes increasingly efficient as predictions improve. The brain's implementation of this strategy relies on a hierarchical generative model where bottom-up sensations are compared to top-down prior beliefs across multiple levels. Each level uses these bottom-up inputs as evidence with which to update its beliefs via Bayesian inference. Each level is also highly interdependent, such that the belief at one level serves as the prior for the level below. The mismatch between evidence and priors, i.e. the prediction error, at that subordinate level then ascends and updates the original belief, at which point the process repeats. This empirical Bayesian approach (Efron and Morris, 1973) enables the brain to estimate priors from data without needing to know them intrinsically, i.e. unsupervised learning; see Fig. 2 and especially (Friston, 2003) for a detailed introduction.

Per Helmholtz (1867), perception is a process of unconscious inference about the causes of sensations. Predictive coding elaborates on this idea by treating perception as a hierarchy of causal inferences, each one constraining the level below (Friston, 2003). This hierarchical generative model—the aptly named Helmholtz machine (Dayan et al., 1995)—thus comprises the brain's internal model of the sensorium. It is important to remember, however, that the brain infers the causes of sensations because the causes themselves are latent (i.e. hidden) variables and thus cannot be directly observed (Vilares and Kording, 2011); the brain must actively infer those latent causes given its sensations and prior beliefs. One implication of this is that inferred causes need not match "true" causes; this will be an important point for when we later describe tinnitus under predictive coding (Section 3).

Because the causes of sensations cannot be directly observed, the brain must instead rely on minimizing prediction error as the major criterion for optimizing its generative model. This means that the brain's ability to recognize which prediction errors carry reliable information about changes in the environment is critical. This distinction between signal and noise is made based on precision. Under predictive coding, beliefs are represented using Gaussian probability distributions, or densities, over a given perceptual dimension such as stimulus intensity (Sedley et al., 2016a). Precision is defined as the inverse variance of such a density; the expectation, or most likely belief, is defined as its mode. A belief is thus defined as both the prediction itself (expectation) as well as the level of confidence in said prediction (precision). Evidence is represented using a likelihood—i.e. the conditional probability of observing the sensory input given the prior belief—and has an expectation and precision of its own. Prediction error then becomes the precision-weighted difference between the prior belief and the evidence; the error itself thus also has an expectation and precision. Recent empirical evidence supports this model, finding distinct neural correlates for prediction errors (weighted by the precision of the prior prediction), updating of predictions, and the precision of predictions (Sedley et al., 2016b). The idea in brief is that each factor (prior belief and evidence) influences inference in proportion to its precision, meaning that only sufficiently precise prediction errors are allowed to significantly alter the generative model and thus affect perception.

Broadly speaking, we can describe changes to the generative model in terms of changes to the following parameters: (1) sensory expectation, (2) sensory precision, (3) prior expectation, and (4) prior precision. The process of updating the generative model may be understood through a framework of perception such as hypothesis testing (Gregory, 1980). Consider first the generative model's default prediction. Conceptually, this default prediction may be considered a null hypothesis (or null percept) that the brain compares with sensory evidence (Sedley et al., 2016a). This null percept does not represent a lack of neural activity per se but rather the pattern of imprecise prediction errors that reflect spontaneous activity. In other words, given a normative generative model, silence is the inferred cause of spontaneous activity in the auditory system (and darkness is the inferred cause of spontaneous activity in the visual system, etc.). Rejecting the null corresponds to updating the model and thus to perception of 'something' over 'nothing.' This occurs whenever the evidence changes such that its expectation and/or precision are high enough to deviate significantly from the null. For example, an acoustic stimulus (e.g. a pure tone) presented at a level above a subject's hearing threshold would elicit a sufficient change in evidence to update their generative model and be perceived. Tinnitus, on the other hand, is by definition not caused by an acoustic stimulus. Sedley et al. (2016a) instead argue that tinnitus is almost always caused by factors that increase the precision of spontaneous activity in the auditory system.

3. Tinnitus in the predictive coding framework

Tinnitus, along with several other signs and symptoms of neuropsychiatric disorders, can be explained under predictive coding as a failure to attenuate sensory precision (Edwards et al., 2012; Friston et al., 2012; Pezzulo et al., 2015; Sedley et al., 2016a). This may be generally understood as a failure to distinguish between prediction errors that contain useful information and those that do not. As mentioned previously (Sedley et al., 2016a), the prediction errors that lead to tinnitus arise from spontaneous activity. Normally, spontaneous prediction errors in sensory systems are imprecise and thus do not significantly deviate from a null percept, e.g. silence. If one or more factors causes precision to increase, however, then spontaneous activity can elicit a prediction error capable of updating the generative model

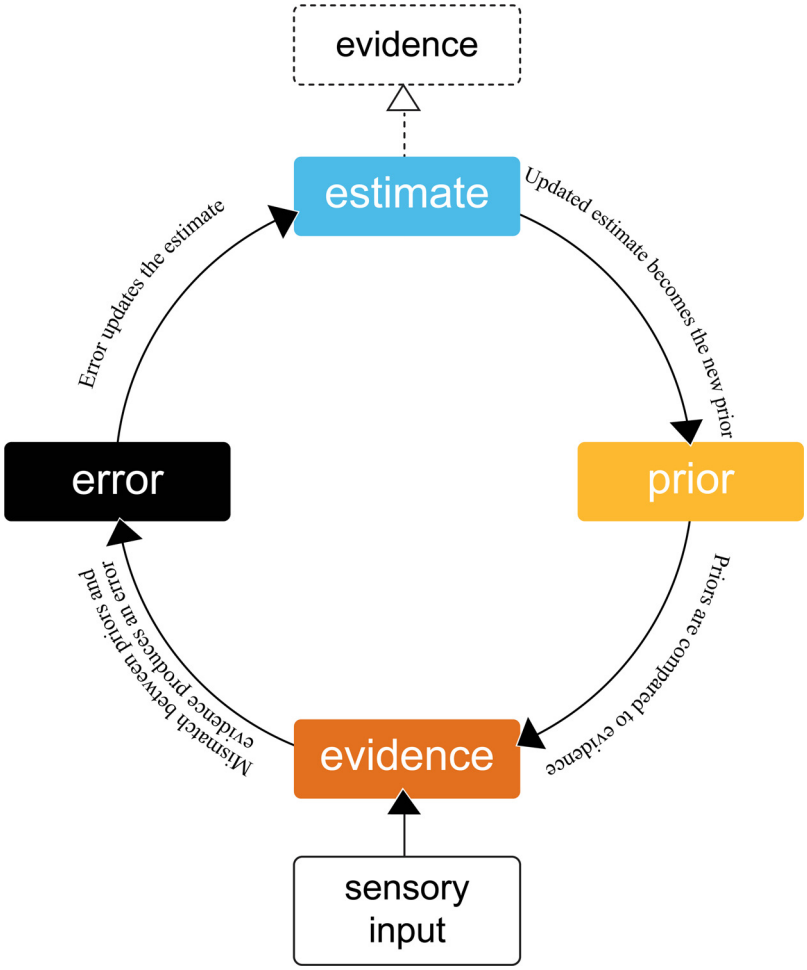

Fig. 2. Diagram of hierarchical predictive coding *via* empirical Bayesian inference. The estimate serves as an empirical prior belief. The evidence is dependent on this empirical prior (top-down input) and well as sensory input (bottom-up input). Comparing the evidence and the prior produces a mismatch signal, i.e. the prediction error, which is used to update the original estimate. This updated estimate becomes the new empirical prior, and the process repeats until the prediction error is minimized. Extending this process over multiple hierarchical levels, the estimate at one level becomes the bottom-up input for the level above (dotted lines).

and therefore will be perceived, i.e. as tinnitus (Fig. 3). These factors encompass many mechanisms of interest in tinnitus, including synaptic gain, subcortical gating, cross-modal connections, neurochemical changes, and both the extent (e.g. tonotopic map plasticity) and synchrony of coherently active neuronal populations. Thus, the first of the two central tenets of this model is that the onset of tinnitus reflects an increase in sensory precision in related frequency channels. In subsequent sections of the manuscript, where we refer to changes in sensory precision at various stages in the emergence and chronification of tinnitus, we refer to any one or combination of these contributory processes to precision.

Research has shown that several bottom-up and top-down factors, either alone or in combination, can lead to tinnitus. The primary bottom-up factor predisposing one to tinnitus would be hearing loss; see especially Section 4 for discussion on this topic. When one acquires hearing loss, their brains expect but no longer receive auditory input. This expectation arises from the fact that sensory systems are always active. Predictive regulation dictates that sensory cells will adjust their sensitivity to match the expected range of inputs (Sterling, 2012; Sterling and Laughlin, 2015). This implies that sensors deprived of input will increase their sensitivity in an attempt to detect missing input, which in turn increases the intensity of spontaneous activity. Ideally, the brain would attenuate sensory precision in these cells and prevent it from updating the generative model. This adaptive reduction in sensory precision following hearing loss reflects the fact that the deafferented cells no longer signal useful information about the auditory environment—or, in the case of partially deafferented cells, contain less reliable information—and thus prevents tinnitus from emerging; see Section 4.1 for more examples of hearing loss without tinnitus.

If the brain fails to attenuate this sensory precision, however, then spontaneous prediction errors do update the generative model and thus do lead to tinnitus. The important thing to note is that the brain is reducing prediction error in both cases. In the adaptive case, i.e. no tinnitus, the brain attenuates sensory precision and thus prevents a (precise) prediction error from arising in the first place. In the maladaptive case, i.e. tinnitus, the brain updates its generative model to explain away the prediction error that it should have attenuated; in other words, it accepts the spontaneous activity as a percept, or deviation from the null hypothesis.

In predictive coding, perception is the weighted average of the sensory signal and the generative model, thus tinnitus can occur when both of these represent an auditory entity. Both of these cases involve reduction of prediction error and therefore can be considered 'adaptive' at some level, the difference being that what we refer to as the maladaptive case reduces error at the lower level of auditory signals (by predicting spontaneous activity), but increases it (at least initially) at the higher level of behaviorally relevant environmental causes (because updating the generative model to encompass a brand new cause entails an increase in free energy; Friston, 2010). This dichotomy potentially provides an alternative or additional explanation as to why only some

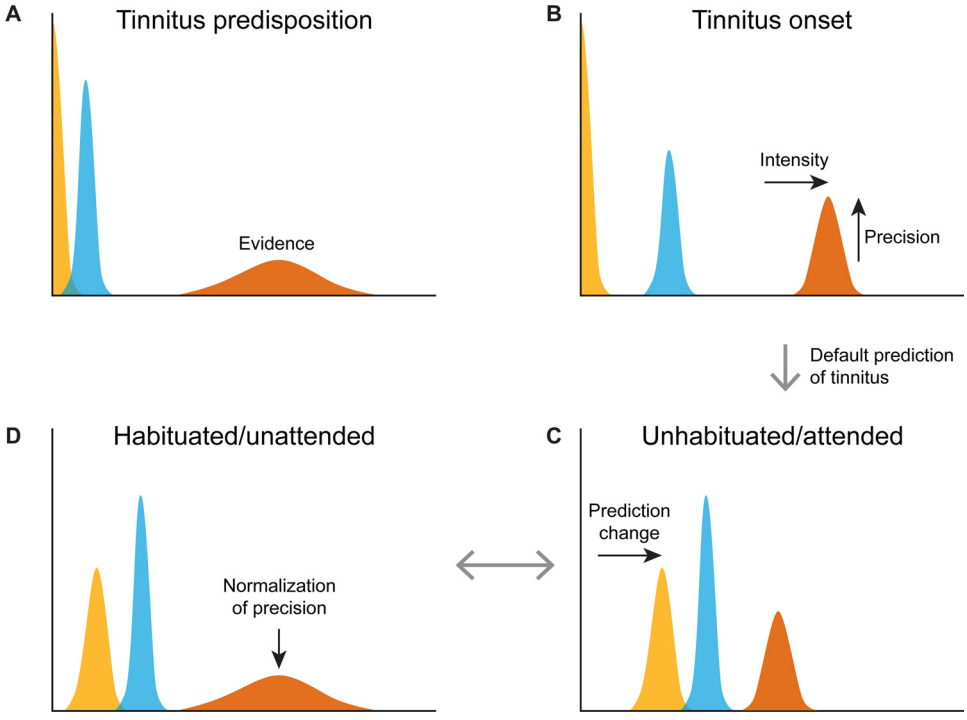

Fig. 3. Illustration of predictive coding in tinnitus. In each plot, the perceived loudness of tinnitus is indicated by the position of the posterior distribution (blue) on the horizontal axis. (A) In hearing loss alone, the sensory evidence for tinnitus, or the tinnitus precursor (orange), has insufficient precision to override the default prediction (yellow) of silence. (B) With increased precision, the precursor influences perception, leading to a revised posterior percept of tinnitus. Potentially, there is a window of reversibility at this stage. (C) If the default prediction is revised to expect tinnitus (generally less intense than the precursor), then the condition becomes chronic (through experience-dependent plasticity). (D) Reduction of the precision of the precursor to its pre-tinnitus level results in habituation, but not cessation of tinnitus—on account of plastic changes to prior predictions. Content in this figure was adapted from (Sedley et al., 2016a), which was made freely available under the Creative Commons Attribution License (CC BY) (https://creativecommons.org/licenses/by/4.0/).

individuals with hearing loss develop tinnitus. The existing explanation in the model is that various factors synergize to increase precision of the tinnitus precursor signal, i.e. the sensory evidence for tinnitus, and crossing a certain threshold results in tinnitus perception. The additional explanation is that differences in how the brain handles the increased precision-weighting (top-down precision attenuation vs. updating the generative model) may determine whether tinnitus occurs, and these differences could encompass both traits and states of the individual's perceptual system. Exactly what these differences are is presently unknown, and could be an important area for future study. Possibilities could include traits relating to propensity for new learning, and states of vigilance for new sounds (for instance, we have encountered multiple patients whose tinnitus first occurred during a pure tone audiogram).

This updating corresponds to the brain learning (i.e. inferring) a causal relationship between spontaneous activity (input) and auditory perception (cause) (Friston, 2003; Hullfish et al., 2018a); see also (Mohan and Vanneste, 2017). Thus, the second of the two central tenets of the model is that perception of tinnitus requires updating its generative model to assign meaning to the underlying spontaneous activity, and the transition from acute to chronic tinnitus occurs when this generative model becomes a default prediction of the auditory environment. Note that simply tolerating the prediction error, i.e. without attenuating its precision or updating the generative model, is not an option because it represents a state of disequilibrium with the environment. In other words, minimizing prediction error is akin to maintaining homeostasis (Friston, 2010).

Active attenuation of sensory precision implies a top-down modulatory effect. This is consistent with existing models of predictive coding. Canonical microcircuits for predictive coding use feedforward (i.e. bottom-up) connections to convey prediction errors while feedback (i.e. top-down) connections convey predictions (Bastos et al., 2012). The top-down connections are inhibitory, such that predicted signals are suppressed and only new information, i.e. prediction error, is transmitted to the next level of the hierarchy. Importantly, these top-down connections exert both driving and modulatory effects (Bastos et al., 2012). One therefore might intuitively consider top-down connections to function as a noise cancelation system under predictive coding. If this noise cancelation system fails, then bottom-up noise will be able to update the generative model. This is consistent with tinnitus models based on aberrations in frontostriatal gating circuits (Hullfish et al., 2019; Rauschecker et al., 2015; Sedley et al., 2016a; Song et al., 2015). We discuss this in more depth in Section 5.

Top-down connections can also actively increase sensory precision, as opposed to simply failing to attenuate it. Precision in this sense is often construed in terms of attention or salience, and not just in the auditory domain (e.g. vision; see Hesselmann et al., 2010). Here, focused (i.e. endogenous, top-down) attention has a modulatory effect on sensory precision (Roberts et al., 2013). The corresponding increase in sensory precision causes spontaneous auditory activity to deviate significantly from the default prediction (the null percept, i.e. silence), a change which the brain infers is caused by auditory input and is thus perceived. The role of focused attention can potentially explain several phenomena, including the habituation of tinnitus as well as the emergence of tinnitus during stress (Sedley et al., 2016a).

All of the above factors and more can predispose one to tinnitus. However, there are systems in place that can still prevent the tinnitus percept from emerging if one system fails. This implies that the individual factors predisposing to tinnitus are less important than the overall balance between top-down and bottom-up systems. For example, if bottom-up losses lead to increased sensory precision, top-down connections can still attenuate that precision and prevent tinnitus. Failing to attenuate sensory precision on the other hand represents an imbalance between top-down and bottom-up connections, i.e. a prediction error. Therefore, what the predictive coding framework establishes is that tinnitus is ultimately the result of a prediction error that arises out of a failure to attenuate sensory precision. Importantly, this applies equally well in the context of many different "subtypes" of tinnitus. The idea here is that tinnitus has a many-to-one mapping—meaning that several different causes can lead to the same outcome, i.e. tinnitus—but predictive coding offers a framework sufficient to explain these various causes. We put this claim to the test in the following sections by exploring several different aspects of tinnitus in the context of predictive coding.

4. The relationship between tinnitus and hearing loss

Tinnitus falls along a broad spectrum of hallucinatory experiences, one that ranges from illusions on one end (e.g. Necker cube, McGurk effect, etc.) all the way to "true" hallucinations, which typically accompany psychosis, at the other (El-Mallakh and Walker, 2010; Wearne and Genetti, 2015). Tinnitus, which occurs in the absence of a corresponding external stimulus and presents without psychosis (but see Linszen et al., 2016), falls somewhere in the middle of this spectrum. It is most commonly considered an auditory phantom percept, akin to phantom limb, but researchers attempting to create a more principled classification call it a pseudohallucination or, more specifically, a parahallucination (El-Mallakh and Walker, 2010; Wearne and Genetti, 2015). Parahallucinations are distinct from other hallucinatory experiences in that they follow some sort of damage to or dysfunction of the peripheral nervous system. Phantom limb, which occurs in 90–98% of limb amputees, is a straightforward example of this; in 70% of cases it also leads to phantom pain, another parahallucination (Ramachandran and Hirstein, 1998). Tinnitus itself most often accompanies hearing loss, especially acquired losses such as those related to age (presbycusis) or overexposure to noise. While these are only three examples of parahallucinations, note that they can occur in any sensory domain (De Ridder et al., 2014c; Mohan and Vanneste, 2017).

There are, however, some instances of tinnitus that may not fit these criteria. For example, 64% of reportedly normal-hearing individuals experienced "tinnitus-like sounds" while partially sense-deprived via 20 min of isolation in a sound booth (Tucker et al., 2005). There are also cases of tinnitus in reportedly normal-hearing individuals—fewer than 20% of total cases (Elgoyhen et al., 2015)—that occur even without experimental manipulation. We review studies on the relationship between tinnitus and hearing loss to explore these and other cases, in particular showing how they all fit within the larger framework of predictive coding.

4.1. Hearing loss without tinnitus

4.1.1. Congenital vs. acquired hearing loss

While more than 80% of tinnitus cases accompany hearing loss (Elgoyhen et al., 2015), the converse is not true; most cases of hearing loss present without tinnitus. One of the factors that influences whether a person with hearing loss is likely to develop tinnitus is their age during its onset. Eggermont and Kral (2016) observed, for example, that cats with congenital single-sided deafness did not have tinnitus. They hypothesized that the absence of auditory experience in that ear could explain the absence of a corresponding auditory phantom.

In a recent study of single-sided deafness in humans, we found evidence that auditory experience is indeed a prerequisite for auditory phantom perception (Lee et al., 2017). The idea is that because subjects with congenital single-sided deafness never learn a causal relationship between auditory perception and neural activity in the deafferented sensory pathway, they never develop the priors necessary to infer the tinnitus percept from activity in that pathway. The same is not true, however, for those with acquired single-sided deafness. Because they had sufficient auditory experience involving the affected ear before losing their hearing, they do have those perceptual priors and thus do have a basis to infer tinnitus. Notably, these observations in tinnitus apply equally well to complex auditory phantoms such as musical hallucinosis (Kumar et al., 2014; Vanneste et al., 2013) as well as to phantoms in other sensory domains, e.g. Charles Bonnet syndrome; see (Mohan and Vanneste, 2017) for a detailed review.

What is not yet clear is how exactly tinnitus risk changes with the age of onset of hearing loss. A simplistic explanation would be that there is a binary difference between congenital losses, which do not cause tinnitus, and acquired losses, which sometimes do. There is however evidence suggesting that there may actually be a spectrum of risk for developing phantoms with increasing age. Specifically, cross-modal plasticity appears to be the primary factor mediating this spectrum. Cross-modal plasticity refers to two mechanisms (Rabinowitch and Bai, 2016): (1) recruitment, where intact sensory systems from other modalities take over cortex from the deprived system to boost their own performance; and (2) compensation, where changes take place in the intact sensory systems themselves to boost the performance of the deprived system. In a predictive coding context, we might interpret cross-modal plasticity as the brain optimizing its generative model such that deprived sensory systems receive useful information as evidence, whether from the same modality (compensation) or from another (recruitment). Importantly, the extent to which the brain engages in cross-modal plasticity changes with age. While there is evidence suggesting the occurrence of cross-modal plasticity in late-onset losses, e.g. blindness (Kujala et al., 1997), it is a much more common phenomenon with congenital and early-onset losses. The putative spectrum of risk with age appears to be determined by a critical period, which is generally agreed to end by early to mid-adolescence (Cohen et al., 1999; Hensch, 2005; Sadato et al., 2002).

A critical period for cross-modal plasticity is entirely consistent with the earlier explanation for why congenitally single-sided deaf patients do not develop tinnitus (Lee et al., 2017). Without a prior expectation of auditory input in the congenitally deprived cortex, the brain would be more likely, for example, to use cross-modal recruitment to ensure that those parts of cortex receive input from elsewhere. That change in evidence precludes the onset of tinnitus because (1) the relevant cortex develops different priors and/or (2) because the evidence for tinnitus—i.e. precise, spontaneous prediction error in the subcortical auditory system (Sedley et al., 2016a)—is no longer used as input to the model. The spectrum of risk in acquired hearing loss thus appears to depend on what constitutes sufficient auditory experience to have developed strong auditory priors, which would presumably prevent cross-modal plasticity. Given that priors are learned empirically, i.e. unsupervised learning (Efron and Morris, 1973; Friston, 2003), it stands to reason that they would take some amount of time to develop. We speculate that this initial learning period for perceptual priors coincides with the critical period for cross-modal plasticity. In other words, as priors develop, the brain becomes increasingly less likely to search for different evidence. This age dimension may have further relevance for the cognitive and emotional aspects of tinnitus, as tinnitus-related distress also appears to depend on the age of tinnitus onset (Schlee et al., 2011).

4.1.2. Dreams

Dreams offer an interesting edge case in tinnitus research. While many people with hearing loss have tinnitus, at least one study indicates that they do not experience it while dreaming (De Ridder et al., 2014a). This is consistent with findings from phantom pain research, with amputees even reporting an intact body schema while dreaming, i.e. all limbs intact and fully functional with no phantom sensations (Alessandria et al., 2011; Mulder et al., 2008). Indeed, hallucinatory experiences in general appear to occur only during the waking state (Waters et al., 2016). This begs the question of how to explain the hallucination-free dream state via predictive coding.

To this end, consider the fact that dreams are entirely detached from ongoing sensory feedback (De Ridder et al., 2014a; Hobson and Friston, 2012). While dreams are constructed from information in the same sensory domains as waking perceptions, they are not constrained by the concurrent state of the external world. Since sensory systems are nevertheless active at all times, the implication is that the sleeping brain strongly attenuates sensory precision in those systems, thus reducing prediction errors. Note however that "strongly attenuate" does not mean "completely attenuate," as external stimuli, e.g. loud noises, can still rouse the brain from sleep. Hence, if the sensory evidence for tinnitus is attenuated while sleeping, the tinnitus percept becomes a much less likely inference.

This explanation is not sufficient to explain why dreams are

tinnitus-free, however. If we take as a given that tinnitus is the product of inference—which we certainly do under predictive coding—then we must also consider the potential impact of priors. For instance, chronic tinnitus reflects a change in the prior belief, i.e. away from the default prediction of silence; see Section 5 for further discussion. With a prior belief of tinnitus, the percept can arise even if the evidence for it returns to normative levels (Sedley et al., 2016a; Fig. 3). The necessary insight here is that priors are constructed from memory and thus depend on context. To paraphrase Denève and Jardri (2016), the context of being in a forest increases the prior probability of seeing leaves, which is validated in turn by sensory inputs of light with wavelengths in the range of 495–570 nm, i.e. corresponding to the inferred color 'green.' The context of being in a clothing store, however, comes with a higher probability that green objects are articles of clothing, as opposed to leaves. As an interesting aside, priors therefore offer an explanation for camouflage: sensory evidence that closely resembles the environment—e.g. forest-patterned clothing seen in a forest—is more likely to be inferred as part of that environment than as an article of clothing because of the context and its associated priors. Returning to dreams, however, the context is that evidence from the external world has inherently low precision. This means that there is a prior expectation that the external world offers little information while sleeping, and so the generative model has learned to attenuate signals from the external world during sleep (Hobson and Friston, 2012). If the model attenuates evidence for tinnitus during sleep, then it is unlikely that the brain will ever infer tinnitus during sleep. This precludes the possibility of the brain updating its priors in that context, since the attenuated prediction errors have too low precision to deviate significantly from those priors. Although this process of updating of priors has occurred, for the individual, in the awake context, we suggest that this updating is context-specific and therefore is not sufficient to affect perceptual inference in the dream state to cause tinnitus. Predictive coding therefore offers an explanation from the perspective of both evidence and prior beliefs for why tinnitus is not a feature of dreaming consciousness.

4.2. Tinnitus without hearing loss?

4.2.1. Experimental manipulation

There are a number of ways in which experimental manipulations can reversibly induce tinnitus (or "tinnitus-like sounds") in healthy, normal-hearing subjects. These generally involve some form of auditory deprivation, such as periods of isolation in a sound booth (Tucker et al., 2005) or an anechoic chamber (Del Bo et al., 2008). These methods, especially the anechoic chamber, simulate profound bilateral hearing loss across the entire audible spectrum, which is not representative of the tinnitus population as a whole. However, an experiment using unilateral earplugs continuously over seven days simulated a more representative hearing loss and also induced reversible auditory phantoms in the majority of participants (Schaette et al., 2012). These experiments establish a causal relationship between auditory deprivation, as opposed to auditory deafferentation per se, and phantom sounds (Schaette et al., 2012). This is consistent with research stating that sensory cells will update their sensitivity to match the expected range of inputs (Sterling, 2012; Sterling and Laughlin, 2015). Explained in terms of predictive coding, these experimentally induced sounds are the result of an increase in sensory expectation (through compensatory gain changes) and/or sensory precision (due to focused attention) of spontaneous activity in response to a reduction of auditory input. In other words, these sounds are the product of a mechanism similar to that of deafferentation, the main differences being that the source of the loss is external and that the loss itself is reversible.

An example of experimental induction of tinnitus-like sounds without auditory deprivation is the Zwicker tone illusion (Franosch et al., 2003; Zwicker, 1979, 1964). A Zwicker tone is a negative auditory afterimage, meaning that it appears immediately after presenting a specific sound and then disappears within seconds. The sounds typically used to elicit Zwicker tones are notched noise stimuli, i.e. broadband noise with a narrow band-stop filter; the Zwicker tone itself has a pitch corresponding to one of the suppressed frequencies. We speculate that Zwicker tones arise after notched-noise stimuli in a manner similar to how tinnitus arises following hearing loss under predictive coding. Recall that hearing loss induces a specific range of auditory deprivation, which defies the prior expectation of auditory input. By upregulating the sensitivity of the deprived neurons in search of missing input, spontaneous activity is also increased. When this activity increases enough to deviate from the prior prediction of silence, the brain infers the tinnitus percept as the cause. In the case of the Zwicker tone, the notched-noise stimulus used to induce it is not something that commonly occurs in the environment. Broadband noise, however, is a similar stimulus that is more common. We hypothesize that when the brain is presented with notched noise, it has a prior prediction that broadband noise is the cause. This mismatch between notched and broadband noise means that the brain expects, but does not receive, auditory input at the cells responsible for transducing frequencies in the range of the notch; in other words, notched noise induces a prediction error. Searching for that expected input, the brain upregulates the sensitivity of those cells in the same way that it does following hearing loss. When the notched noise stops, that sensitivity remains upregulated. While this is the case, the relatively high precision of spontaneous activity leads to an inference of a Zwicker tone. Since input is not actually missing, however, the sensitivity of the affected cells quickly returns to match the normal range of auditory input levels, at which point the Zwicker tone disappears.

4.2.2. Hidden hearing loss

While experimental manipulation can induce auditory phantoms in normal hearing subjects, tinnitus can also occur outside of such controlled settings. For instance, as many as one in five tinnitus cases present with normal pure-tone hearing thresholds (Elgoyhen et al., 2015). This is often taken to mean that tinnitus can arise without peripheral hearing loss. However, a normal pure-tone audiogram does not preclude cochlear damage. Schaette and McAlpine (2011) demonstrate, for example, that supposedly normal-hearing tinnitus subjects have a deficit in auditory nerve function, which indicates deafferentation of high-threshold fibers. While this may not affect pure-tone hearing thresholds, it does affect the processing of sounds more complex than pure tones, such as speech, especially in noisy environments (Liberman, 2015). Schaette and McAlpine (2011) argue that this is compensated for in the auditory brainstem, consistent with the view that homeostatic mechanisms acting to normalize neural activity levels in the central auditory system promote tinnitus, i.e. the central gain theory. This compensatory mechanism increases spontaneous firing rates in the deafferented auditory pathway (Mohan and Vanneste, 2017), which in predictive coding terms corresponds to an increase in the sensory expectation (Sedley et al., 2016a). The posterior expectation, i.e. the most likely inferred cause, also increases as a result. The inference in this case would be of perceptual intensity, specifically tinnitus loudness (Sedley et al., 2016a; Fig. 3). This means that many cases of tinnitus presenting with normal pure-tone audiograms may nevertheless be the result of (hidden) hearing loss. Tinnitus without any degree of hearing loss at all thus increasingly appears to be a rarity.

4.2.3. Somatic tinnitus

Somatic (craniocervical) tinnitus is a term coined by Levine (1999) to describe patients with normal audiograms—both pure-tone and speech, the latter making hidden hearing loss less likely—that have several clinical features in common, including: (1) somatic disorder of the head and/or upper neck; (2) unilateral tinnitus localized to the ear ipsilateral to the somatic disorder; (3) no vestibular complaints; and (4) no neurological abnormalities on examination. He concluded that there was a physiological mechanism for somatosensory interaction with the dorsal cochlear nucleus leading to the onset of tinnitus in these patients

(Levine, 1999); see also (Shore et al., 2008; Young et al., 1995). In another study, Levine et al. (2003) found that 80% of non-clinical tinnitus subjects who perceived their tinnitus at the time of testing could modulate their tinnitus via head and neck contractions. Furthermore, 60% of those who did not perceive their tinnitus at the time of testing could elicit an auditory phantom via those same head and neck contractions. Similar results were also reported in profoundly deaf subjects. They ultimately concluded that the central auditory and somatosensory systems interact and that somatosensory inputs to the auditory system are sufficient for tinnitus-like auditory perception (Levine et al., 2003).

This finding is interesting in the context of our present framework because it demonstrates that, while the auditory system (and auditory priors; recall Section 4.1.1) may be necessary to infer tinnitus, the evidence used to make that inference need not originate there. This harkens back to Helmholtz's insight that the brain does not innately know the causes of sensory inputs (Helmholtz, 1867). In other words, if an unfamiliar pattern of somatosensory input to a node of the auditory system—in this case, the dorsal cochlear nucleus (Levine, 1999)—resembles a familiar pattern of auditory input to that same node, then the brain would naturally infer that both inputs have the same cause, i.e. auditory perception.

Per Sedley et al. (2016a), tinnitus is the inferred cause of spontaneous activity in the auditory system that has a sufficiently high expectation and/or precision to update the generative model (Fig. 3). Spontaneous activity, however, is not a unique feature of auditory neurons. If such activity in the somatosensory system reached the auditory system, then the resulting inference would be the same as if it had originated from within the auditory system itself: tinnitus, if the expectation/precision were sufficiently high, or silence if not. Indeed, somatosensory inputs to the DCN can increase the firing rates of auditory neurons (Shore, 2005), which corresponds to an increase in sensory expectation. Somatic tinnitus may therefore be the exception that proves the rule of parahallucinations. On one hand, it is an example of tinnitus that appears to truly not require peripheral auditory damage (cf. hidden hearing loss; Section 4.2.2). Nevertheless, there is still damage to the peripheral nervous system, i.e. in the somatosensory domain, which is sufficient to satisfy the criterion for parahallucinations (El-Mallakh and Walker, 2010; Wearne and Genetti, 2015). It is however also likely that the co-occurrence of hearing loss and somatic disorders is more likely to result in tinnitus than the presence of either factor alone. Indeed, animal evidence indicates that one of the functions of somatosensory input to the auditory system is to suppress responses to expected sounds and to conversely enhance responses to unexpected sounds (Shore and Zhou, 2006). It therefore stands to reason that somatosensory inputs could further amplify the spontaneous prediction errors that already have an increased expectation/precision as a result of hearing loss, which would suggest that somatic factors play a role in tinnitus above and beyond what is currently acknowledged.

5. From acute to chronic tinnitus via learning

Under predictive coding, learning and inference are equivalent (Friston, 2003). This is relevant here because there is an emerging consensus that the transition from acute to chronic phantom perception is a learning process, albeit a maladaptive one. This is perhaps best established in pain research. To paraphrase Apkarian et al. (2011), given that chronic pain is a state of continuous pain perception, it must also be a state of continuous acquisition of associations between environmental events and pain perception. In other words, the perception of pain in different contexts leads to associating pain with those contexts. Extinguishing these associations inversely requires experiencing those same contexts repeatedly without pain. While pain perception persists, however, the brain never gets this opportunity and so acute pain perception eventually becomes chronic. Their model for the transition from acute to chronic pain (see Fig. 14 of Apkarian et al., 2011) defines the primary driver of that transition as pain's ability to induce long-term memories via reinforcement learning. Note that it may be possible to model this transition with active inference exclusively, i.e. without recourse to reinforcement learning (Friston et al., 2009), but this is not a matter of consensus; in either case, researchers have often used reinforcement learning in tandem with Bayesian inference to model brain function, demonstrating that they are compatible frameworks (see e.g. Donoso et al., 2014; Koechlin, 2016; Mathys, 2011; Payzan-LeNestour et al., 2013; Payzan-LeNestour and Bossaerts, 2011).

The Apkarian et al. (2011) model for chronic pain is sufficiently general such that it can apply equally well to other phantoms. This is especially true for tinnitus, which has well-established parallels with pain (De Ridder et al., 2011; Rauschecker et al., 2015; Tonndorf, 1987). One paper proposes a model of both disorders as persisting, aversive memory networks (De Ridder et al., 2011). This model emphasizes the importance of the salience network (Seeley et al., 2007), which among other things influences the sensory–limbic interactions necessary for the continuous experience of phantom perception. These interactions are reportedly mediated by subgenual anterior cingulate cortex–nucleus accumbens (i.e. a frontostriatal circuit) and the amygdala, which in turn modulate the thalamic reticular nucleus (Rauschecker et al., 2015). Memory mechanisms involving the para-/hippocampus play a role in the persistence of both phantom perception as well as the phantom–distress association (Husain and Schmidt, 2014; Sedley et al., 2015; Vanneste and De Ridder, 2016, 2011); we speculate that these areas are responsible for encoding and relaying auditory priors, e.g. the altered prior belief of tinnitus. The phantom–distress association is reflected in a general distress network comprising the parahippocampus, the salience network, and the amygdala (De Ridder et al., 2011). A key insight of the De Ridder et al. (2011) model is that multiple networks operating in parallel contribute to the pathologies of both pain and tinnitus; see also (De Ridder et al., 2014d, 2016). As implied earlier, this model is also broadly consistent with the pain model by Apkarian et al. (2011). A recent opinion paper corroborates both models, proposing that a dysfunctional frontostriatal network plays a role in the perceptual and emotional aspects of both tinnitus and chronic pain (Rauschecker et al., 2015).

These models are not only consistent with each other but also with subsequent empirical studies. Researchers behind the model in (Apkarian et al., 2011), for example, ran a longitudinal study of subacute back pain patients and found that greater functional connectivity in frontostriatal circuits (i.e. nucleus accumbens with prefrontal cortex, incl. anterior cingulate cortex), which are important for learning, predicted whether patients would transition to a state of chronic pain (Baliki et al., 2012); see also (Baliki et al., 2010, 2006; Hashmi et al., 2013; Mutso et al., 2012). Furthermore, a recent study using resting-state fMRI found that frontostriatal circuits, and nucleus accumbens in particular, exhibit increased functional connectivity in tinnitus subjects relative to healthy controls (Hullfish et al., 2019). That same study found that connectivity involving nucleus accumbens correlates with clinical measures of tinnitus; most notably, connectivity between nucleus accumbens and ventromedial prefrontal cortex correlated with the duration of tinnitus symptoms while connectivity between nucleus accumbens and parahippocampal cortex correlated with both tinnitus loudness and the duration of tinnitus symptoms (Hullfish et al., 2019). An earlier EEG study further establishes a link between anterior midcingulate cortex (a.k.a. "rostral anterior cingulate cortex"; Song et al., 2015) and dysfunctional noise cancelation in tinnitus patients, in that activity of this area inversely correlated with the proportion of the time subjects were consciously aware of their tinnitus. Multiple, independent lines of evidence thus appear to converge in favor of extended frontostriatal networks comprising a top-down noise cancelation system that is dysfunctional in at least two phantoms, namely tinnitus and chronic pain. Based on this, we hypothesize that future studies of phantoms in other domains (e.g. visual, olfactory, etc.) will uncover

similar findings, implying that the analogy between tinnitus and chronic pain extends to all phantoms.

Despite important questions remaining for future studies, there is at least a consensus that frontostriatal circuits play a role in learning. This comes not only from the disorder studies already discussed but also from studies of normative brain function (see e.g. Donoso et al., 2014; Koechlin, 2016; Mathys, 2011; Payzan-LeNestour et al., 2013; Payzan-LeNestour and Bossaerts, 2011). These studies tend to focus on learning in the context of decision-making and behavior, especially with regards to reward prediction error signaling in midbrain dopaminergic systems (incl. nucleus accumbens). While this is the canonical role of these systems, there are several studies making an interesting case for a re-appraisal of that role. We have mentioned one of these studies already (Friston et al., 2009), which claims that dopamine may not actually encode prediction errors of value (i.e. reward prediction errors) but rather the value of prediction errors; in other words, Friston et al. (2009) claim that dopamine encodes precision. Precise prediction errors, which measure surprise, are what drive both perception and action, i.e. active inference (Friston et al., 2009).

Recent studies using other frameworks have also proposed a more nuanced view of dopaminergic error signaling. One line of research in particular has shown that these dopaminergic neurons respond not only to changes in value, i.e. reward prediction errors, but also to value-neutral changes in the sensory features of a predicted reward, i.e. sensory prediction errors (Takahashi et al., 2017). They conclude that these neurons do respond to changes in the value of predicted rewards but that this role does not fully explain their functional profile; see also (Gardner et al., 2017), which discusses this concept of sensory prediction errors in more depth. The key takeaway—regardless of whether the reader prefers reinforcement learning, hierarchical Bayesian approaches (Figs. 1,2), or some combination of the two—is that dopaminergic signaling can drive value-neutral perceptual learning as well as reward learning. This is consistent with existing frameworks of dopamine's role in learning, e.g. nucleus accumbens confirming new task sets as the default (Donoso et al., 2014; Koechlin, 2016), and with the previously proposed role of frontostriatal circuits in the transition from acute to chronic phantom perception. By signaling sensory/perceptual prediction errors, dopaminergic circuits may facilitate the maladaptive updating of the generative model's default prediction from a null percept (silence, darkness, no pain, etc.) to a phantom percept, as predicted in the Sedley et al. (2016a) tinnitus model; see also (De Ridder et al., 2014c). In other words, by updating its default prediction (Fig. 3C), the brain is learning to perceive the phantom by default, i.e. chronically.

Our present discussion has converged on three roles for dopamine and frontostriatal circuits in tinnitus: (1) top-down noise cancelation, or more aptly the failure thereof, (2) salience, and (3) perceptual learning, i.e. updating the default prediction from 'no tinnitus' to 'tinnitus.' Not only are these potential roles not mutually exclusive, they indeed appear to be closely linked. The failure to attenuate sensory precision corresponds to the brain rating that evidence as salient by default, given the already discussed links between precision and attention (Section 3). Salient perceptual information, in turn, is information that the brain is more likely to incorporate into memory—i.e. to learn—for the sake of improving the accuracy of future predictions. In brief, if the frontostriatal network fails to suppress irrelevant sensory signals, they are able to enter conscious awareness, which then enables the resulting percepts to be learned (at a rate possibly determined by their salience/precision). It is no surprise, then, that the extended frontostriatal network that reportedly mediates top-down noise cancelation overlaps so conspicuously with both the salience network and regions involved in learning.

6. Testing predictive coding in tinnitus

As with any theory, our present account of tinnitus is only useful to the extent that it holds up to empirical testing. Predictive coding theories in particular are often met with skepticism on the grounds that they tend to be overly conceptual without making clearly articulated and falsifiable hypotheses. We offer the following as a means of avoiding this pitfall. Beyond the broad claim that predictive coding describes the mechanism of action in tinnitus, our account thus far offers two key hypotheses: (1) tinnitus reflects an increase in sensory precision in related frequency channels, at least in the acute phase, and (2) chronic tinnitus reflects a change in the brain's default prediction. We shall discuss each of these hypotheses in turn, describing not only concrete options for confirming them but also the evidence that could refute them. Finally, we shall offer some thoughts about testing the predictive coding model for tinnitus more generally.

6.1. Tinnitus reflects increased sensory precision

To reiterate our first hypothesis, tinnitus—at least in the acute phase—reflects an increase in sensory precision. Furthermore, this increase should occur specifically in tinnitus-related frequency channels, which necessitates methods that are capable of observing said channels. One way to test this claim, for instance, would be to look for evidence of increased cholinergic action on these tonotopic channels. Such an experiment would also need to include appropriate controls for attention since refocused attention can modulate sensory precision in a top-down fashion, as we have previously discussed. Note that it may be possible to assess basal forebrain cholinergic function via a weighted combination of EEG sensor space power and coherence metrics (Johannsson et al., 2015). This is a particularly promising avenue for tinnitus research, given that the literature has already proposed a role for the basal forebrain cholinergic system in the generation and modulation of tinnitus (Roberts et al., 2013).

Another option would be to conduct psychophysical experiments designed to look for evidence of increased weighting of these channels when performing perceptual inference operations. A study of evoked brain responses could offer a third option, provided said responses indicated precision in some way, e.g. precision-weighted prediction error (Sedley et al., 2016b). If the proposed studies found no differences in these measures between tinnitus subjects and matched controls, then this would stand as evidence against the increased sensory precision aspect of our model, at least in the phase of tinnitus tested.

It is important to note, however, that this is the more difficult claim to test. We hypothesize only that increased sensory precision is necessary in the acute phase; it may well return to normal later on. Testing would therefore ideally incorporate subjects in both the acute and chronic phases of tinnitus. Recent evidence however suggests that by the time tinnitus is present for four weeks, there is an approximately 90% persistence rate at six months (Wallhäusser-Franke et al., 2017). If we take this to mean that the acute phase of tinnitus is generally not longer than four weeks, then recruiting acute tinnitus subjects becomes problematic. The problem could be alleviated somewhat by first performing experiments to verify whether tinnitus-like sounds—such as those induced by notched-noise stimuli, i.e. Zwicker tones, or by partial auditory deprivation (recall Section 4.2.1)—constitute good neural models of acute tinnitus. While these experiments themselves would likely not be able to avoid the problem of recruiting acute tinnitus patients, they could form a body of evidence sufficient to justify inducing tinnitus-like sounds as proxies for acute tinnitus in future experiments.

6.2. Chronic tinnitus reflects a new default prediction

Our second hypothesis is that chronic tinnitus reflects a change in the brain's default prediction. By definition, this change should persist—and is, in fact, more likely to be seen—in chronic tinnitus. This is therefore the easier claim to test. Indeed, we already have evidence of persistent brain changes in chronic tinnitus, e.g. from comparing the evoked responses of tinnitus subjects to tinnitus-matched vs. control

sound stimuli (Hullfish et al., 2018a, 2018b) and from comparing the resting-state brain activity of tinnitus subjects to that of healthy controls (Hullfish et al., 2019). The challenge for future research lies in attributing these changes specifically to a change in the default prediction.

If chronic tinnitus does, in fact, reflect a change in the default prediction, then this change should be observable in tinnitus-related frequency channels. This should skew the predicted intensity—and perhaps also to some extent the frequency—of stimuli towards the predicted value. There are two possibilities here; either the default prediction has a lower expectation than the evidence for tinnitus (Sedley et al., 2016a) or it reaches the same intensity/expectation. In either case, stimuli of greater intensity than the tinnitus percept would necessarily exceed the intensity of the default prediction and would therefore evoke a prediction error. This is amenable to direct testing via evoked or induced responses that indicate prediction errors, e.g. mismatch negativity. Other indicators of prediction error could also be valid ways to test this hypothesis, including psychophysical measures or even eye tracking; see e.g. (Chait et al., 2012). Alternatively, a psychophysical paradigm designed to measure subjects' predictions themselves could also work. These experiments would refute our hypothesis if they either failed to demonstrate the proposed changes in predictions and/or prediction errors or if the effect of tinnitus was shown to be no different than the effect of presenting external, tinnitus-like stimuli to matched controls; the latter would indicate that the effect is merely the consequence of a competing low-intensity sound.

6.3. Testing the model in general

We have thus far discussed how to test the specific claims we make in our predictive coding account of tinnitus. The remaining question is whether predictive coding is the correct framework for tinnitus in the first place. Perhaps the most straightforward answer is to use Occam's razor; if another, simpler model such as central gain or neural synchrony can be demonstrated to explain tinnitus fully, then the predictive coding model becomes unnecessary and therefore unlikely. However, as already discussed at great length (Sedley et al., 2016a), extant tinnitus models have all failed to account for at least one aspect of tinnitus, if not several. A recent review further concludes that there is limited evidence linking neural firing rates and synchrony specifically to tinnitus—over and above hearing loss, to which they clearly do relate—and that present evidence does not support classification of individual subjects' tinnitus status based on any such metric (Sedley, 2019).

The predictive coding model is the first that can claim to resolve all of these aspects and unite them under a single explanatory framework, though some of these claims do still need to be tested, hence the present discussion. Refuting a predictive coding model for tinnitus outright would therefore require a straightforward neural metric that can explain the presence (or absence) of tinnitus. The present lack of such evidence is mainly due to a lack of appropriately controlled and performed studies to seek it, and so it might well turn out that some metric of neural synchrony (by far the better candidate than firing rate, or even bursting) does prove to be a sufficient basis for tinnitus. In this case, there would be no need for a predictive coding model of tinnitus. With that said, there is an emerging consensus that all perception does involve predictive coding; refuting the present account would therefore merely establish that there is nothing special about how predictive coding behaves in the case of tinnitus.

7. Implications for intervention

Let us assume for the moment that our present account of tinnitus is, in fact, the correct account. This begs the question of what it portends for the development of clinical interventions. First and foremost, our account implies that the time to intervene is in the acute phase of tinnitus. The increased sensory precision that is necessary to infer tinnitus in the first place is only necessary at this stage, i.e. before the default prediction changes and tinnitus becomes chronic. This is an ostensibly bleak outlook, especially given that the acute phase of tinnitus may be less than four weeks (Wallhäusser-Franke et al., 2017). However, our model does not imply that this is the only time window for effective intervention, merely that it may be the most effective. The question of whether tinnitus interventions are more effective early on is one that warrants formal scientific study. If it is indeed the case that tinnitus interventions have the highest success rates early on, then this would mean that health services would need to restructure and reprioritize to focus on acute tinnitus. Even so, our model does suggest other possibilities for intervention, even in chronic tinnitus.

Some possibilities for tinnitus interventions in the chronic stage that are supported by our model target the sensory evidence for tinnitus directly. These are not necessarily new ideas, as existing treatments that aim to suppress neural firing rates are essentially working to reduce the intensity of the evidence for tinnitus. Another option would be to target the precision of that evidence. This is in some ways also not a new idea, as just about every current treatment option that shows promise works to reduce neural synchrony (De Ridder et al., 2014b; Marks et al., 2018; Okamoto et al., 2010; Tass et al., 2012), which is one way of reducing precision.

The predictive coding model for tinnitus also offers new ideas for intervention. One such possibility also deals with precision, namely by teaching the brain that the evidence for tinnitus is unimportant or unreliable (and should therefore downregulate its precision in a top-down fashion). Such an intervention would effectively amount to guiding a tinnitus patient's brain towards a similar state as the brain of someone who has hearing loss and has successfully adapted to it, i.e. without developing tinnitus.

It may also be possible to eliminate the prediction of tinnitus outright. Recent evidence, for example, showed that physical disruption of white matter connecting the auditory and parahippocampal cortices led to a permanent cessation of tinnitus in the contralateral ear (Sedley et al., 2015). This is consistent with EEG observations of increased parahippocampal gamma oscillations coinciding with tinnitus in the contralateral ear (De Ridder et al., 2015; Vanneste and De Ridder, 2016). Furthermore, recent fMRI evidence shows that functional connectivity between nucleus accumbens and parahippocampal cortex correlates with tinnitus intensity as reported using a 0–10 numeric rating scale, including covariates to control for the effects of age and hearing loss (Hullfish et al., 2019). Taken together, this represents at least three independent lines of evidence suggesting that the parahippocampus is a key mediator of the tinnitus prediction and therefore an important potential target for intervention. Such intervention would not have to be as drastic as surgical disruption or even noninvasive stereotactic lesioning. Possible methods could include noninvasive brain stimulation or even EEG neurofeedback based on targeting connectivity between auditory and parahippocampal cortex.

8. Conclusion

Predictive coding asserts that tinnitus is ultimately an inference about the cause of spontaneous activity in the subcortical auditory system with an excess of precision. The notion of tinnitus subtypes, in terms of its causes, therefore arises naturally from the various factors that can increase sensory precision. Bottom-up factors are those that deprive the brain of expected auditory input, which mainly includes hearing loss, both naturally occurring and experimentally simulated. These factors spur an adaptation of sensitivity in the deprived neurons to search for these expected but missing inputs via predictive regulation (Sterling, 2018, 2012, 2004; Sterling and Laughlin, 2015). Top-down factors include failures to attenuate sensory precision as well as actively augmenting it, i.e. via endogenous attention. Furthermore, the phenomenon of somatic tinnitus demonstrates that lateral connections to

the auditory system are also important.

While these factors can all affect the evidence used to infer tinnitus, that inference also requires the presence of auditory priors. This conclusion is drawn from studies of congenitally single-sided deaf patients, who do not experience tinnitus in the deaf ear (Lee et al., 2017). This also makes intuitive sense given Bayes' rule, which takes the product of evidence and priors to produce a posterior belief. The idea is that if the prior is essentially equal to zero—i.e. an expectation of silence with a high precision—then the inference will never change regardless of the evidence; in other words, zero times anything is still zero. The powerful impact of priors on inference is also evident in the transition from acute to chronic tinnitus. The perception of tinnitus as the inferred cause of increased sensory precision can eventually shift the prior belief such that the default expectation is that of tinnitus. With this new prior, the brain can infer tinnitus even if the original evidence returns to normative levels (Fig. 3C, D).

Note especially that all of these factors are not affecting the mechanism of inference per se but rather the generative model on which such inference is based, i.e. evidence and priors. This ironically implies that tinnitus is an optimal inference, albeit one based on a suboptimal generative model; see (Schwartenbeck et al., 2015) for further discussion of this concept in the context of addictive behaviors. The critical insight here is that the brain is always working to reduce prediction error. Given the sensory evidence for tinnitus—i.e. spontaneous activity in the subcortical auditory system with an excess of precision—the brain can respond in one of two ways. The first option is to attenuate the excess precision via top-down noise cancelation. This prevents the evidence for tinnitus from updating the generative model and thus the tinnitus percept does not emerge. Should this process fail, e.g. due to dysfunction in extended frontostriatal networks (Hullfish et al., 2019; Rauschecker et al., 2015; Song et al., 2015), the only recourse for the brain is to update the generative model to account for this new evidence. This updating directly corresponds to the perception of tinnitus. While the former option is adaptive and the latter is maladaptive, both strategies minimize prediction error and thus both represent optimal inferences under predictive coding. The implication here is that the brain will tolerate hallucinatory percepts more readily than it will prediction errors, because these errors indicate homeostatic disequilibrium (Friston, 2010).

By taking the most current predictive coding model of tinnitus (Sedley et al., 2016a), unpacking it in terms of depth of detail and breadth of application, and deploying it to explain wider aspects of the phenomenon than just the initial emergence of tinnitus, we seek to lay the groundwork for more widespread acceptance of predictive coding as a framework for tinnitus specifically and for brain (dys)function generally. Furthermore, we newly introduce an account of the role of dopaminergic systems in tinnitus, i.e. as contributors to tinnitus salience, awareness, and chronification. The predictive coding framework has the potential to serve as the Rosetta Stone for the neurosciences by offering a means to translate findings across seemingly disparate lines of research. Already established parallels, such as those between tinnitus, chronic pain, and addiction, are an early indication that predictive coding can serve as the basis for a unifying theory.

Acknowledgements

Research by Jeffrey Hullfish is supported by the Eugene McDermott Graduate Fellowship [201501]. The authors would like to thank Anusha Mohan for her insights regarding cross-modal plasticity as well as Nikki Simmons for her help with figure generation and aesthetics. The authors declare no competing financial interests.

References

Alessandria, M., Vetrugno, R., Cortelli, P., Montagna, P., 2011. Normal body scheme and absent phantom limb experience in amputees while dreaming. Conscious. Cogn. 20, 1831–1834. https://doi.org/10.1016/j.concog.2011.06.013.

Apkarian, A.V., Hashmi, J.A., Baliki, M.N., 2011. Pain and the brain: Specificity and plasticity of the brain in clinical chronic pain. Pain 152. https://doi.org/10.1016/j.pain.2010.11.010.

Baliki, M.N., Chialvo, D.R., Geha, P.Y., Levy, R.M., Harden, R.N., Parrish, T.B., Apkarian, A.V., 2006. Chronic pain and the emotional brain: specific brain activity associated with spontaneous fluctuations of intensity of chronic back pain. J. Neurosci. 26, 12165–12173. https://doi.org/10.1523/JNEUROSCI.3576-06.2006.

Baliki, M.N., Geha, P.Y., Fields, H.L., Apkarian, A.V., 2010. Predicting value of pain and analgesia: nucleus accumbens response to noxious stimuli changes in the presence of chronic pain. Neuron 66, 149–160. https://doi.org/10.1016/j.neuron.2010.03.002.

Baliki, M.N., Petre, B., Torbey, S., Herrmann, K.M., Huang, L., Schnitzer, T.J., Fields, H.L., Apkarian, A.V., 2012. Corticostriatal functional connectivity predicts transition to chronic back pain. Nat. Neurosci. 15, 1117–1119. https://doi.org/10.1038/nn.3153.

Bastos, A.M., Usrey, W.M., Adams, R.A., Mangun, G.R., Fries, P., Friston, K.J., 2012. Canonical microcircuits for predictive coding. Neuron 76, 695–711. https://doi.org/10.1016/j.neuron.2012.10.038.

Chait, M., Ruff, C.C., Griffiths, T.D., McAlpine, D., 2012. Cortical responses to changes in acoustic regularity are differentially modulated by attentional load. Neuroimage 59, 1932–1941. https://doi.org/10.1016/j.neuroimage.2011.09.006.

Clark, A., 2013. Whatever next? Predictive brains, situated agents, and the future of cognitive science. Behav. Brain Sci. 36, 181–253. https://doi.org/10.1017/S0140525X12000477.

Cohen, L.G., Weeks, R.A., Sadato, N., Celnik, P., Ishii, K., Hallett, M., 1999. Period of susceptibility for cross-modal plasticity in the blind. Ann. Neurol. 45, 451–460. https://doi.org/10.1002/1531-8249(199904)45:4<451::AID-ANA6>3.0.CO;2-B.

Dayan, P., Hinton, G.E., Neal, R.M., Zemel, R.S., 1995. The helmholtz machine. Neural Comput. 7, 889–904. https://doi.org/10.1162/neco.1995.7.5.889.

De Ridder, D., Elgoyhen, A.B., Romo, R., Langguth, B., 2011. Phantom percepts: tinnitus and pain as persisting aversive memory networks. Proc. Natl. Acad. Sci. U. S. A. 108, 8075–8080. https://doi.org/10.1073/pnas.1018466108.

De Ridder, D., Joos, K., Vanneste, S., 2014a. The enigma of the tinnitus-free dream state in a Bayesian world. Neural Plast. 2014 (5). https://doi.org/10.1155/2014/612147.

De Ridder, D., Vanneste, S., Engineer, N.D., Kilgard, M.P., 2014b. Safety and efficacy of vagus nerve stimulation paired with tones for the treatment of tinnitus: a case series. Neuromodulation Technol. Neural Interface 17, 170–179. https://doi.org/10.1111/ner.12127.

De Ridder, D., Vanneste, S., Freeman, W., 2014c. The Bayesian brain: phantom percepts resolve sensory uncertainty. Neurosci. Biobehav. Rev. 44, 4–15. https://doi.org/10.1016/j.neubiorev.2012.04.001.

De Ridder, D., Vanneste, S., Weisz, N., Londero, A., Schlee, W., Elgoyhen, A.B., Langguth, B., 2014d. An integrative model of auditory phantom perception: tinnitus as a unified percept of interacting separable subnetworks. Neurosci. Biobehav. Rev. 44, 16–32. https://doi.org/10.1016/j.neubiorev.2013.03.021.

De Ridder, D., Congedo, M., Vanneste, S., 2015. The neural correlates of subjectively perceived and passively matched loudness perception in auditory phantom perception. Brain Behav. 5, 1–12. https://doi.org/10.1002/brb3.331.

De Ridder, D., Vanneste, S., Gillett, G., Manning, P., Glue, P., Langguth, B., 2016. Psychosurgery reduces uncertainty and increases free will? A review. Neuromodulation 19, 239–248. https://doi.org/10.1111/ner.12405.

Del Bo, L., Forti, S., Ambrosetti, U., Costanzo, S., Mauro, D., Ugazio, G., Langguth, B., Mancuso, A., 2008. Tinnitus aurium in persons with normal hearing: 55 years later. Otolaryngol. - Head Neck Surg. 139, 391–394. https://doi.org/10.1016/j.otohns.2008.06.019.

Denève, S., Jardri, R., 2016. Circular inference: mistaken belief, misplaced trust. Curr. Opin. Behav. Sci. 11, 40–48. https://doi.org/10.1016/j.cobeha.2016.04.001.

Donoso, M., Collins, A.G.E., Koechlin, E., 2014. Foundations of human reasoning in the prefrontal cortex. Science 344, 1481–1486. https://doi.org/10.1126/science.1252254.

Edwards, M.J., Adams, R.A., Brown, H., Pareés, I., Friston, K.J., 2012. A Bayesian account of "hysteria.". Brain 135, 3495–3512. https://doi.org/10.1093/brain/aws129.

Efron, B., Morris, C., 1973. Stein's estimation rule and its competitors—an empirical bayes approach. J. Am. Stat. Assoc. 68, 117–130. https://doi.org/10.1080/01621459.1973.10481350.

Eggermont, J.J., Kral, A., 2016. Somatic memory and gain increase as preconditions for tinnitus: insights from congenital deafness. Hear. Res. 333, 37–48. https://doi.org/10.1016/j.heares.2015.12.018.

Elgoyhen, A.B., Langguth, B., De Ridder, D., Vanneste, S., 2015. Tinnitus: perspectives from human neuroimaging. Nat. Rev. Neurosci. 16, 632–642. https://doi.org/10.1038/nrn4003.

El-Mallakh, R.S., Walker, K.L., 2010. Hallucinations, psuedohallucinations [sic], and parahallucinations. Psychiatry Interpers. Biol. Process. 73, 34–42. https://doi.org/10.1521/psyc.2010.73.1.34.

Franosch, J.M.P., Kempter, R., Fastl, H., van Hemmen, J.L., 2003. Zwicker tone illusion and noise reduction in the auditory system. Phys. Rev. Lett. 90, 4. https://doi.org/10.1103/PhysRevLett.90.178103.

Friston, K., 2003. Learning and inference in the brain. Neural Netw. 16, 1325–1352. https://doi.org/10.1016/j.neunet.2003.06.005.

Friston, K., 2010. The free-energy principle: a unified brain theory? Nat. Rev. Neurosci. 11, 127–138. https://doi.org/10.1038/nrn2787.

Friston, K., 2018. Does predictive coding have a future? Nat. Neurosci. 21, 1019–1021. https://doi.org/10.1038/s41593-018-0200-7.

Friston, K.J., Daunizeau, J., Kiebel, S.J., 2009. Reinforcement learning or active inference? PLoS One 4. https://doi.org/10.1371/journal.pone.0006421.

Friston, K.J., Shiner, T., FitzGerald, T., Galea, J.M., Adams, R., Brown, H., Dolan, R.J., Moran, R., Stephan, K.E., Bestmann, S., 2012. Dopamine, affordance and active

inference. PLoS Comput. Biol. 8. https://doi.org/10.1371/journal.pcbi.1002327.
Gardner, M.P.H., Schoenbaum, G., Gershman, S.J., 2017. Rethinking dopamine prediction errors. bioRxiv 1–21. https://doi.org/10.1101/239731.
Gregory, R.L., 1980. Perceptions as hypotheses. Philos. Trans. R. Soc. Lond., B, Biol. Sci. 290, 181–197. https://doi.org/10.1098/rstb.1980.0090.
Hashmi, J.A., Baliki, M.N., Huang, L., Baria, A.T., Torbey, S., Hermann, K.M., Schnitzer, T.J., Apkarian, A.V., 2013. Shape shifting pain: chronification of back pain shifts brain representation from nociceptive to emotional circuits. Brain 136, 2751–2768. https://doi.org/10.1093/brain/awt211.
Helmholtz, H., 1867. Handbuch der Physiologischen Optik. Voss.
Hensch, T.K., 2005. Critical period plasticity in local cortical circuits. Nat. Rev. Neurosci. 6, 877–888. https://doi.org/10.1038/nrn1787.
Hesselmann, G., Sadaghiani, S., Friston, K.J., Kleinschmidt, A., 2010. Predictive coding or evidence accumulation? False inference and neuronal fluctuations. PLoS One 5, 1–5. https://doi.org/10.1371/journal.pone.0009926.
Hobson, J.A., Friston, K.J., 2012. Waking and dreaming consciousness: neurobiological and functional considerations. Prog. Neurobiol. 98, 82–98. https://doi.org/10.1016/j.pneurobio.2012.05.003.
Hullfish, J., Abenes, I., Kovacs, S., Sunaert, S., De Ridder, D., Vanneste, S., 2018a. Functional brain changes in auditory phantom perception evoked by different stimulus frequencies. Neurosci. Lett. 683, 160–167. https://doi.org/10.1016/J.NEULET.2018.07.043.
Hullfish, J., Abenes, I., Kovacs, S., Sunaert, S., De Ridder, D., Vanneste, S., 2018b. Functional connectivity analysis of fMRI data collected from human subjects with chronic tinnitus and varying levels of tinnitus-related distress. Data Br. 21, 779–789 doi:S2352340918312617.
Hullfish, J., Abenes, I., Yoo, H.Bin, De Ridder, D., Vanneste, S., 2019. Frontostriatal network dysfunction as a domain-general mechanism underlying phantom perception. Hum. Brain Mapp. 1–11. https://doi.org/10.1002/hbm.24521.
Husain, F.T., Schmidt, S.A., 2014. Using resting state functional connectivity to unravel networks of tinnitus. Hear. Res. 307, 153–162. https://doi.org/10.1016/j.heares.2013.07.010.
Johannsson, M., Snaedal, J., Johannesson, G.H., Gudmundsson, T.E., Johnsen, K., 2015. The acetylcholine index: an electroencephalographic marker of cholinergic activity in the living human brain applied to Alzheimer's disease and other dementias. Dement. Geriatr. Cogn. Disord. 39, 132–142. https://doi.org/10.1159/000367889.
Koechlin, E., 2016. Prefrontal executive function and adaptive behavior in complex environments. Curr. Opin. Neurobiol. 37, 1–6. https://doi.org/10.1016/j.conb.2015.11.004.
Kujala, T., Alho, K., Huotilainen, M., Ilmoniemi, R.J., Lehtokoski, A., Leinonen, A., Rinne, T., Salonen, O., Sinkkonen, J., Standertskjöld-Nordenstam, C.G., Näätänen, R., 1997. Electrophysiological evidence for cross-modal plasticity in humans with early- and late-onset blindness. Psychophysiology 34, 213–216. https://doi.org/10.1111/j.1469-8986.1997.tb02134.x.
Kumar, S., Sedley, W., Barnes, G.R., Teki, S., Friston, K.J., Griffiths, T.D., 2014. A brain basis for musical hallucinations. Cortex 52, 86–97. https://doi.org/10.1016/j.cortex.2013.12.002.
Lee, S.-Y., Nam, D.W., Koo, J.-W., De Ridder, D., Vanneste, S., Song, J.-J., 2017. No auditory experience, no tinnitus: lessons from subjects with congenital- and acquired single-sided deafness. Hear. Res. 354, 9–15. https://doi.org/10.1016/j.heares.2017.08.002.
Levine, R.A., 1999. Somatic (craniocervical) tinnitus and the dorsal cochlear nucleus hypothesis. Am. J. Otolaryngol. - Head Neck Med. Surg. 20, 351–362. https://doi.org/10.1016/S0196-0709(99)90074-1.
Levine, R.A., Abel, M., Cheng, H., 2003. CNS somatosensory-auditory interactions elicit or modulate tinnitus. Exp. Brain Res. 153, 643–648. https://doi.org/10.1007/s00221-003-1747-3.
Liberman, M.C., 2015. Hidden hearing loss. Sci. Am. 313, 48–53. https://doi.org/10.1038/scientificamerican0815-48.
Linszen, M.M.J., Brouwer, R.M., Heringa, S.M., Sommer, I.E., 2016. Increased risk of psychosis in patients with hearing impairment: review and meta-analyses. Neurosci. Biobehav. Rev. 62, 1–20. https://doi.org/10.1016/j.neubiorev.2015.12.012.
Marks, K.L., Martel, D.T., Wu, C., Basura, G.J., Roberts, L.E., Schvartz-Leyzac, K.C., Shore, S.E., 2018. Auditory-somatosensory bimodal stimulation desynchronizes brain circuitry to reduce tinnitus in guinea pigs and humans. Sci. Transl. Med. 10, eaal3175. https://doi.org/10.1126/scitranslmed.aal3175.
Mathys, C., 2011. A Bayesian foundation for individual learning under uncertainty. Front. Hum. Neurosci. 5, 1–20. https://doi.org/10.3389/fnhum.2011.00039.
Mohan, A., Vanneste, S., 2017. Adaptive and maladaptive neural compensatory consequences of sensory deprivation—from a phantom percept perspective. Prog. Neurobiol. 153, 1–17. https://doi.org/10.1016/j.pneurobio.2017.03.010.
Mulder, T., Hochstenbach, J., Dijkstra, P.U., Geertzen, J.H.B., 2008. Born to adapt, but not in your dreams. Conscious. Cogn. 17, 1266–1271. https://doi.org/10.1016/j.concog.2007.04.001.
Mutso, A.A., Radzicki, D., Baliki, M.N., Huang, L., Banisadr, G., Centeno, M.V., Radulovic, J., Martina, M., Miller, R.J., Apkarian, A.V., 2012. Abnormalities in hippocampal functioning with persistent pain. J. Neurosci. 32, 5747–5756. https://doi.org/10.1523/JNEUROSCI.0587-12.2012.
Okamoto, H., Stracke, H., Stoll, W., Pantev, C., 2010. Listening to tailor-made notched music reduces tinnitus loudness and tinnitus-related auditory cortex activity. Proc. Natl. Acad. Sci. U. S. A. 107, 1207–1210. https://doi.org/10.1073/pnas.0911268107.
Payzan-Lenestour, E., Bossaerts, P., 2011. Risk, unexpected uncertainty, and estimation uncertainty: Bayesian learning in unstable settings. PLoS Comput. Biol. 7. https://doi.org/10.1371/journal.pcbi.1001048.
Payzan-LeNestour, E., Dunne, S., Bossaerts, P., O'Doherty, J., 2013. The neural representation of unexpected uncertainty during value-based decision making. Neuron 79, 191–201. https://doi.org/10.1016/j.neuron.2013.04.037.
Pezzulo, G., Rigoli, F., Friston, K., 2015. Active Inference, homeostatic regulation and adaptive behavioural control. Prog. Neurobiol. 134, 17–35. https://doi.org/10.1016/j.pneurobio.2015.09.001.
Rabinowitch, I., Bai, J., 2016. The foundations of cross-modal plasticity. Commun. Integr. Biol. 9, 1–3. https://doi.org/10.1080/19420889.2016.1158378.
Ramachandran, V.S., Hirstein, W., 1998. The perception of phantom limbs. The D.O. Hebb lecture. Brain J. Neurol. 121, 1603–1630.
Rao, R.P.N., Ballard, D.H., 1999. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. Nat. Neurosci. 2, 79–87. https://doi.org/10.1038/4580.
Rauschecker, J.P., May, E.S., Maudoux, A., Ploner, M., 2015. Frontostriatal gating of tinnitus and chronic pain. Trends Cogn. Sci. 19, 567–578. https://doi.org/10.1016/j.tics.2015.08.002.
Roberts, L.E., Husain, F.T., Eggermont, J.J., 2013. Role of attention in the generation and modulation of tinnitus. Neurosci. Biobehav. Rev. 37, 1754–1773. https://doi.org/10.1016/j.neubiorev.2013.07.007.
Sadato, N., Okada, T., Honda, M., Yonekura, Y., 2002. Critical period for cross-modal plasticity in blind humans: a functional MRI study. Neuroimage 16, 389–400. https://doi.org/10.1006/nimg.2002.1111.
Schaette, R., McAlpine, D., 2011. Tinnitus with a normal audiogram: physiological evidence for hidden hearing loss and computational model. J. Neurosci. 31, 13452–13457. https://doi.org/10.1523/JNEUROSCI.2156-11.2011.
Schaette, R., Turtle, C., Munro, K.J., 2012. Reversible induction of phantom auditory sensations through simulated unilateral hearing loss. PLoS One 7, 1–6. https://doi.org/10.1371/journal.pone.0035238.
Schlee, W., Kleinjung, T., Hiller, W., Goebel, G., Kolassa, I.T., Langguth, B., 2011. Does tinnitus distress depend on age of onset? PLoS One 6, 4–7. https://doi.org/10.1371/journal.pone.0027379.
Schwartenbeck, P., FitzGerald, T.H.B., Mathys, C., Dolan, R., Wurst, F., Kronbichler, M., Friston, K., 2015. Optimal inference with suboptimal models: addiction and active Bayesian inference. Med. Hypotheses 84, 109–117. https://doi.org/10.1016/j.mehy.2014.12.007.
Sedley, W., 2019. Tinnitus: Does gain explain? Neuroscience. https://doi.org/10.1016/J.NEUROSCIENCE.2019.01.027.
Sedley, W., Gander, P.E., Kumar, S., Oya, H., Kovach, C.K., Nourski, K.V., Kawasaki, H., Howard, M.A., Griffiths, T.D., 2015. Intracranial mapping of a cortical tinnitus system using residual inhibition. Curr. Biol. 25, 1208–1214. https://doi.org/10.1016/j.cub.2015.02.075.
Sedley, W., Friston, K.J., Gander, P.E., Kumar, S., Griffiths, T.D., 2016a. An integrative tinnitus model based on sensory precision. Trends Neurosci. 39, 799–812. https://doi.org/10.1016/j.tins.2016.10.004.
Sedley, W., Gander, P.E., Kumar, S., Kovach, C.K., Oya, H., Kawasaki, H., Howard, M.A., Griffiths, T.D., 2016b. Neural signatures of perceptual inference. Elife 5, 1–14. https://doi.org/10.7554/eLife.11476.
Seeley, W.W., Menon, V., Schatzberg, A.F., Keller, J., Glover, G.H., Kenna, H., Reiss, A.L., Greicius, M.D., 2007. Dissociable intrinsic connectivity networks for salience processing and executive control. J. Neurosci. 27, 2349–2356. https://doi.org/10.1523/JNEUROSCI.5587-06.2007.
Shore, S.E., 2005. Multisensory integration in the dorsal cochlear nucleus: unit responses to acoustic and trigeminal ganglion stimulation. Eur. J. Neurosci. 21, 3334–3348. https://doi.org/10.1111/j.1460-9568.2005.04142.x.
Shore, S.E., Koehler, S., Oldakowski, M., Hughes, L.F., Syed, S., 2008. Dorsal cochlear nucleus responses to somatosensory stimulation are enhanced after noise-induced hearing loss. Eur. J. Neurosci. 27, 155–168. https://doi.org/10.1111/j.1460-9568.2007.05983.x.
Shore, S.E., Zhou, J., 2006. Somatosensory influence on the cochlear nucleus and beyond. Hear. Res. 216–217, 90–99. https://doi.org/10.1016/j.heares.2006.01.006.
Song, J.J., Vanneste, S., De Ridder, D., 2015. Dysfunctional noise cancelling of the rostral anterior cingulate cortex in tinnitus patients. PLoS One 10, 1–18. https://doi.org/10.1371/journal.pone.0123538.
Sterling, P., 2004. Principles of allostasis: optimal design, predictive regulation, pathophysiology, and rational therapeutics. In: Schulkin, J. (Ed.), Allostasis, Homeostasis, and the Costs of Physiological Adaptation. Cambridge University Press, pp. 17–64. https://doi.org/10.1017/CBO9781316257081.004.
Sterling, P., 2012. Allostasis: a model of predictive regulation. Physiol. Behav. 106, 5–15. https://doi.org/10.1016/j.physbeh.2011.06.004.
Sterling, P., 2018. Predictive regulation and human design. Elife 7, 1–12. https://doi.org/10.7554/eLife.36133.
Sterling, P., Laughlin, S., 2015. Principles of Neural Design, Principles of Neural Design. MIT Presshttps://doi.org/10.1017/CBO9781107415324.004.
Takahashi, Y.K., Batchelor, H.M., Liu, B., Khanna, A., Morales, M., Schoenbaum, G., 2017. Dopamine neurons respond to errors in the prediction of sensory features of expected rewards. Neuron 95, 1395–1405. https://doi.org/10.1016/j.neuron.2017.08.025. e3.
Tass, P.A., Adamchic, I., Freund, H.J., Von Stackelberg, T., Hauptmann, C., 2012. Counteracting tinnitus by acoustic coordinated reset neuromodulation. Restor. Neurol. Neurosci. 30, 137–159. https://doi.org/10.3233/RNN-2012-110218.
Tonndorf, J., 1987. The analogy between tinnitus and pain: a suggestion for a physiological basis of chronic tinnitus. Hear. Res. 28, 271–275. https://doi.org/10.1016/0378-5955(87)90054-2.
Tucker, D.A., Phillips, S.L., Ruth, R.A., Clayton, W.A., Royster, E., Todd, A.D., 2005. The effect of silence on tinnitus perception. Otolaryngol. - Head Neck Surg. 132, 20–24. https://doi.org/10.1016/j.otohns.2005.08.016.
Vanneste, S., De Ridder, D., 2011. Bifrontal transcranial direct current stimulation modulates tinnitus intensity and tinnitus-distress-related brain activity. Eur. J.

Neurosci. 34, 605–614. https://doi.org/10.1111/j.1460-9568.2011.07778.x.
Vanneste, S., De Ridder, D., 2016. Deafferentation-based pathophysiological differences in phantom sound: tinnitus with and without hearing loss. Neuroimage 129, 80–94. https://doi.org/10.1016/j.neuroimage.2015.12.002.
Vanneste, S., Song, J.J., De Ridder, D., 2013. Tinnitus and musical hallucinosis: the same but more. Neuroimage 82, 373–383. https://doi.org/10.1016/j.neuroimage.2013.05.107.
Vilares, I., Kording, K., 2011. Bayesian models: the structure of the world, uncertainty, behavior, and the brain. Ann. N. Y. Acad. Sci. 1224, 22–39. https://doi.org/10.1111/j.1749-6632.2011.05965.x.
Wallhäusser-Franke, E., D'Amelio, R., Glauner, A., Delb, W., Servais, J.J., Hörmann, K., Repik, I., 2017. Transition from acute to chronic tinnitus: predictors for the development of chronic distressing tinnitus. Front. Neurol. 8, 1–13. https://doi.org/10.3389/fneur.2017.00605.
Waters, F., Blom, J.D., Dang-Vu, T.T., Cheyne, A.J., Alderson-Day, B., Woodruff, P., Collerton, D., 2016. What is the link between hallucinations, dreams, and hypnagogic-hypnopompic experiences? Schizophr. Bull. 42, 1098–1109. https://doi.org/10.1093/schbul/sbw076.
Wearne, D., Genetti, A., 2015. Pseudohallucinations versus hallucinations: wherein lies the difference? Australas. Psychiatry 23, 254–257. https://doi.org/10.1177/1039856215586150.
Young, E.D., Nelken, I., Conley, R.A., 1995. Somatosensory effects on neurons in dorsal cochlear nucleus. J. Neurophysiol. 73, 743–765.
Zwicker, E., 1964. "Negative afterimage" in hearing. J. Acoust. Soc. Am. 36, 2413–2415. https://doi.org/10.1121/1.1919373.
Zwicker, E., 1979. A model describing nonlinearities in hearing by active processes with saturation at 40 dB. Biol. Cybern. 35, 243–250. https://doi.org/10.1007/BF00344207.