BURDENING THE PREDICTIVE MIND: A PREDICTIVE PROCESSING APPROACH TO HEALTH-RELATED BURDENS

Paul Smart
Electronics and Computer Science
University of Southampton
Southampton, SO17 1BJ, UK
ps02v@ecs.soton.ac.uk

Nic Fair
Electronics and Computer Science
University of Southampton
Southampton, SO17 1BJ, UK
N.S.Fair@soton.ac.uk

Simon D. S. Fraser
School of Primary Care
Population Science and Medical Education
Faculty of Medicine
University of Southampton
Southampton, SO17 1BJ, UK
S.Fraser@soton.ac.uk

Michael Boniface
Electronics and Computer Science
University of Southampton
Southampton, SO17 1BJ, UK
M.J.Boniface@soton.ac.uk

ABSTRACT

The notion of burden features as a central aspect of research into the challenges faced by patients and their carers, especially in regard to long-term health conditions and multimorbidity. Research in this area has considered the burdens that stem from the presence of disease (e.g., symptom burden) as well as the burden associated with healthcare interventions (e.g., treatment burden). While there have been a number of attempts to theorize burden, there is, at present, little consensus on how burdens ought to be understood. It is, in particular, unclear what makes something a burden, why certain things are perceived as burdensome, what forces and factors moderate the experience of burden, and how burdensome experiences relate to other experiential constructs, such as wellbeing, despair, and suffering. The present paper seeks to advance our understanding of burden by drawing on predictive processing accounts of brain function. All burdens, it is suggested, have their origins in a reduced capacity to fulfill neurally-realized expectations (or predictions). This is marked by a hypothesized increase in a particular form of prediction error, dubbed expected prediction error. In addition to providing a unitary theoretical approach to burden, the present account supports the effort to apply predictive processing to a wider array of clinical and health-related phenomena.

Keywords Symptom Burden · Treatment Burden · Active Inference · Predictive Processing · Adherence · Epidemiology · General Practice

1 Introduction

Within the sciences of the mind and brain, there has been growing interest in the idea that the brain is an engine of probabilistic prediction [Clark, 2013, 2016, Friston, 2010, Hohwy, 2013]. This idea—commonly referred to as "predictive processing"—depicts the brain as a hierarchically-organized system that is constantly striving to minimize the error associated with its own predictions. Remarkably, this simple imperative yields an integrated approach to understanding a rich array of cognitive phenomena, including attention, memory, action, perception, imagination, emotion, and even aspects of conscious experience [Clark, 2016, 2019, 2023, Hohwy and Seth, 2020]. Predictive processing has also yielded insights into the nature of neural processing, providing new ways of understanding the computational significance of neurochemical systems [Schwartenbeck et al., 2015], anatomical structures [Parr et al., 2020], and connectional principles [Bastos et al., 2012, Shipp, 2016].

While predictive processing is intended as a general theoretical account of neural functioning, its explanatory successes have spawned significant interest in its application to more specialized areas, such as our understanding of health-related phenomena. Recent work in this area has sought to develop predictive processing accounts of placebo and nocebo effects [Büchel et al., 2014, Clark, 2023, Ongaro and Kaptchuk, 2019], medical adherence decisions [Smith et al., 2021], symptom perception [Pezzulo et al., 2019, Van den Bergh et al., 2017], depression [Arnaldo et al., 2022, Fabry, 2020, Stephan et al., 2016], neuropsychiatric conditions [Kiverstein et al., 2019, Montague et al., 2012], developmental disorders [Hohwy, 2015], and the role of social support in determining health outcomes [Smith et al., 2019].

As suggested by these examples, there is considerable interest in applying the predictive processing framework to the phenomenological aspects of illness (e.g., the experience of pain, fatigue, or other symptoms). This establishes a potent point of contact with work exploring the experiential (or phenomenological) aspects of health and disease [see Burch, 2023, Carel, 2011, 2016]. This includes, of course, a concern with the nature of the illness experience (e.g., the way that patients experience a change in bodily functioning); but it also extends to additional issues, such as the experiential response to healthcare interventions, and the more general role of lived experience in charting the course of health-related trajectories, both prior to the onset of disease and subsequent to diagnosis.

In this paper, we seek to apply the predictive processing framework to a particular aspect of patient experience, namely the experience of burden. The study of burden has a long history in biomedical research. While early research was directed to the more general notion of illness burden (or the burden of illness) [e.g., Downes, 1942], work of a more recent vintage has specialized in the study of particular types of burden, such as the burden imposed by symptoms (symptom burden) or the burden that comes with the performance of healthcare tasks (e.g., treatment burden). Additional forms of specialization occur in respect of particular diseases, as when researchers seek to understand the challenges associated with diseases such as cancer [e.g., Gapstur, 2007, Spratt et al., 2021], diabetes [Bohlen et al., 2012, Eton et al., 2013, Ludman et al., 2004], stroke [Gallacher et al., 2018], psychotic disorders [Iversen et al., 2018], and COVID-19 [Sykes et al., 2021].

Unfortunately, despite the considerable interest in burden, our conceptual understanding of burden remains somewhat poor. It is, in particular, unclear what makes something a burden, why certain things are perceived as burdensome, what forces and factors moderate the experience of burden, and how burdensome experiences relate to other aspects of human experience (e.g., fatigue, boredom, wellbeing, happiness, suffering, and despair). Addressing all these issues is clearly beyond the scope of a single paper; nevertheless, by drawing on the resources of the predictive processing framework we aim to provide an account that connects burdensome experiences to the error-minimizing dynamics of the predictive brain. This, we suggest, opens the door to research that seeks to connect the study of burden to a wider array of clinical and health-related phenomena.

The structure of the paper is as follows: A brief overview of predictive processing and its application to burden is presented in Section 2. Here, we outline the basic structure of a predictive processing account of burden, what we call the expectation-based account. Subsequent sections apply this account to three types of burden that have been discussed in the medical literature. Symptom burden is discussed in Section 3, treatment burden in Section 4, and side-effects burden in Section 5. The paper concludes with a discussion of directions for future research (Section 6).

2 Predictive Processing

The term "predictive processing" refers to a recent account of brain function that depicts the brain as a hierarchically-organized prediction machine [Clark, 2013, 2016, Friston, 2010, Hohwy, 2013]. According to predictive processing, the brain is engaged in a ceaseless, multi-layered effort to predict its own activity. Higher layers in the hierarchy issue *predictions* about the activity of lower layers, and the mismatch between predicted and actual activity is returned to the higher levels in the form of *prediction error*. This process unfolds at each and every layer in the hierarchy, terminating at the organism's sensory surfaces—the point at which sensory information is converted to neural signals. Given that the primary source of perturbation for the brain is the incoming sensory information, the overall predictive success of the brain is determined by its ability to anticipate moment-by-moment shifts in the incoming sensory stream. This is achieved via the acquisition of a neurally-realized *generative model* that embodies the causal dynamics of the sensorium—i.e., the causal forces and factors that give rise to sensory input [Friston, 2010, Parr et al., 2022].[1] If the generative model succeeds in latching onto these causal forces and factors, then it is in a position to predict how the current sensory situation will evolve into the future.

According to predictive processing, the overarching computational imperative of the brain is to minimize prediction error. Over the longer term, this can be achieved by ensuring that generative models are suitably aligned with the

[1] As characterized by Clark [2016, p. 21]: "A generative model [...] aims to capture the statistical structure of some set of observed inputs by inferring a causal matrix able to give rise to that very structure."

statistical structure of the sensory environment in which the generative model is situated. Such forms of adjustment and adaptation occur as the result of learning, where the prediction error signal is used to progressively refine the parameters of a generative model.

In addition to learning, there are two further ways the brain can minimize prediction error. The first is what is sometimes called *perceptual inference* [Badcock et al., 2019, Fabry, 2020, Kiverstein et al., 2022]. In this case, higher-level predictions are adjusted so as to better accommodate the incoming sensory signal. In short, the brain adjusts its predictions to best account for (or 'explain away') the moment-to-moment shifts in the sensory stream. This is what we commonly understand as perception in all its forms (e.g., visual perception, auditory perception, interoception, and so on).

The second route to minimizing prediction error is what is called *active inference* [Friston et al., 2017, Linson et al., 2018, Parr et al., 2022].[2] In this case, the brain relies on action to minimize prediction error. Rather than change predictions to match the world (as is the case with perceptual inference), active inference relies on action to change the sensory signal in a way that conforms to predictions. This is what we commonly understand as action in all its forms (e.g., somatomotor and visceromotor responses).

While perceptual inference is no doubt relevant to our understanding of burden, our focus in the present paper lies with active inference. We are, in particular, concerned with the application of active inference to matters of action selection and evaluation. As noted by Smith et al. [2022a], there is an important distinction between the application of active inference to matters of motor control (what Smith et al. call *motor active inference*) and its application to matters of action selection (what Smith et al. call *decision active inference*). In the present paper, we will be concerned with the latter form of active inference (i.e., decision active inference).

The general idea behind decision active inference is that action policies are selected based on their capacity to minimize prediction errors *in the future*. In short, action policies are scored based on their predicted potential to bring about sensory states that the brain optimistically expects itself to be in. Within the predictive processing literature, these sensory states are known by many names. They have been variously referred to as target priors [Clark, in press], prior preferences [Peters et al., 2017, Smith et al., 2022b], prior expectations [Friston, 2010], preferred outcomes [Arnaldo et al., 2022], expected states [Kiverstein and Miller, 2023], desired outcomes [Van de Cruys and Van Dessel, 2021], and attracting states [Friston et al., 2013]. For present purposes, we will simply refer to these expected/predicted sensory states as *expectations*. An expectation is just a particular sort of prediction, one that encodes the sensory states-of-affairs that an organism expects itself to be in. As was noted above, these are *optimistic* expectations (or optimistic predictions); they are not predictions pertaining to the organism's anticipated success or failure in fulfilling expectations—that is another sort of prediction (see below). The reason for this optimistic bias becomes apparent when we consider expectations pertaining to basic bodily states, such as states of hydration or the availability of metabolic resources. As noted by Van de Cruys et al. [2020]:

> Our predictive models can (must) be optimistically biased, in that the distribution of expected states is realized, when we act upon the world. For example, the model may include interoceptive expectations on adequate glucose levels throughout a famine. These allostatic predictions are deep set, ingrained in low-level structural mechanisms (and underwrite foraging for food) … [Van de Cruys et al., 2020, p. 680]

The optimistic bias of expectations is thus important when it comes to their role in controlling action (specifically, the selection of action policies). If the brain expected itself *not* to be hydrated, for example, then it would fail to select actions that ensured its future hydrational status. By contrast, if the brain expects itself to be hydrated (regardless of the prevailing conditions), then action policies can be evaluated relative to their success in bringing about the expected outcome. The overarching objective here is to ensure that one's actions are coordinated with respect to one's expectations, such that actual sensory futures are aligned with expected sensory futures. If an action policy is predicted to culminate in these expected sensory states (i.e., to fulfill expectations), then it will have a greater likelihood of being selected for implementation. This is because the best action policies are simply those that succeed in fulfilling expectations (or matching optimistic predictions). This is just another way of saying that the best action policies are those that have the best chance of minimizing prediction error in the future.[3]

[2] It should be noted that the term "active inference" is often used in a way that subsumes predictive processing.

[3] Technically, the best action policy is the one that minimizes an information-theoretic quantity known as *expected free energy* [Friston et al., 2015, Parr et al., 2022]. Expected free energy refers to the likelihood that a given action policy will terminate in an expected state. Expected free energy is minimized when the predicted outcome of an action policy matches the expected state. Mismatches, by contrast, yield progressively greater levels of expected free energy. In this sense, expected free energy represents the error that would be obtained if a given action policy were to be performed, where this error is calculated relative to what we are calling an expectation.

As should be clear from the earlier appeal to interoceptive expectations, there is no reason why expectations ought to be understood in the same way we understand consciously-accessible (or personal-level) expectations (i.e., expectations that we are aware of). For the most part, expectations are understood to be sub-personal (neural) predictions that operate below the level of conscious awareness. At the same time, there is no reason why we cannot think of expectations as akin to goals, objectives, desires, and other conative constructs.[4] Indeed, there have been a number of attempts to situate these aspects of the human mental economy within the predictive processing framework [see Clark, 2020, Smith et al., 2022a]. From this perspective, a goal is just a particular sort of expectation (or optimistic prediction). A goal encodes the target state-of-affairs that we want to produce as the result of performing certain actions, and our actions are selected on the basis of their perceived potential to produce this target state-of-affairs.

There is one further aspect of the predictive processing story that needs to be covered before we turn our attention to burden. This relates to the notion of *precision*. At a general level, precision refers to the accuracy or certainty of predictions. When applied to expectations, precision can be understood to represent the relative importance or priority of expectations. The greater the precision of an expectation, the more important it is for action policies to fulfill that expectation, and thus the greater the control the expectation exerts over overt behavior. This gives us a means of understanding the role of context in shaping goal-directed behavior, especially in regard to fluctuating motivational imperatives [Pezzulo et al., 2015, 2018]. As the precision of expectations pertaining to hydrational status increases, for example, so the greater the influence of these expectations in selecting action policies (e.g., those associated with drinking).

Precision thus gives us a means of understanding the moment-to-moment shifts in behavior—why some action policies might be abandoned in favor of others. It is, however, important to remember the hierarchical aspect of predictive processing. Predictive processing, recall, depicts the brain as a hierarchical (multi-layer) system, with different layers trading in predictions of varying spatial and temporal scale. This can sometimes lead to conflicts, as when the pursuit of a higher-level expectation entails the violation of a lower-level expectation. A nice example of this is provided by Kiverstein and Miller [2023] [drawing on work by Smith et al., 2022c]:

> Keeping expected prediction errors to a minimum over time will sometimes require prediction errors to remain high at some lower layers of the generative model so that the expected prediction errors can be minimized at higher levels. Smith and colleagues give the example of high expenditure of metabolic energy during busy periods in one's work life [Smith et al., 2022c]. This expenditure of energy may well lead to high-allostatic load in terms of the stress and fatigue the person experiences. However, they may also find a sense of purpose and meaning in work that allows for the person to act in ways that are consistent with higher-level self-beliefs about what a good life means for them. Thus, even though errors are sustained at high levels at some layers of the model, this is the price to be paid for reducing expected prediction errors arising from higher-level self-beliefs that concern what a meaningful life looks like for them in the long-run. [Kiverstein and Miller, 2023, p. 3]

Kiverstein and Miller [2023] are mainly concerned with the notion of 'human flourishing'—what it means for someone to lead a 'good life', or to lead a life that is (one presumes) dominated by largely positive experiences. The focus of the present paper is, of course, somewhat different. Nevertheless, the appeal to *expected prediction error* features as a central element of the present attempt to provide a predictive processing account of burden. This is what we call the *expectation-based account*.

The expectation-based account rests on the idea that burdensome experiences are tied to increases in expected prediction error. Expected prediction error is the error that we (or rather our brains) expect to encounter in the future given the pursuit of particular action policies. Action policies, recall, are selected according to their potential to minimize future forms of prediction error. As noted by Kiverstein and Miller [2023], however, there is no reason why we should expect all forms of (future) prediction error to be minimized by an action policy. Sometimes the pursuit of a high-level personal goal requires us to violate lower-level expectations pertaining to (e.g.) levels of energy expenditure. At other times, the fulfillment of a lower-level expectation entails the violation of a higher-level (and usually longer-term) goal (such as when we are on a diet, but find ourselves succumbing to a tasty, but supposedly forbidden, treat). Then there are those occasions where we simply lack the capacity to fulfill an expectation. The seasoned runner may expect themselves to go for a daily jog, but then find themselves unable to run as the result of a bodily injury. A similar problem can arise when one has multiple high-precision expectations, not all of which can be fulfilled at the same time. In such situations, the individual may find themselves in a "damned if I do, damned if I don't" type situation—a situation where the effort to fulfill one expectation precludes the possibility of fulfilling another.

[4] Friston [2010, p. 134], for example, suggests that "goals can be considered as prior expectations that an action is obliged to fulfil."

All these situations, we suggest, involve a future state-of-affairs in which one or more expectations (i.e., goals, prior preferences, desired outcomes, etc.) are being violated. The anticipation (or prediction) of those violations is simply the quantity we have labeled expected prediction error. In this sense, then, expected prediction error is not so much a form of prediction error as it is a prediction about prediction error—it is the error we expect ourselves to encounter in the future as a result of the situations in which we find ourselves. Crucially, however, expected prediction error is not merely a prediction about there being troubled (or perhaps uncertain[5]) times ahead. It is more that our expectations are violated as the result of the expectations we have, and the way we pursue certain action policies at the expense of others. If one is carrying a heavy backpack, for example, then one may experience the backpack as a burden on account of the fact that certain expectations (e.g., those pertaining to interoceptive parameters) are being violated. But the reason these expectations are being violated is because one has opted to continue carrying the backpack. In such cases, the sense of burden is tied to our persistence with a certain course of action. But the reason why we stay the action-related course is often because we have no genuine alternative—the selected policy is simply the least worst option. The best we can do in such situations is simply bear the 'load'.

The details of the expectation-based account will become clearer as the discussion progresses. In subsequent sections, we seek to apply the account to three types of burden, namely, symptom burden, treatment burden, and side-effects burden. Each of these types of burden provides opportunities to explore different aspects of the expectation-based account.

3 Symptom Burden

Symptom burden is one of a number of types of burden that have been discussed in the medical and psychological literature [Cleeland, 2007, Gapstur, 2007, Gill et al., 2012]. It refers to the burden caused by (or at least associated with) symptoms of one sort or another. While there is some disagreement about the nature of symptoms [see Cox et al., 2014, Scheuermann et al., 2009], symptoms are typically understood to be subjective experiences, especially ones that indicate the presence of a disease or illness [e.g., Dodd et al., 2001]. Common examples of symptoms include pain, fatigue, anxiety, breathlessness, emotional distress, nausea, and perceived alterations in sensorimotor function. As noted by Cleeland [2007, p. 17], such experiences are seldom deemed to be particularly pleasant—as something we would like to experience. In this respect, symptoms might be seen to have a valence, similar to that attributed to emotional and affective states [see Carruthers, 2018].[6] The extent to which all symptoms can be seen to have a negative valence remains unclear; nevertheless, a consideration of valence (i.e., the perceived unpleasantness of symptoms) may be relevant to our understanding of symptom burden.[7]

Within the existing literature, there have been a number of attempts to define symptom burden. One definition is owed to Gapstur [2007]. Based on an extensive review of the psychological, medical, and nursing literature, Gapstur offers the following definition of symptom burden:

> Symptom burden is defined as the subjective, quantifiable prevalence, frequency, and severity of symptoms placing a physiologic burden on patients and producing multiple negative, physical, and emotional patient responses. [Gapstur, 2007, p. 673]

A second, not altogether unrelated, definition is proposed by Cleeland [2007]:

> Symptom burden can be thought of as the sum of the severity and impact of symptoms reported by a significant proportion of patients with a given disease or treatment. [Cleeland, 2007, p. 17]

Unfortunately, neither of these definitions are devoid of problems. Gapstur's definition, for example, involves an appeal to physiologic (or physiological) burden, but it is not entirely clear what is meant by the notion of physiologic burden.[8] The problem, here, is that the definitional emphasis has shifted from one sort of burden ("What is a symptom burden?")

[5] For reasons of simplicity, we refrain from a more detailed discussion of uncertainty. Note, however, that uncertainty (and uncertainty minimization) are central aspects of the predictive processing framework [e.g., Parr and Friston, 2017, Peters et al., 2017]. Uncertainty is also a recurring feature of the burden-related literature [e.g., Etkind et al., 2022].

[6] The concept of "valence" refers to the positive or negative ("good" or "bad") felt character of an affective state.

[7] While we do not explore the links between valence and burden in the present work, it is worth noting that valence has been a feature of recent predictive processing work [Fernandez Velasco and Loev, 2021, Hesp et al., 2020, Joffily and Coricelli, 2013, Miller et al., 2022]. Understanding the relationship between these accounts and the present expectation-based account is thus an important area for future research.

[8] Gapstur [2007, p. 676] connects the notion of physiologic burden to "a change in the normal physical functioning of the human body." While such changes are undoubtedly a common (perhaps universal) feature of disease processes, it remains unclear what it is about such changes that makes them burdensome.

to another sort of burden ("What is a physiological burden?"), but in the absence of a better understanding of what it means for something to qualify as a burden, we are arguably no further forward.[9]

While Cleeland's definition doesn't suffer from this particular problem, it also leaves us with a number of unanswered questions. According to Cleeland, symptom burden can be thought of as the sum of two quantities(?), namely symptom severity and symptom impact. The problem is that the relationship between these two quantities is left obscure. One assumes that symptom severity must be distinct from symptom impact; otherwise, there would be little point in adding the two together to yield an overall measure of symptom burden. It would, for example, make little sense to say that symptom impact and symptom severity are the same quantity, or that they are different manifestations of the same underlying quantity. This raises a question about the relationship between symptom impact and symptom severity. Are there certain types of symptoms that can be rated as severe, but which have no discernible impact on the affected individual? If so, what are these symptoms? Conversely, is it possible for a symptom to have both high impact and low severity? The answer to these questions is, unfortunately, unclear.

The more general problem with the definitional efforts of both Gapstur and Cleeland is that they fail to tackle the foundational notion of burden. That is to say, they do not tell us what a burden is or why something like a symptom would be the source of a burdensome experience. Quite plausibly, all symptoms are burdensome, at least to some extent—they are, after all, experiences that we would prefer not to have. But what is it that makes some symptoms more burdensome than others? And how does this particular aspect of patient phenomenology (i.e., the way patients experience symptoms) relate to a wider set of forces and factors, including other aspects of an individual's lived experience.

The expectation-based account seeks to answer these questions by providing an overarching account of burden, one that seeks to identify the common features of burden-related scenarios (i.e., the situations in which burdensome experiences are apt to arise). In Section 2, we suggested that all burdens could be understood in more or less the same way—as a form of disruption or interference with an individual's capacity to minimize future forms of prediction error. In respect of symptom burden, then, the expectation-based account suggests that symptoms become burdensome when they interfere with our capacity to fulfill expectations, thereby yielding an increase in expected prediction error. The greater this interference, the greater the prediction error, and thus the more burdensome the symptom.

As a means of making this clearer, it will help to connect the discussion of the expectation-based account to the features of particular scenarios (or cases). Our first scenario is centered on pain. Call it the toothache case.

> **The Toothache Case**
> Anna is a student who is confronted with an important essay deadline. Unfortunately, Anna begins to experience a toothache. The toothache becomes progressively worse, to the point where Anna is no longer able to concentrate on her studies. She books an appointment with her local dentist, but she nevertheless needs to wait a couple of days before she can be seen. In the meantime, Anna does her best to work despite the pain. She takes some painkillers, which help to alleviate the pain. When she goes to bed, however, the pain seems to worsen, and Anna finds it difficult to sleep. The next day, Anna tries to return to her studies. Unfortunately, she is now tired as well as in pain. Realizing that there is little chance of meeting the original deadline, Anna appeals to her lecturer for a deadline extension.

In considering this case, we can ask ourselves why the relevant symptom (i.e., pain) might be perceived as burdensome.[10] There are a couple of things that stand out here. The most salient feature of the case is that Anna is experiencing pain, and this pain is interfering with something that Anna is required to do. Specifically, the pain is interfering with Anna's ability to concentrate on her work. Relative to the expectation-based account, Anna can be seen to have an expectation pertaining to the completion of a task. This expectation corresponds to the goal of completing the essay, and it is this

[9] A similar issue can sometimes be found in discussions of treatment burden (see Section 4). Sav et al. [2013a, p. 320], for example, define treatment burden as "a person's subjective and objective overall estimation of the dynamic and multidimensional burden that their treatment regimen for chronic illness has imposed on them and on their family members." The problem with this definition is that it makes the definition of one sort of burden (treatment burden) contingent on another sort of burden (multidimensional burden). In the absence of an understanding of what burdens are, however, the definition risks being uninformative.

[10] Note that it is probably not enough to say that the mere experience of pain is sufficient for burdensomeness. The reason for this is that it is possible to imagine situations where an individual claims to experience pain, but the pain is not deemed to be particularly burdensome. Consider, for example, the case of pain asymbolia, where an individual claims to experience pain without feeling particularly bothered by it [Gerrans, 2020, Klein, 2015]. In this case, it is plausible that we encounter a dissociation between the mere presence of a symptom and the symptom's overall impact on the individual concerned. In many cases, pain is apt to be a potent source of burdensome experiences, but it need not always be so. Much will depend on how the individual perceives the pain relative to the sorts of things they expect themselves to do, or (equivalently) the sorts of sensory futures they expect themselves to bring about as a result of their own actions.

goal that defines the intentional or goal-directed nature of Anna's essay-writing efforts. The problem is that the presence of the pain is threatening to derail Anna's capacity to achieve this goal. (Equivalently, we might say that the pain is interfering with Anna's capacity to perform actions that ensure the timely fulfillment of an expectation that is accorded high precision.) Anna clearly expects herself to complete the relevant task, and the task is of sufficient importance to her that the corresponding expectation is able to entrain Anna's actions in the present moment. But the persistent presence of the pain is undermining the chances that Anna will be able to fulfill her expectation. In normal circumstances, Anna might have little problem with the essay. Now, however, the pain (the symptom) is threatening to throw her plans into disarray.

There is another feature of the toothache case that is worth highlighting. This concerns the idea that Anna is, in some sense, 'obliged' to participate in whatever process is responsible for the subjective experience of pain. Anna cannot simply decide whether or not she is in pain. Instead, the pain is something that happens to her. From an ontological standpoint, we can understand this as a form of participatory relationship between Anna and a pain-related process.[11] While the ontic status of pains (and other experiences) remains unclear, we will assume that all subjective experiences are rooted in processes of one sort of another (e.g., neural or bodily processes).[12] Given the earlier characterization of symptoms as subjective experiences, it follows that all symptoms involve the instantiation of processes. These are what we will call *symptom processes*. As long as Anna experiences the pain she is, we suggest, a participant in one of these processes.

There is, of course, another sort of process in which Anna is involved. This is the essay-writing task—the task that Anna is trying to complete. There is, however, an important difference between the essay-writing task and the pain-related symptom process. This difference centers on the extent to which Anna can choose to be a participant in each of the relevant processes. When it comes to the essay-writing task, Anna can exert a degree of control over whether or not she is a participant in the orthographic process. She can, if she so decides, cease to perform the essay-writing task, and this is, indeed, what happens when she decides to go to bed. In this sense, Anna's participation in the essay-writing process is a particular form of participation. It is what we might call *facultative participation*.

It should be relatively clear that Anna's participation in the symptom process is not of the facultative variety. Anna cannot simply choose to suspend her participation in the symptom process when it is time to go to bed. Instead of her participation being of the facultative kind, she is instead obliged to participate in this process. This is what we might call *obligatory participation*.

Obligatory participation is, as far as we can tell, a universal feature of symptom processes. Unlike something that we might be asked to do, we cannot choose to participate in a symptom process. If we are in pain, then we are obliged to endure the pain, and this is the case no matter how much we might wish that we were not in pain. The same is true of any number of other symptoms, such as feelings of nausea or fatigue or dizziness. While we might wish that we were not a participant in these processes, it is not so easy for us to, in effect, simply 'walk away' (or disengage) from them. We cannot, in general, refuse to experience pain, or disconnect ourselves from the experience of pain. And if such forms of disconnection were, in fact, possible, then it is likely that we would experience pains (and other symptoms) as much less burdensome than they so often are.[13]

The toothache case is intended to establish a basic point of connection with the expectation-based account of burden. It is, of course, no more than a toy example, but the general features of the case are applicable to a range of health-related scenarios. According to the expectation-based account, symptoms are burdensome when they interfere with the fulfillment of expectations. In the toothache case, the relevant expectation corresponds to the goal of writing the essay, and the symptom (the pain) is undermining Anna's capacity to fulfill this expectation. Given that we can understand essay-writing activities as, at root, just another sort of process, the toothache case can be reduced to one in which we have one sort of process (a symptom process) that is interfering with another sort of process. It is this interference, we suggest, that lies at the heart of symptom burden. In general, a symptom will be perceived as burdensome when a symptom process (*P1*) interferes with another process (*P2*), where the performance of (*P2*) is required to fulfill one or more (high precision) expectations. The upshot will be an increase in expected prediction error, which reflects the

[11] From an ontological perspective, participation is one of the main relationships between continuants (e.g., human individuals) and occurrents (e.g., processes). Intuitively, continuants are participants in occurrents. For more on the semantics of participation relationships, see Rodrigues and Abel [2019].

[12] Smith et al. [2011], for example, represent pain as a bodily process.

[13] As we will see, this idea of being obliged to participate in processes is one that is common to other types of burden (see Section 4). This makes sense when one considers the general nature of burdens. As we noted in Section 2, a heavy backpack is a burden just so long as one is obliged to carry it. If the backpack can be abandoned at any point (without the risk of violating other expectations), then it is hard to see why it would be regarded as a burden. What makes the backpack burdensome (at least in part) is the fact that we are obliged to carry it.

fact that the probability of achieving a certain sort of outcome has been undermined as the result of one's obligatory participation in a symptom process.

While this way of thinking about symptom burden may seem a little abstract, at least relative to the way symptom burden has been understood in the medical and psychological literature, there is no reason to regard the expectation-based account as incompatible with the earlier definitional efforts. Consider the earlier notion of symptom impact. From an expectation-based perspective, the impact of a symptom is simply the extent to which a symptom interferes with the performance of processes that are required to fulfill expectations. Accordingly, the greater the interference, the greater the impact. Symptoms with minimal impact will be those that do not interfere with our capacity to fulfill expectations. In essence, they will not undermine estimates of our own capacity to bring about certain sensory futures (i.e., our capacity to fulfill high precision expectations courtesy of our own actions). The more a symptom interferes with our capacity to bring about these sensory futures, the greater the impact, and thus the more burdensome the symptom.[14]

As a means of clarifying the link between interference and impact, let us consider two additional cases. These cases are intended to be highly similar. Indeed, the only difference between the cases relates to the use of particular resources to perform an orthographic task.

> **The Keyboard Novelist Case**
> Bob is a left-handed novelist with a mild case of Dupuytren's contracture. Bob does not experience any pain as a result of the condition. He does, however, experience some changes in motor functioning. As it happens, Bob uses a keyboard to support his novel-writing activities, and such activities are not affected by the Dupuytren's contracture. Bob is aware of the Dupuytren's contracture, and he suspects that at some point he may need to bring the condition to the attention of his physician. For the time being, however, the disease does not interfere with Bob's orthographic efforts.

> **The Pen-and-Paper Novelist Case**
> Charlie is a left-handed novelist with a mild case of Dupuytren's contracture. Charlie does not experience any pain as a result of the condition. He does, however, experience some changes in motor functioning. As it happens, Charlie uses pen and paper resources to support his novel-writing activities, and such activities are seriously affected by the Dupuytren's contracture. Unable to write novels with the ease and efficiency that he is accustomed to, Charlie seeks the advice of his physician, requesting that something be done to address the situation.

These two cases are intended to be examples of what are dubbed "thought experiments" [see Frappier et al., 2013]. In essence, they are hypothetical scenarios that are intended to press our intuitions into some sort of useful philosophical service. For present purposes, it should be noted that the two cases are identical in *all* respects, except for the fact that Bob's orthographic efforts involve the use of a conventional computer keyboard, while Charlie's efforts rely on the use of pen and paper resources. This sort of (imaginary) manipulation is akin to what, in other contexts, might be dubbed an 'independent variable'—it is something that we are changing as part of an experimental manipulation while attempting to hold all other things constant. The 'dependent variable' is then the experiences of two protagonists. These experiences, it should be clear, are not the same. Bob is largely unaffected by the disease and its associated symptoms, and he thus sees no reason to consult a physician. Charlie, by contrast, is sufficiently perturbed by the symptoms to seek immediate medical attention. We thus have a difference in the way the two protagonists experience their symptoms. This is despite the fact that both the disease and the symptoms are the same in both cases. What is more, both protagonists are aware of the symptoms—they both recognize that there is a health-related problem of sorts and that this problem is interfering with at least some of their daily activities.

What these cases reveal is that our notions of symptom burden are unlikely to be well-served by a simple focus on the nature of a disease or its associated symptoms. Instead, what seems to be important is the way in which a symptom *interferes* with the sorts of things that an individual expects themselves to do. Again, it is worth emphasizing that the difference between Bob and Charlie is not their expectations, the nature of the disease, or the mere presence of the symptoms. These things, recall, are being held constant across the two cases. What differs between the two cases is the fact that the same symptoms are exerting different *impacts* on the two protagonists. This impact can be understood in precisely the same way we understand the notion of impact in the toothache case; i.e., as a form of interference or disruption with the performance of activities that are required to ensure that a given individual delivers on their (high precision) expectations.[15]

[14] It should also be noted that the expectation-based account speaks to at least some of the ways that patients, themselves, understand the term "symptom burden." In one study, Gill et al. [2012], asked patients to define symptom burden. The most common response—reported by 59% of the patients—was an inability to perform usual activities.

[15] As is suggested by the Charlie case, such interference may be one of the reasons people seek medical assistance. This tallies with the views of Demain et al. [2015, p. 4], who suggest that: "... people seek healthcare not simply to relieve physical or emotional

There is, of course, an important distinction between the toothache case and the Bob/Charlie cases. While the motor impairments associated with Dupuytren's contracture are typically of long duration, the pain associated with a dental infection is (one hopes) a much more transient affair. This durational distinction is important, for it should be relatively clear that the longer a symptom persists, the greater the chances of it interfering with our capacity to fulfill expectations.[16] The symptoms associated with a common cold can, of course, be a source of frustration, especially if they require us to put our lives on hold for a few days. In general, however, few of us would probably gloss these experiences as particularly burdensome. As in the toothache case, we can usually renegotiate the demands that are placed upon us by modifying our own expectations as to what we will do, as well as the expectations others have about us. Matters are not so straightforward if symptoms should persist over the longer-term. While the symptoms associated with an acute COVID-19 infection can be unpleasant but bearable, the symptoms associated with Long-COVID can be much more irksome [e.g., Sykes et al., 2021]. From an expectation-based perspective, we can understand this difference relative to the way in which chronic health conditions threaten to derail a number of longer-term (perhaps lifelong) projects.[17] In the toothache case, we saw how pain might interfere with the performance of a single task—that of writing an essay. But chronic health conditions are apt to interfere with our efforts to construct a much more protracted and personal narrative, one that both reflects and informs our sense of who and what we are as individuals. This is where our understanding of burden begins to connect with the growing interest in predictive processing approaches to the self [Hohwy and Michael, 2017, Friston, 2018, Letheby and Gerrans, 2017, Seth and Tsakiris, 2018]. As noted in earlier work, chronic illness can lead to biographical disruptions that challenge one's sense of self and/or identity [Bury, 1982, Demain et al., 2015]. There is no reason why we cannot connect this to the expectation-based account. Quite plausibly, a biographical disruption represents a disturbance to the sort of things that we expect ourselves to do, challenging our basic conception of who we are. From a predictive processing perspective, the self (understood as a constituent of a neurally-realized generative model—e.g., a deeply nested hidden causal variable) is a potent source of predictions pertaining to future states-of-affairs, especially those that can be brought about as a result of our own actions. A chronic illness threatens to throw such predictive efforts into disarray, forcing us to revise generative models that may have taken a lifetime to acquire. In time, an individual's expectations may adjust to reflect the new 'reality'. There is, however, no reason to think this process of adjustment will be particularly straightforward. Nor is there any reason to think that the process of adjustment (or adaptation) will terminate in a successful conclusion (i.e., the normalization of prediction error).[18] As long as there is a mismatch between the sort of person one expects oneself to be and the sort of person one is allowed to be (as the result of injury or disease), then there will always be a degree of prediction error, and (according to the expectation-based account) a persistent sense of burden.

For the most part, there ought to be little problem in applying the expectation-based account to symptom burden. According to the expectation-based account, symptoms become burdensome when they interfere with our capacity to bring about sensory states-of-affairs that match our expectations—what predictive processing theorists dub target states or prior preferences (see Section 2). The application to symptoms, especially those of the chronic variety, ought to be straightforward, for it is natural to think of symptoms as causing some sort of disruption to the things we would otherwise expect ourselves to do. The persistent presence of fatigue, for example, will invariably have some sort of bearing on our capacity to fulfill a number of expectations, including those that stem from our occupation of certain social roles (e.g., our status as a parent, an employee, or caregiver).[19] Interestingly, however, the expectation-based account is also able to cater for situations in which symptoms are mostly absent. To help us understand this, consider the following case.

> **The Ménière's Disease Case**
> Darya suffers from Ménière's disease, which is a chronic condition affecting the inner ear. For Darya, the primary symptoms are vertigo and nausea. While Darya finds these symptoms to be highly unpleasant, she has learned that they are also somewhat controllable. In particular, Darya

symptoms but because those symptoms stop them from doing what they want (e.g. running or hiking) and being who they want to be (a professional athlete, a member of the rambling club or an optimistic person)."

[16] It should be noted that illness-related burdens are typically studied in the context of *chronic* health conditions. In this sense, the toothache case might be seen as problematic, focusing as it does on an acute health condition. There is, to be sure, an important difference between acute and chronic conditions, but it is not clear that this requires us to advert to a radically different account of burden. From a predictive processing perspective, the difference between acute and chronic conditions may relate to the nature of the expectations that stand to be violated or the persistent nature of expectation-related errors. While these differences are important, they do not mark a substantive difference in the mechanism that is deemed to be responsible for burdensome experiences.

[17] This problem is apt to be compounded by multimorbidities, where multiple diseases (and thus multiple symptoms) are present at the same time. See Willadsen et al. [2021], for more on this.

[18] See Charmaz [1995], for a discussion of some of the challenges posed by long-term health conditions.

[19] This establishes an interesting point of connection with the way that social roles have been understood in the sociological literature. According to one proposal, social roles are defined by expectations that shape behavior in role-appropriate ways [see Biddle, 1979, 1986].

> has learned that the symptoms tend to be triggered by noisy environments, such as nightclubs. She has also learned that she can no longer travel on airplanes, for this also triggers the symptoms. Prior to the onset of the disease, Darya regularly went to nightclubs, and she enjoyed foreign travel. Now, however, she is obliged to refrain from such activities. For much of the time, Darya does not experience either vertigo or nausea, and yet she still feels burdened by her illness.

What makes this case interesting is the idea that symptoms can be a source of burden even though the symptoms, themselves, are seldom experienced. This poses a potential problem for the idea that symptoms (*qua* subjective experiences) are interfering with one's capacity to fulfill expectations. For if a symptom is not being experienced for much of the time, then it must be mostly absent. But if a symptom is absent, then it is hard to see how it could interfere with one's capacity to fulfill expectations.[20]

The puzzle is, of course, illusory, for while Darya succeeds in controlling the occurrence of symptoms, she does so at the expense of doing things that she would otherwise do given the absence of the disease. In this case, then, the prediction error stems from the fact that Darya is engaging in a course of action that violates her expectations pertaining to (e.g.) travel and nightclubs. It is the fact that Darya is obliged to refrain from doing the things she *wants* to do that leads to the experience of burden, for this entails the violation of her (high-precision) expectations. Absent such expectations, and our intuitions about the presence of burden begin to shift. If Darya hated nightclubs and had no interest in airborne travel, then there is no sense in which her symptom-avoiding actions could violate her expectations. In this case, the prediction error would be greatly attenuated, as would (we suggest) the feeling of burden.

4 Treatment Burden

In contrast to symptom burden, treatment burden focuses on the treatments (or other interventions) that are deployed to manage a disease (or, perhaps, to prevent the onset of a disease) [Eton et al., 2022, Hounkpatin et al., 2022, May et al., 2014, Rosbach and Andersen, 2017, Sav et al., 2013a, Tran et al., 2015]. Alsadeh et al. [2020] provide a useful summary of definitional efforts in this area, identifying a total of 16 definitions for treatment burden. A subset of these definitions is as follows:

> [. . .] we defined BOT [burden of treatment] as treatment related effects that limit the patient's ability to participate in activities and tasks that are crucial to his or her quality of life and that are not attributable to underlying disease. [Bohlen et al., 2012, p. 47]

> Treatment burden refers to the personal workload of healthcare, including treatment and self-management of chronic health conditions, and the impact of this workload on patient functioning and well-being. [Eton et al., 2017, p. 490]

> Treatment burden is defined here as a patient's perception of the aggregate weight of the actions and resources they devote to their health care, including difficulty, time, and out-of-pocket costs dedicated to health care tasks such as adhering to medications, dietary recommendations, and self-monitoring. [Boyd et al., 2014, p. 2]

> [Treatment burden is] defined as the impact of the 'work of being a patient' on functioning and well-being. This work includes drug management, self-monitoring, visits to the doctor, laboratory tests, lifestyle changes, and other actions that take place in addition to the other work patients and their caregivers must do as part of life. [Tran et al., 2015, pp. 1–2]

We will not attempt to critique these existing definitional efforts; nor will we seek to single out one definition as being preferable to the others. Rather than subject existing definitions to critical scrutiny, we will attempt to show how the expectation-based account provides a common approach to understanding these definitions. Our aim, in short, is to show how a predictive processing approach to burden might provide an opportunity for theoretical integration, enabling us to accommodate multiple definitions of treatment burden within an overarching theoretical framework.

The starting point for our analysis centers on the notion of patient workload. Workload is a recurring feature of treatment burden definitions. It is generally understood as the work that must be performed by patients as part of a treatment

[20] Note that this is also one of the reasons why we resist the temptation to equate the experience of burden with the aversive quality (or negative valence) of symptoms (what we might call the mere experience view of burden). If burdensome experiences are tied to the mere unpleasantness of symptoms, then the absence of symptoms should yield a correlative reduction in the experience of burden. The Ménière's disease case, however, highlights the way in which a sense of burden might persist, even in the absence of symptoms.

regimen. In the above definitions, for example, we see an appeal to "health care tasks," the "work of being a patient," and the "personal workload of healthcare." The notion of workload thus refers to the various actions and activities that must be performed by the patient as part of a treatment (or, more generally, healthcare) regimen. For the sake of terminological convenience, we will refer to these actions and activities as *treatment actions*.[21] Some examples of treatment actions are evident in the above definitions of treatment burden. They include the likes of taking medication, maintaining medical appointments, monitoring health status, and engaging in physical therapy. A treatment action is, in short, any action that is performed by a patient as a means of treating or managing a health-related problem.[22] Treatment actions may, of course, be performed by individuals other than the patient (e.g., family members or healthcare professionals). In the present context, however, we are interested in the actions performed by patients, as opposed to the actions performed by other individuals—this is how we understand the notion of *patient* workload as opposed to, let's say, the notion of caregiver workload.

Whatever else we might say about treatment actions, they are clearly actions of one sort of another. This establishes a point of contact with the analysis of symptom burden. In Section 3, we suggested that symptoms are tied to symptom processes, and symptom processes are key to our understanding of symptom burden. Symptom processes, it should be clear, are a particular sort of process. This makes symptom processes a member of the metaphysical category of occurrents—a category that includes the likes of processes, states, and events [see Rodrigues and Abel, 2019]. Interestingly, treatment actions also belong to this metaphysical category. Indeed, from an ontological standpoint, we can think of actions as particular types of processes. Actions are what we might call *intentional processes*—they are implemented by agents with the express purposes of achieving some sort of outcome or goal. This marks an important difference between treatment actions and symptom processes. While treatment actions are a form of goal-directed behavior, symptom processes are not. Despite this difference, however, treatment actions and symptom processes are of a similar kind: they are both processes. This raises a question about the extent to which our understanding of symptom burden might be applied to treatment burden. Can we understand treatment burden in more or less the same way as we understood symptom burden?

In order to answer this question, we will need to consider treatment actions from a predictive processing perspective. In particular, we will need to consider the goals towards which treatment actions are oriented. These goals, recall, are just representations of the sensorily defined state-of-affairs that a particular action is intended to bring about (see Section 2). They are what we have been referring to as expectations. In respect of treatment actions, these expectations are what we will call *treatment expectations*.

In one sense, a treatment expectation is just like any other sort of expectation—it is a neurally-realized prediction that works to entrain an individual's actions, as per predictive processing accounts of goal-directed (or motivated) behavior [e.g., Pezzulo et al., 2018]. There are, however, a couple of features of treatment expectations that make them of particular interest in the present context.

The first feature relates to the socially-inflected nature of treatment expectations. Treatment expectations are an example of what we might call *social expectations*—they are expectations that are typically held by multiple individuals, or they are at least expectations that arise as the result of social encounters. In a clinical context, treatment expectations are often established as part of a healthcare encounter involving both a healthcare provider and a patient. A physician, for example, might discuss the details of a treatment plan with a patient, thereby establishing joint or mutually-acknowledged expectations[23] about patient behavior. There is, of course, plenty of room for miscommunication and misunderstanding here. Just because a healthcare provider expects a patient to perform a given treatment action, this does not mean the patient will possess the *same* expectation as that of the healthcare provider. Assuming all goes well, however, the patient and healthcare provider will arrive at a common understanding of what the patient must do to care for their own health. Insofar as the patient adopts the treatment expectation, they will attempt to do what their healthcare provider expects them to do. They will, in short, attempt to fulfill treatment expectations via the performance of treatment actions.

A second feature of treatment expectations relates to their importance or priority. In predictive processing, the importance or priority of an expectation is marked by its precision, with higher precision corresponding to greater importance (see

[21] We acknowledge that the term "treatment activity" may be more appropriate here. For the purposes of the present paper, however, we will overlook the distinction between activities and actions.

[22] Note that the performance of treatment actions may require one to perform a number of ancillary tasks. In order to take one's medication, for example, one must first acquire the medication. This may require the patient to visit a pharmacist and perhaps pay for a prescription. Accordingly, the act of taking medication requires one to participate in additional actions, such as attending a pharmacist (or other healthcare facility) and paying for one's medication. These additional actions add to the workload that is imposed on patients, and they should thus be included in the notion of patient workload.

[23] For more on mutually-acknowledged expectations, see Cibik [2018]. According to Cibik, mutually-acknowledged expectations correspond to what we typically refer to as commitments or obligations.

Section 2). There are a number of reasons to think that treatment expectations will be assigned high precision.[24] A disease or illness represents a departure in the normal (expected) functioning of the body, and this serves as a potent source of prediction error. Pursuing a course of treatment provides a means of resolving this error, or at least resolving some of the uncertainties that come with a clinical diagnosis. Given this, it is reasonable to assume that the majority of treatment expectations will be accorded high precision, at least initially. This does not mean that patients will always do what they are expected to do, of course. Sometimes a patient may reject medical advice, or they may regard treatment actions as unimportant. In such cases, the precision of treatment expectations will be low, and non-adherence is the likely result. Note, however, that non-adherence is not a surefire way of avoiding burdensome experiences. To be sure, non-adherence may be an effective means of reducing treatment burden—if one is not committed to performing treatment actions, then it is hard to see how one could regard those actions as particularly burdensome. The problem is that treatments are intended to treat diseases, and diseases often come with symptoms. Refraining from treatment actions may help to reduce (or even eliminate) treatment burden, but it will not help to minimize one's exposure to symptom burden.

The precision of treatment expectations is important, for it establishes a point of contact with the earlier discussion of symptom burden. Recall that in discussing symptom burden, we encountered two forms of participatory relationship that an individual may have with regard to processes. These are what we referred to as obligatory participation and facultative participation. In the case of symptom processes, we suggested that individuals have no choice but to participate in the relevant process, and for this reason the relevant form of participation is of the obligate variety.

When it comes to treatment actions, the appeal to obligatory participation doesn't really work. The reason for this is that treatment actions are a form of *intentional* process, and it is always possible (at least in principle) for the individual to suspend such actions or refuse to perform them altogether. In this sense, then, we cannot approach an individual's participation in treatment actions in the same way as we approached their participation in symptom processes. Rather than being obligate in nature, an individual's participation in treatment actions is better understood from the perspective of what we called facultative participation.

While the appeal to facultative participation is technically correct, matters need not be so clear-cut from the patient's perspective. If treatment expectations are assigned high precision, the patient will be under considerable pressure to perform treatment actions. A failure to perform these actions will lead to an escalation of expected prediction error, simply because the patient is, in effect, failing to do something that they expect themselves to do. For this reason, high-precision treatment expectations mandate the performance of treatment actions in a way that is not well-served by the notion of facultative participation. The relevant form of participation here is what we might call *effectively obligate*. In principle, there is no reason why a patient is obliged to participate in treatment actions (by performing treatment actions). In practice, however, the options may be somewhat limited. As long as treatment expectations are assigned high precision, then the patient will attempt to fulfill these expectations by performing treatment actions. If such expectations cannot be fulfilled (for whatever reason), then there will be an increase in expected prediction error. Such increases could be avoided if the precision assigned to treatment expectations were to be reduced. In practice, however, such adjustments may be difficult, and, in any case, the failure to fulfill treatment expectations may lead to an overall increase in expected prediction error, especially if non-adherence leads to an accentuation of symptoms (and thus symptom burden).

Once we recognize the constraints associated with treatment expectations, the similarity between treatment actions and symptom processes starts to come into sharper focus. In particular, we can begin to see how the pressure to participate in treatment actions might interfere with the fulfillment of expectations in more or less the same way as that seen in the case of symptom processes. In particular, a treatment action may detract from the performance of other actions, if only because one can only do one thing at a time, and there is only so much time available to do all the things that must be done to ensure one's expectations are fulfilled. If the individual had a choice, they would no doubt prefer to renege on their commitment to performing treatment actions. Given the nature of treatment expectations, however, this is not straightforward. The result is that the individual is effectively stuck with something that 'must' be done, but which they would prefer not to do. In this respect, they are in roughly the same sort of position as the individual who is obliged to endure the pain of a persistent toothache. In both cases, we confront a situation in which an individual is compelled to participate in a certain sort of process (a symptom process or treatment action), but this participation leads to an overall increase in expected prediction error due to the way in which one process interferes with another.

It seems, then, that we can understand treatment burden in more or less the same way as we understood symptom burden. There are, however, some peculiarities of treatment burden that are worth exploring. The following case will help us explore these peculiarities:

[24] See Theriault et al. [2021], for a predictive processing approach to social pressure, what Theriault et al. dub a "sense of should." For an active inference approach to shared expectations and social conformity, see Constant et al. [2019].

The Chronic Pain Case

Ella has been experiencing lower back pain for many months. Deciding that something must be done, she visits a physiotherapist who recommends that she attend weekly appointments and perform a number of exercises at home. Ella begins attending the appointments and performing the exercises. As time passes, she begins to see a gradual improvement in her symptoms.

Despite its simplicity, this case illustrates many of the concepts we have encountered thus far. Ella is, of course, the patient in this scenario; the physiotherapist is the health provider. The recommendations of the physiotherapist correspond to the treatment expectations that are communicated to Ella as part of her healthcare encounter. These treatment expectations then entail the performance of actions, such as attending appointments and performing exercises. These actions are what we have called treatment actions. Assuming that Ella heeds the advice of the physiotherapist, she will adopt the physiotherapist's expectations as her *own* expectations. She will, in short, expect herself to do what her physiotherapist expects her to do. Patient workload then corresponds to the treatment actions that Ella is expected to perform; i.e., the actions that are required to fulfill the treatment expectations. In this case, we are assuming that Ella is in a position to perform the treatment actions and that she does, indeed, perform the actions. This ensures that the treatment expectations are fulfilled. The problem is that the performance of treatment actions can interfere with the performance of other actions, and if this results in a failure to fulfill other expectations, the result will be an increase in expected prediction error, just as was the case with symptom burden.

There ought to be nothing controversial about the mere idea that treatment actions can interfere with the performance of other actions. All actions consume resources of one sort of another, even if it is only the time it takes to perform an action. The time that Ella spends with the physiotherapist may or may not be well-spent; but it is not time that can be re-invested in an alternative course of action. In addition to issues of time, the performance of treatment actions may incur other resource-related costs. Perhaps, for example, Ella is required to pay an hourly rate to the physiotherapist. If so, then the performance of treatment actions comes with both a temporal and financial cost. The problem is that such costs can quickly escalate to the point where other actions become difficult or perhaps impossible to perform.[25] The money that Ella invests in her treatment cannot be reassigned to something else. Money is a fungible resource, but it is not one that is inexhaustible.[26] If Ella must fulfill treatment expectations while failing to fulfill other expectations (e.g., those associated with an occupational or parental role), then she will probably experience the treatment actions as somewhat burdensome. Note that the reason for this is *not* because she is incapable of fulfilling treatment expectations. Nor is it the case that her treatment expectations go unfulfilled. It is more that the fulfillment of one set of expectations (treatment expectations) is coming into conflict with another set of expectations. The upshot is that Ella encounters a form of 'dissonance' or tension that stems from her inability to service multiple goals, and thus fulfill multiple expectations.[27] As with the notion of symptom burden, we arrive at the idea that some processes (in this case, intentional processes or actions) are apt to interfere with other processes, and this degrades the individual's capacity to fulfill multiple expectations. The end result is the same as that seen with symptom burden: an increase in expected prediction error.

In the chronic pain case, we assumed that Ella had no problem performing the relevant treatment actions.[28] There is, however, no reason to assume that treatment actions are always straightforward. In some cases, a patient may struggle to perform a treatment action, and thus 'live up to' treatment expectations. Some treatment actions may require specialist knowledge and expertise, or they may rely on resources that are either absent or difficult to acquire. In such cases, the patient may deem themselves to be incapable of performing treatment actions, or they may at least harbor doubts about their ability to fulfill treatment expectations. This may go some way towards helping us understand the results of studies that have explored the relationship between health literacy, self-efficacy, and treatment burden. According to the results of these studies, higher treatment burden is reported by individuals who score lower in both health literacy and self-efficacy [Eton et al., 2017, 2022, Herzig et al., 2019, Hounkpatin et al., 2022]. Such results make a great deal of sense if we assume that health literacy and self-efficacy are associated with a greater degree of subjective confidence in one's capacity to fulfill treatment expectations. Individuals who are health literate, for example, may have greater certainty about what they need to do to fulfill treatment expectations, while those who are blessed with high self-efficacy are perhaps more inclined to believe they are capable of fulfilling treatment expectations. Given this, it would be

[25] This will obviously depend on how much time and money Ella has to invest in her treatment actions. The more resources one has at one's disposal, the easier it is to perform actions without introducing conflicts or becoming overwhelmed.

[26] For a recent study exploring the temporal costs of treatment actions, see Hounkpatin et al. [2022]. Financial considerations also figure prominently in the treatment burden literature [Eton et al., 2015, 2017, Sav et al., 2013b,a]. The temporal and financial costs of performing treatment actions are sometimes discussed under the heading of temporal burdens and financial burdens.

[27] See Arnaldo et al. [2022, p. 7], for a brief discussion of dissonance within a predictive processing context. Connolly [2018] also provides an interesting account of the role of conflict (and elevated expected free energy) in a Freudian psychoanalytic context.

[28] This is not to say that Ella does not experience any treatment burden; it is merely to note that whatever treatment burden Ella does experience cannot be attributed to the fact that she is failing to fulfill *treatment* expectations.

no surprise to discover that individuals with higher levels of health literacy and self-efficacy are less susceptible to treatment burden—such results are precisely what we would expect if individuals were estimating the likelihood of treatment expectations being fulfilled, courtesy of their own actions.

There may be additional reasons why patients find it difficult to fulfill treatment expectations. In all the cases we have considered thus far, we have assumed that the relevant individual is dealing with a *single* health problem (a dental infection, Dupuytren's contracture, or a bad back). While these cases are not irrelevant when it comes to our understanding of burden in a medical context, they undoubtedly overlook many of the challenges associated with multiple, long-term health conditions. There may, for example, be multiple treatment expectations for multiple diseases. There may also be multiple symptoms. There is also the potential for interactions between treatments and symptoms. Consider, for example, that the treatment actions performed in respect of one treatment expectation may interfere with one's capacity to fulfill another treatment expectation, either as the result of resource constraints (the performance of one action reduces the resources available for another) or because there is an incompatibility between treatments and diseases, such that the performance of a treatment action (for one disease) is rendered problematic due to the symptoms of another disease (e.g., the performance of a treatment action for one disease exacerbates the symptoms of another disease).

By now, it should be clear that the notion of patient workload presents little in the way of a problem for the expectation-based account. As was noted earlier, patient workload is typically understood as the work that must be done by a patient, or the healthcare tasks that 'must' be performed as a result of one's occupation of the patient role. Courtesy of the appeal to treatment expectations (and treatment actions), we are able to see why the performance of healthcare tasks might be a source of burdensome experiences. We are also in a better position to identify the various forces and factors (e.g., resource constraints, pre-existing commitments, and the presence of socially-significant others[29]) that might inform this particular aspect of the patient's lived experience.

Does this mean the analytic effort is complete? One reason to think that it might not be complete is that definitions of treatment burden often feature an appeal to the impacts (or effects) of treatment (in addition to workload). As noted by Eton et al. [2022, p. 3], "[w]orkload and impact represent two fundamental components of treatment burden." In view of this, it seems that a theoretical account of treatment burden ought to do more than just accommodate the notion of patient workload. In addition to workload, it ought to tell us something about the impact of treatments, or, as we shall call it, *treatment impact*. How, then, does the expectation-based account accommodate the notion of treatment impact?

In responding to this challenge, it is worth noting that definitions of treatment burden frequently refer to two sorts of impact, namely, the impact on patient functioning and the impact on patient wellbeing [Eton et al., 2015, 2017, 2022, Hounkpatin et al., 2022, Kyle et al., 2020, Tran et al., 2015]. While these impacts are undoubtedly connected, we will consider them separately.

In respect of patient wellbeing, it is important to note that subjective wellbeing, happiness, and positive affective experiences have all been the focus of recent predictive processing work [Kiverstein and Miller, 2023, Hesp et al., 2020, Miller et al., 2022, Smith et al., 2022c]. For reasons of space, we will not attempt to discuss the details of this work. Suffice to say, however, the expectation-based approach is broadly compatible with a predictive processing approach to wellbeing. Smith et al. [2022c], for example, suggest that subjective wellbeing is tied to the quality of generative models and thus the capacity to minimize prediction error (or free energy) over time. It is perfectly plausible, in our view, that this state-of-affairs is marked by a reduction in expected prediction error. There is, at least, no reason to think a persistent elevation in expected prediction error is compatible with the quality of a generative model. In this sense, then, there is a relatively straightforward link between the expectation-based account and predictive processing approaches to patient wellbeing: if the performance of treatment actions entails a persistent increase in prediction error, then it is plausible that neural estimates pertaining to the quality of a generative model would be downgraded. This, Smith et al. [2022c] suggest, is precisely the sort of situation in which we encounter a decline in subjective wellbeing.

What about patient functioning? The primary problem here relates to the meaning of the term "patient functioning." What does it mean for a patient to function well (or poorly) given the presence of treatment actions? One thing that it could mean for a patient to function well is for the patient to be in a position where they are able to fulfill their expectations, thereby keeping expected prediction error under control. In this case, patient functioning would be high and treatment burden would be low. Conversely, if treatment actions were to interfere with a patient's capacity to fulfill expectations, then patient functioning would be low and treatment burden would be high.

There is, of course, nothing here that is inconsistent with the idea that treatment burden is, in part, constituted by treatment impact. At the same time, however, the expectation-based account tends to paper over the distinction between patient workload and treatment impact, especially when it comes to the notion of patient functioning. According to

[29]Of particular interest is the way that social proximity moderates the perception of physiological and cognitive effort, given baseline expectations of social contact. See Beckes and Sbarra [2022], for more on this.

the expectation-based account, if a treatment action causes an elevation in expected prediction error then it will be perceived as burdensome. But the only way a treatment action could cause an increase in expected prediction is if it were to reduce patient functioning. For all that patient functioning seems to mean (at least from a predictive processing perspective) is that there is some sort of interference with a patient's capacity to fulfill their expectations.

5 Side-Effects Burden

While symptom burden and treatment burden are typically seen as different types of burden [Sav et al., 2013a], there is one area where this distinction might be called into question. This concerns the side-effects produced by certain treatments; e.g., those stemming from medications. For the most part, side-effects are studied as part of research into treatment burden, and it is thus tempting to view side-effects as a component of treatment burden, or perhaps a specialized form of treatment burden. This is, at least, how side-effects have been understood by those investigating treatment burden [Demain et al., 2015, Eton et al., 2013, Sav et al., 2013b]. Sav et al. [2013b], for example, identify medication burden as one of the "components" of treatment burden. Medication burden, they suggest, stems from the "side effects and adverse events from medication use, polypharmacy (multiple medication use), the inconvenience of organising medications and the stigma associated with taking medication" [Sav et al., 2013b, p. 671].

There is, of course, an important (causal) link between side effects and treatments [see Due, 2023]. But should side-effects be understood as a component of treatment burden, or, perhaps, a particular sort of treatment burden? One reason to doubt this stems from work into symptom burden. Gapstur [2007], for example, identifies a particular form of burden, dubbed *side-effects burden* [see also Iversen et al., 2018]. According to Gapstur [2007, p. 677], side-effects burden refers to the "burden of treatment-related symptoms." What is interesting about this characterization is not so much the connection between side-effects and treatments; it is more the idea that side effects can be understood as a particular sort of *symptom* (i.e., a treatment-related symptom). If this is correct, then there seems to be no good reason why side-effects burden could not qualify as a specialized type of symptom burden.

While there are reasons to doubt the status of side-effects as symptoms (see below), it is interesting to note that side-effects burden is more easily understood from the standpoint of symptom burden. The reason for this is that side-effects have much in common with symptoms. There is clearly a degree of correspondence between the fatigue/nausea produced by a certain drug, and the fatigue/nausea produced by a certain disease. Indeed, from a purely phenomenological standpoint, these experiences might be indistinguishable: the fatigue we feel as the result of taking a drug may be much the same as the fatigue we feel as the result of a disease. Given this, it would not be particularly surprising if we were to discover that side-effects burden could be understood in more or less the same way as we understood symptom burden.

As before, it will help to orient the discussion around a particular case:

> **The Side-Effect Case**
> Flynn is a business analyst who works in a high-paced, competitive environment. Flynn was recently diagnosed with depression and has just begun a course of (oral) antidepressant medication. Unfortunately, Flynn is experiencing a number of side-effects as a result of taking the medication. In addition to a mild, albeit persistent, sense of fatigue, he feels that he is no longer able to think clearly at work. Given the importance that Flynn assigns to his work, he wonders whether he might be better off abandoning the medication.

By now, it should be relatively clear why the side-effects in this case might be the source of a burdensome experience. Flynn deems the side-effects to be interfering with his work, and thus the side-effects are interfering with Flynn's (perceived) capacity to fulfill the expectations associated with his occupational role. As with symptoms, we can understand side-effects as being tied to processes in which Flynn is a participant. What is more, the nature of Flynn's participation in these processes is of the obligate variety, just as is the case with symptoms. Accordingly, we can draw on our existing understanding of symptom burden to support our understanding of side-effects burden. Side effects are burdensome in precisely the same way that symptoms are burdensome—they are burdensome because they reduce the chances that a given individual will be in a position to fulfill the expectations that the individual deems to be important.

None of this means that side-effects ought to be understood as symptoms, of course. Nor does it mean that side-effects burden ought to be understood as a specialized type of symptom burden. Notwithstanding the phenomenological similarity between side-effects and symptoms, there are reasons to think that side effects are conceptually distinct from symptoms. While symptoms are tied to diseases, side-effects are tied to treatments. This difference, it should be clear, centers on the way in which an ostensibly similar (perhaps identical) symptom-like process is brought into existence as the result of certain causal relationships. Side-effects are caused by deliberate attempts to treat a disease. They stem from actions that are intended to ameliorate symptoms, retard the progression of a disease process, or (in the best cases) terminate the disease process and return the individual to a state of health. It is thus the performance of what we have

called treatment actions that distinguishes side-effects from symptoms: Side-effects are caused by treatment actions; symptoms, by contrast, are not.

Does this mean that side-effects burden ought to be subsumed under the heading of treatment burdens, as opposed to symptom burden? The answer to this question is, unfortunately, unclear. From a predictive processing perspective, we can understand the burdensomeness of side-effects relative to an increase in expected prediction error. By itself, however, this increase does not tell us anything about the taxonomic status of a burden. The burden in question could be a symptom burden, a treatment burden, or a side-effects burden. The difference between these burden types probably depends on the thing that is responsible (in a causal sense) for the increase in expected prediction error. In the case of symptom burden, the cause is a symptom (or symptoms); in the case of treatment burden, it is a treatment action (or actions); and, in the case of side-effects burden it is a side-effect (or side-effects). If this is how we are to distinguish between different types of burden, then side-effects burden is quite plausibly neither a symptom burden nor a treatment burden; it is, perhaps, a distinct type of burden. This makes sense given the features of the side-effects case. Consider, for example, that the treatment action in the side-effects case is the taking of the medication (i.e., the swallowing of a pill). Flynn performs this action so as to fulfill a particular treatment expectation. But there is no reason to think that the mere performance of the action should be particularly problematic for Flynn. There is, in particular, no reason to think that Flynn is particularly challenged by the treatment action, that he lacks the capacity to perform this action, or that the treatment action is contributing to a substantive depletion of available material, temporal, and financial resources. On the whole, the effort associated with the relevant treatment action seems to be minimal, the resources all seem to be available, and there is no sense in which Flynn is incapable of doing what his physician expects him to do. In this sense, then, it is somewhat hard to see why the side-effects case would qualify as a *bona fide* case of treatment burden. There is, of course, a sense in which the relevant treatment action is part of the causal chain that leads to an increase in expected prediction error: It is by taking the pill that Flynn comes to feel the fatigue, and it is the fatigue that is then responsible for the perceived impairment in Flynn's daily functioning. The problem here is that the fatigue (*qua* side-effect) is the more proximal (or direct) cause of this impairment. The pill-taking action, by contrast, is further up the causal chain. We could, perhaps, expand the notion of treatment burden to include more indirect forms of causation, thereby subsuming side-effects burden within treatment burden. The problem with this, however, is that treatments can lead to iatrogenic diseases, which come with their own symptoms (and thus the potential for symptom burden). Accordingly, if we allow the concept of treatment burden to include downstream causal effects, then we need to be sure that these effects are limited to a certain sort of effect. We cannot allow such effects to include the triggering of disease processes and thus the triggering of symptoms processes; for this would risk obscuring the seemingly crisp conceptual distinction between treatment burden and symptom burden.

How might an individual respond to the presence of burdensome side-effects? One answer is suggested by the side-effect case: If Flynn deems the current situation to be unsustainable, then he may decide to abandon the course of treatment, resulting in non-adherence. In this respect, a consideration of side-effects burden is likely to be relevant to our understanding of the forces and factors that contribute to adherence-related decisions.[30] From a predictive processing perspective, it is possible that non-adherence stems from changes to the precision assigned to treatment expectations. Such adjustments may arise as the result of efforts to minimize one's exposure to burdens (either of the side-effect or treatment variety).[31] By reducing the precision assigned to treatment expectations, the importance of those expectations is diminished, and thus the failure to fulfill treatment expectations is of little consequence (at least in regards to prediction error). In one sense, then, non-adherence is a perfectly rational (or at least understandable) response to burden—it reflects the efforts of the patient to curtail heightened levels of expected prediction error. This looks to be particularly important when one considers the purported health-related effects of sustained increases in prediction error [Arnaldo et al., 2022, Barrett and Simmons, 2015, Peters et al., 2017, Smith et al., 2019].

While adjustments to the precision of treatment expectations may culminate in non-adherence, it is also possible to imagine situations in which an otherwise low-precision treatment expectation is preserved courtesy of the particular social environment in which an individual is embedded. Imagine, for the sake of example, that Flynn is in a long-term romantic relationship, and that his partner insists that he continue with the course of treatment. Insofar as Flynn is inclined to adopt his partner's expectations as his own expectations—that is to say, he does what his partner expects him to do—then he will have an additional incentive to stay the treatment-related course. The reason for this is that the costs of non-adherence (understood as a failure to fulfill expectations) have increased due to the presence of a social

[30] Interestingly, recent work has sought to understand adherence-related decisions from a predictive processing perspective [Smith et al., 2021]. The side-effect case is, in fact, based on the same situation as that studied by Smith et al. [2021]—a situation in which (virtual) patients are exposed to the side-effects associated with antidepressant medication.

[31] See Hartwig et al. [2022], for a discussion of how changes in precision-related adjustments may be relevant to our understanding of stress habituation.

expectation, namely, that of Flynn's partner. If Flynn strives to fulfill his partner's expectation, then he will persist with the course of treatment. If, by contrast, he is alone, then the story may have a very different ending.[32]

6 Conclusion

In the present paper, we sought to expand our understanding of burdens in a healthcare context by drawing on the predictive processing framework. The result of this analysis was a theoretical account, dubbed the expectation-based account. According to this account, burdens are tied to an increase in expected prediction error, which arises as the result of a mismatch between high precision expectations (variously known as goals, desires, target priors, prior preferences, etc.) and one's (subjective/perceived) capacity to fulfill these expectations. In effect, an elevation in expected prediction error indexes a perceived mismatch between a desired future (all expectations fulfilled) and an attainable future (one or more expectations violated). In a healthcare context, these shifts in expected prediction error can occur as the result of an individual's (obligatory) participation in processes that are either connected to a disease (e.g., symptom processes) or processes that are intended to treat/manage a disease (e.g., treatment actions).

One of the virtues of the expectation-based account is that it provides a common approach to understanding burden. Thus, while symptom, treatment, and side-effects burden have (for the most part) been understood as distinct kinds of burden, each deserving of separate theoretical and empirical attention, our analysis suggests that all these burdens can be understood in more or less the same way; i.e., as a form of interference with one's capacity to fulfill expectations. In this sense, the expectation-based account provides an important opportunity for theoretical integration. At the same time, however, the expectation-based account does not require us to lose sight of the fact that there are multiple types of burden. We can thus continue to recognize the status of (e.g.) symptom and treatment burden as distinct kinds of burden. What is more, there is no reason to regard the expectation-based account as undermining existing efforts to conceptualize burdens. The notion of treatment burden, for example, can still be understood as a form of disruption that reaches across the biological, biographical, and relational domains [see Demain et al., 2015]. Rather than seeking to replace or supplant existing conceptual or definitional efforts, the expectation-based approach is best seen as an attempt to make sense of these efforts—to tell us something about the common features of burden-related experiences, or at least the situations in which these experiences are apt to arise.

The expectation-based account is intended as the basis for future theoretical and empirical work. Some parts of the account are clearly speculative, in the sense that they equate burdensome experiences with a hypothetical increase in expected prediction error. There is clearly more work to be done to map this notion of expected prediction error to the various parameters of the predictive processing framework (especially those that feature as part of research in active inference). There is also much work to be done to connect the present account with existing bodies of predictive processing work, especially in respect of the following areas:

- social inequality and health outcomes [Kelly et al., 2021]
- subjective well-being and happiness [Miller et al., 2022, Smith et al., 2022c]
- biopsychosocial approaches to disease [Smith et al., 2019]
- stress [Arnaldo et al., 2022, Hartwig et al., 2022, Peters et al., 2017]
- fatigue [Greenhouse-Tucknott et al., 2022, Stephan et al., 2016]
- anxiety [McGovern et al., 2022]
- self-efficacy [Bottemanne and Friston, 2021, Stephan et al., 2016]
- medical non-adherence [Smith et al., 2021]
- symptom perception [Kube et al., 2020, Maisto et al., 2021, Van den Bergh et al., 2017]
- subjective effort [Parr et al., 2023, Zénon et al., 2019]
- mental distress [Van de Cruys and Van Dessel, 2021]

Finally, the scope of the present analysis is limited to a consideration of three types of burden, namely, symptom, treatment, and side-effects burden. In future work, it will be important to consider a more expanded array of burden

[32] Note that the primary outcome, here, is one of adherence or non-adherence; it is *not* the attenuation of treatment burden. As noted in earlier research, social relationships may be something of a 'double-edged sword' when it comes to matters of treatment burden [Boehmer et al., 2016, Eton et al., 2022]. In some cases, social relationships may help to ameliorate burdens, but every social relationship comes with certain expectations, and there is no reason why these expectations cannot, on occasion, work to accentuate burden-related experiences.

types, including, for example, the burden experienced by informal and formal caregivers, what is often referred to as caregiver burden [Adelman et al., 2014, Bastawrous, 2013, Liu et al., 2020].

Notwithstanding the scope and scale of this future work, we hope to have convinced the reader that a predictive processing approach to burden is at least deserving of further scrutiny. In addition to providing a common approach to understanding burden, the present account opens the door to computational models that seek to understand the forces and factors that influence burdensome experiences. It also provides us with an important opportunity to connect the study of burdens with the burgeoning work on predictive processing approaches to stress, symptom perception, social inequality, subjective well-being, mental health, fatigue, and disease etiology. The further exploration of these links may help to both refine the present account and connect the study of burden to a wider array of clinical and health-related phenomena.

Acknowledgments

This paper is independent research funded by the National Institute for Health Research Artificial Intelligence for Multiple Long-Term Conditions (AIM) programme (Multidisciplinary Ecosystem to study Lifecourse Determinants and Prevention of Early-onset Burdensome Multimorbidity (MELD-B), award identifier NIHR203988). The views expressed in this publication are those of the author(s) and not necessarily those of the NHS, the National Institute for Health Research or the Department of Health and Social Care.

References

Andy Clark. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3):181–204, 2013.

Andy Clark. *Surfing Uncertainty: Prediction, Action and the Embodied Mind*. Oxford University Press, New York, New York, USA, 2016.

Karl Friston. The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010.

Jakob Hohwy. *The Predictive Mind*. Oxford University Press, Oxford, UK, 2013.

Andy Clark. Consciousness as generative entanglement. *The Journal of Philosophy*, 116(12):645–662, 2019.

Andy Clark. *The Experience Machine: How Our Minds Predict and Shape Reality*. Pantheon Books, New York, New York, USA, 2023.

Jakob Hohwy and Anil Seth. Predictive processing as a systematic basis for identifying the neural correlates of consciousness. *Philosophy and the Mind Sciences*, 1(II):1–34, 2020.

Philipp Schwartenbeck, Thomas H B FitzGerald, Christoph Mathys, Ray Dolan, and Karl Friston. The dopaminergic midbrain encodes the expected certainty about desired outcomes. *Cerebral Cortex*, 25(10):3434–3445, 2015.

Thomas Parr, Rajeev Vijay Rikhye, Michael M Halassa, and Karl J Friston. Prefrontal computation as active inference. *Cerebral Cortex*, 30(2):682–695, 2020.

Andre M Bastos, W Martin Usrey, Rick A Adams, George R Mangun, Pascal Fries, and Karl J Friston. Canonical microcircuits for predictive coding. *Neuron*, 76(4):695–711, 2012.

Stewart Shipp. Neural elements for predictive coding. *Frontiers in Psychology*, 7(Article 1792):1–21, 2016.

Christian Büchel, Stephan Geuter, Christian Sprenger, and Falk Eippert. Placebo analgesia: A predictive coding perspective. *Neuron*, 81(6):1223–1239, 2014.

Giulio Ongaro and Ted J Kaptchuk. Symptom perception, placebo effects, and the Bayesian brain. *Pain*, 160(1):1–4, 2019.

Ryan Smith, Sahib S Khalsa, and Martin P Paulus. An active inference approach to dissecting reasons for nonadherence to antidepressants. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, 6(9):919–934, 2021.

Giovanni Pezzulo, Domenico Maisto, Laura Barca, and Omer Van den Bergh. Symptom perception from a predictive processing perspective. *Clinical Psychology in Europe*, 1(4):1–14, 2019.

Omer Van den Bergh, Michael Witthöft, Sibylle Petersen, and Richard J Brown. Symptoms and the body: Taking the inferential leap. *Neuroscience & Biobehavioral Reviews*, 74(Part A):185–203, 2017.

Irene Arnaldo, Andrew W Corcoran, Karl J Friston, and Maxwell JD Ramstead. Stress and its sequelae: An active inference account of the etiological pathway from allostatic overload to depression. *Neuroscience & Biobehavioral Reviews*, 135(Article 104590):1–16, 2022.

Regina E Fabry. Into the dark room: a predictive processing account of major depressive disorder. *Phenomenology and the Cognitive Sciences*, 19(4):685–704, 2020.

Klaas E Stephan, Zina M Manjaly, Christoph D Mathys, Lilian A E Weber, Saee Paliwal, Tim Gard, Marc Tittgemeyer, Stephen M Fleming, Helene Haker, and Anil K Seth. Allostatic self-efficacy: A metacognitive theory of dyshomeostasis-induced fatigue and depression. *Frontiers in Human Neuroscience*, 10(550):1–27, 2016.

Julian Kiverstein, Erik Rietveld, Heleen A Slagter, and Damiaan Denys. Obsessive compulsive disorder: A pathology of self-confidence? *Trends in Cognitive Sciences*, 23(5):369–372, 2019.

P Read Montague, Raymond J Dolan, Karl J Friston, and Peter Dayan. Computational psychiatry. *Trends in Cognitive Sciences*, 16(1):72–80, 2012.

Jakob Hohwy. Prediction error minimization, mental and developmental disorder, and statistical theories of consciousness. In Rocco J Gennaro, editor, *Disturbed Consciousness: New Essays on Psychopathology and Theories of Consciousness*, pages 293–324. MIT Press, Cambridge, Massachusetts, USA, 2015.

Ryan Smith, Karen L Weihs, Anna Alkozei, William D S Killgore, and Richard D Lane. An embodied neurocomputational framework for organically integrating biopsychosocial processes: An application to the role of social support in health and disease. *Psychosomatic Medicine*, 81(2):125–145, 2019.

Matthew Burch. Phenomenology's place in the philosophy of medicine. *Theoretical Medicine and Bioethics*, 44: 209–227, 2023.

Havi Carel. Phenomenology and its application in medicine. *Theoretical Medicine and Bioethics*, 32:33–46, 2011.

Havi Carel. *Phenomenology of Illness*. Oxford University Press, Oxford, UK, 2016.

Jean Downes. Illness in the chronic disease family. *American Journal of Public Health and the Nations Health*, 32(6): 589–600, 1942.

Roxanna Gapstur. Symptom burden: A concept analysis and implications for oncology nurses. *Oncology Nursing Forum*, 34(3):673–680, 2007.

Daniel E Spratt, Neal Shore, Oliver Sartor, Dana Rathkopf, and Kara Olivier. Treating the patient and not just the cancer: Therapeutic burden in prostate cancer. *Prostate Cancer and Prostatic Diseases*, 24(3):647–661, 2021.

Krista Bohlen, Elizabeth Scoville, Nathan D Shippee, Carl R May, and Victor M Montori. A videographic analysis of how patients with type 2 diabetes and clinicians articulate and address treatment burden during clinical encounters. *Diabetes Care*, 35(1):47–49, 2012.

David T Eton, Tarig A Elraiyah, Kathleen J Yost, Jennifer L Ridgeway, Anna Johnson, Jason S Egginton, Rebecca J Mullan, Mohammad Hassan Murad, Patricia J Erwin, and Victor M Montori. A systematic review of patient-reported measures of burden of treatment in three chronic diseases. *Patient Related Outcome Measures*, 4:7–20, 2013.

Evette J Ludman, Wayne Katon, Joan Russo, Michael Von Korff, Gregory Simon, Paul Ciechanowski, Elizabeth Lin, Terry Bush, Edward Walker, and Bessie Young. Depression and diabetes symptom burden. *General Hospital Psychiatry*, 26(6):430–436, 2004.

Katie I Gallacher, Carl R May, Peter Langhorne, and Frances S Mair. A conceptual model of treatment burden and patient capacity in stroke. *BMC Family Practice*, 19(1):1–15, 2018.

Trude Seselie Jahr Iversen, Nils Eiel Steen, Ingrid Dieset, Sigrun Hope, Ragni Mørch, Erlend Strand Gardsjord, Kjetil Nordbø Jørgensen, Ingrid Melle, Ole A Andreassen, Espen Molden, and Erik G Jönsson. Side effect burden of antipsychotic drugs in real life—impact of gender and polypharmacy. *Progress in Neuro-Psychopharmacology and Biological Psychiatry*, 82:263–271, 2018.

Dominic L Sykes, Luke Holdsworth, Nadia Jawad, Pumali Gunasekera, Alyn H Morice, and Michael G Crooks. Post-COVID-19 symptom burden: What is Long-COVID and how should we manage it? *Lung*, 199:113–119, 2021.

Thomas Parr, Giovanni Pezzulo, and Karl J Friston. *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. The MIT Press, Cambridge, Massachusetts, USA, 2022.

Paul B Badcock, Karl J Friston, and Maxwell J D Ramstead. The hierarchically mechanistic mind: A free-energy formulation of the human psyche. *Physics of Life Reviews*, 31:104–121, 2019.

Julian Kiverstein, Michael D Kirchhoff, and Mick Thacker. An embodied predictive processing theory of pain experience. *Review of Philosophy and Psychology*, 13:973–998, 2022.

Karl Friston, Thomas FitzGerald, Francesco Rigoli, Philipp Schwartenbeck, and Giovanni Pezzulo. Active inference: A process theory. *Neural Computation*, 29(1):1–49, 2017.

Adam Linson, Andy Clark, Subramanian Ramamoorthy, and Karl Friston. The active inference approach to ecological perception: General information dynamics for natural and artificial embodied cognition. *Frontiers in Robotics and AI*, 5(21):1–22, 2018.

Ryan Smith, Maxwell J D Ramstead, and Alex Kiefer. Active inference models do not contradict folk psychology. *Synthese*, 200(Article 81):1–37, 2022a.

Andy Clark. Extending the predictive mind. *Australasian Journal of Philosophy*, pages 1–12, in press.

Achim Peters, Bruce S McEwen, and Karl Friston. Uncertainty and stress: Why it causes diseases and how it is mastered by the brain. *Progress in Neurobiology*, 156:164–188, 2017.

Ryan Smith, Karl J Friston, and Christopher J Whyte. A step-by-step tutorial on active inference and its application to empirical data. *Journal of Mathematical Psychology*, 107(Article 102632):1–60, 2022b.

Julian Kiverstein and Mark Miller. Experiencing well-being playfulness and the meaningful life: an active inference perspective. *Neuroscience of Consciousness*, 2023(1):1–10, 2023.

Sander Van de Cruys and Pieter Van Dessel. Mental distress through the prism of predictive processing theory. *Current Opinion in Psychology*, 41:107–112, 2021.

Karl Friston, Philipp Schwartenbeck, Thomas FitzGerald, Michael Moutoussis, Timothy Behrens, and Raymond J Dolan. The anatomy of choice: active inference and agency. *Frontiers in Human Neuroscience*, 7(Article 598):1–18, 2013.

Sander Van de Cruys, Karl Friston, and Andy Clark. Controlled optimism: Reply to Sun and Firestone on the dark room problem. *Trends In Cognitive Sciences*, 24(9):680–681, 2020.

Karl Friston, Francesco Rigoli, Dimitri Ognibene, Christoph Mathys, Thomas Fitzgerald, and Giovanni Pezzulo. Active inference and epistemic value. *Cognitive Neuroscience*, 6(4):187–214, 2015.

Andy Clark. Beyond desire? Agency, choice, and the predictive mind. *Australasian Journal of Philosophy*, 98(1):1–15, 2020.

Giovanni Pezzulo, Francesco Rigoli, and Karl Friston. Active inference, homeostatic regulation and adaptive behavioural control. *Progress in Neurobiology*, 134:17–35, 2015.

Giovanni Pezzulo, Francesco Rigoli, and Karl J Friston. Hierarchical active inference: A theory of motivated control. *Trends In Cognitive Sciences*, 22(4):294–306, 2018.

Ryan Smith, Lav R Varshney, Susumu Nagayama, Masahiro Kazama, Takuya Kitagawa, and Yoshiki Ishikawa. A computational neuroscience perspective on subjective wellbeing within the active inference framework. *International Journal of Wellbeing*, 12(4):102–131, 2022c.

Thomas Parr and Karl J Friston. Uncertainty, epistemics and active inference. *Journal of The Royal Society Interface*, 14(136):20170376, 2017.

Simon Noah Etkind, Jiaqi Li, John Louca, Sarah A Hopkins, Isla Kuhn, Anna Spathis, and Stephen I G Barclay. Total uncertainty: a systematic review and thematic synthesis of experiences of uncertainty in older people with advanced multimorbidity, their informal carers and health professionals. *Age and Ageing*, 51(8):1–11, 2022.

Charles S Cleeland. Symptom burden: multiple symptoms and their impact as patient-reported outcomes. *Journal of the National Cancer Institute Monographs*, 2007(37):16–21, 2007.

Ashlinder Gill, Anita Chakraborty, and Debbie Selby. What is symptom burden: A qualitative exploration of patient definitions. *Journal of Palliative Care*, 28(2):83–89, 2012.

Alexander P Cox, Patrick L Ray, Mark Jensen, and Alexander D Diehl. Defining 'sign' and 'symptom'. In Selja Seppälä, Patrick Ray, and Alan Ruttenberg, editors, *The Second International Workshop on Definitions in Ontologies*, volume 1309, pages 42–48, Houston, Texas, USA, 2014. CEUR-WS.org.

Richard H Scheuermann, Werner Ceusters, and Barry Smith. Toward an ontological treatment of disease and diagnosis. *Summit on Translational Bioinformatics*, 2009:116–120, 2009.

Marylin Dodd, Susan Janson, Noreen Facione, Julia Faucett, Erika S Froelicher, Janice Humphreys, Kathryn Lee, Christine Miaskowski, Kathleen Puntillo, and Sally Rankin. Advancing the science of symptom management. *Journal of Advanced Nursing*, 33(5):668–676, 2001.

Peter Carruthers. Valence and value. *Philosophy and Phenomenological Research*, 97(3):658–680, 2018.

Pablo Fernandez Velasco and Slawa Loev. Affective experience in the predictive mind: a review and new integrative account. *Synthese*, 198(11):10847–10882, 2021.

Casper Hesp, Ryan Smith, Thomas Parr, Micah Allen, Karl J Friston, and Maxwell JD Ramstead. Deeply felt affect: The emergence of valence in deep active inference. *Neural Computation*, 33(2):398–446, 2020.

Mateus Joffily and Giorgio Coricelli. Emotional valence and the free-energy principle. *PLoS Computational Biology*, 9 (6):e1003094, 2013.

Mark Miller, Julian Kiverstein, and Erik Rietveld. The predictive dynamics of happiness and well-being. *Emotion Review*, 14(1):15–30, 2022.

Adem Sav, Michelle A King, Jennifer A Whitty, Elizabeth Kendall, Sara S McMillan, Fiona Kelly, Beth Hunter, and Amanda J Wheeler. Burden of treatment for chronic illness: a concept analysis and review of the literature. *Health Expectations*, 18(3):312–324, 2013a.

Philip Gerrans. Pain asymbolia as depersonalization for pain experience. An interoceptive active inference account. *Frontiers in Psychology*, 11(Article 523710):1–10, 2020.

Colin Klein. What pain asymbolia really shows. *Mind*, 124(494):493–516, 2015.

Fabrício Henrique Rodrigues and Mara Abel. What to consider about events: A survey on the ontology of occurrents. *Applied Ontology*, 14(4):343–378, 2019.

Barry Smith, Werner Ceusters, Louis J Goldberg, and Richard Ohrbach. Towards an ontology of pain. In Okada Mitsuhiro, editor, *Proceedings of the Conference on Ontology and Analytical Metaphysics*, pages 23–36, Tokyo, Japan, 2011. Keio University Press.

Mélanie Frappier, Letitia Meynell, and James Robert Brown, editors. *Thought Experiments in Philosophy, Science, and the Arts*. Routledge, New York, New York, USA, 2013.

Sara Demain, Ana-Carolina Goncalves, Carlos Areia, Ruben Oliveira, Ana Jorge Marcos, Alda Marques, Ranj Parmar, and Katherine Hunt. Living with, managing and minimising treatment burden in long term conditions: A systematic review of qualitative research. *PLoS One*, 10(5):e0125457, 2015.

Tora Grauers Willadsen, Volkert Siersma, Dagny Ros Nicolaisdottir, Dorte Jarbol, Ann Dorrit Guassora, Susanne Reventlow, and Rasmus Køster-Rasmussen. Symptom burden in multimorbidity: a population-based combined questionnaire and registry study from Denmark. *BMJ Open*, 11(4):e041877, 2021.

Jakob Hohwy and John Michael. Why should any body have a self? In Frédérique de Vignemont and Adrian J T Alsmith, editors, *The Subject's Matter: Self-Consciousness and the Body*, pages 363–391. MIT Press, Cambridge, Massachusetts, USA, 2017.

Karl Friston. Am I Self-Conscious? (Or Does Self-Organization Entail Self-Consciousness?). *Frontiers in Psychology*, 9(Article 579):1–10, 2018.

Chris Letheby and Philip Gerrans. Self unbound: ego dissolution in psychedelic experience. *Neuroscience of Consciousness*, 2017(1):1–11, 2017.

Anil K Seth and Manos Tsakiris. Being a beast machine: The somatic basis of selfhood. *Trends in Cognitive Sciences*, 22(11):969–981, 2018.

Michael Bury. Chronic illness as biographical disruption. *Sociology of Health & Illness*, 4(2):167–182, 1982.

Kathy Charmaz. The body, identity, and self: Adapting to impairment. *The Sociological Quarterly*, 36(4):657–680, 1995.

Bruce J Biddle. *Role Theory: Expectations, Identities, and Behaviors*. Academic Press, New York, New York, USA, 1979.

Bruce J Biddle. Recent developments in role theory. *Annual Review of Sociology*, 12(1):67–92, 1986.

David T Eton, Roger T Anderson, Jennifer L St Sauver, Elizabeth A Rogers, Mark Linzer, and Minji K Lee. Longitudinal trajectories of treatment burden: A prospective survey study of adults living with multiple chronic conditions in the midwestern United States. *Journal of Multimorbidity and Comorbidity*, 12:1–14, 2022.

Hilda O Hounkpatin, Paul Roderick, Scott Harris, James E Morris, Dianna Smith, Bronagh Walsh, Helen C Roberts, Hajira Dambha-Miller, Qian Yue Tan, and Forbes Watson. Change in treatment burden among people with multimorbidity: a follow-up survey. *British Journal of General Practice*, 72(724):e816–e824, 2022.

Carl R May, David T Eton, Kasey Boehmer, Katie Gallacher, Katherine Hunt, Sara MacDonald, Frances S Mair, Christine M May, Victor M Montori, and Alison Richardson. Rethinking the patient: using Burden of Treatment Theory to understand the changing dynamics of illness. *BMC Health Services Research*, 14(1):1–11, 2014.

Michael Rosbach and John Sahl Andersen. Patient-experienced burden of treatment in patients with multimorbidity—A systematic review of qualitative data. *PLoS One*, 12(6):e0179916, 2017.

Viet-Thi Tran, Caroline Barnes, Victor M Montori, Bruno Falissard, and Philippe Ravaud. Taxonomy of the burden of treatment: a multi-country web-based qualitative study of patients with chronic conditions. *BMC Medicine*, 13(1): 1–15, 2015.

Ahmed Alsadah, Tiny van Merode, Riyadh Alshammari, and Jos Kleijnen. A systematic literature review looking for the definition of treatment burden. *Heliyon*, 6(4):e03641, 2020.

David T Eton, Kathleen J Yost, Jin-shei Lai, Jennifer L Ridgeway, Jason S Egginton, Jordan K Rosedahl, Mark Linzer, Deborah H Boehm, Azra Thakur, and Sara Poplau. Development and validation of the Patient Experience with Treatment and Self-management (PETS): A patient-reported measure of treatment burden. *Quality of Life Research*, 26(2):489–503, 2017.

Cynthia M Boyd, Jennifer L Wolff, Erin Giovannetti, Lisa Reider, Carlos Weiss, Qian-li Xue, Bruce Leff, Chad Boult, Travonia Hughes, and Cynthia Rand. Health care task difficulty among older adults with multimorbidity. *Medical Care*, 52(03):S118–S125, 2014.

Matej Cibik. Expectations and obligations. *Ethical Theory and Moral Practice*, 21(5):1079–1090, 2018.

Jordan E Theriault, Liane Young, and Lisa Feldman Barrett. The sense of should: A biologically-based framework for modeling social pressure. *Physics of Life Reviews*, 36:100–136, 2021.

Axel Constant, Maxwell J D Ramstead, Samuel P L Veissière, and Karl Friston. Regimes of expectations: An active inference model of social conformity and decision making. *Frontiers in Psychology*, 10(Article 679):1–15, 2019.

David T Eton, Jennifer L Ridgeway, Jason S Egginton, Kristina Tiedje, Mark Linzer, Deborah H Boehm, Sara Poplau, Djenane Ramalho de Oliveira, Laura Odell, and Victor M Montori. Finalizing a measurement framework for the burden of treatment in complex patients with chronic conditions. *Patient Related Outcome Measures*, 6:117–126, 2015.

Adem Sav, Elizabeth Kendall, Sara S McMillan, Fiona Kelly, Jennifer A Whitty, Michelle A King, and Amanda J Wheeler. 'You say treatment, I say hard work': treatment burden among people with chronic illness and their carers in Australia. *Health & Social Care in the Community*, 21(6):665–674, 2013b.

Patrick Connolly. Expected free energy formalizes conflict underlying defense in Freudian psychoanalysis. *Frontiers in Psychology*, 9(Article 1264):1–15, 2018.

Lilli Herzig, Andreas Zeller, Jérôme Pasquier, Sven Streit, Stefan Neuner-Jehle, Sophie Excoffier, and Dagmar M Haller. Factors associated with patients' and GPs' assessment of the burden of treatment in multimorbid patients: a cross-sectional study in primary care. *BMC Family Practice*, 20(Article 88):1–11, 2019.

Lane Beckes and David A Sbarra. Social baseline theory: State of the science and new directions. *Current Opinion in Psychology*, 43:36–41, 2022.

John Kyle, Dimitris Skleparis, Frances S Mair, and Katie I Gallacher. What helps and hinders the provision of healthcare that minimises treatment burden and maximises patient capacity? A qualitative study of stroke health professional perspectives. *BMJ Open*, 10(3):e034113, 2020.

Austin Due. What are side effects? *European Journal for Philosophy of Science*, 13(Article 16):1–21, 2023.

Mattis Hartwig, Anjali Bhat, and Achim Peters. How stress can change our deepest preferences: Stress habituation explained using the free energy principle. *Frontiers in Psychology*, 13(Article 865203):1–14, 2022.

Lisa Feldman Barrett and W Kyle Simmons. Interoceptive predictions in the brain. *Nature Reviews Neuroscience*, 16 (7):419–429, 2015.

Kasey R Boehmer, Michael R Gionfriddo, Rene Rodriguez-Gutierrez, Abd Moain Abu Dabrh, Aaron L Leppin, Ian Hargraves, Carl R May, Nathan D Shippee, Ana Castaneda-Guarderas, and Claudia Zeballos Palacios. Patient capacity and constraints in the experience of chronic disease: a qualitative systematic review and thematic synthesis. *BMC Family Practice*, 17(1):1–23, 2016.

Michael P Kelly, Carol Brayne, Ann Louise Kinmonth, Natasha Kriznik, John Ford, and Paul C Fletcher. Inequalities in mental health: predictive processing and social life. *Current Opinion in Psychiatry*, 34(2):171–176, 2021.

Aaron Greenhouse-Tucknott, J B Butterworth, James G Wrightson, Nicholas J Smeeton, H D Critchley, Jeanne Dekerle, and Neil A Harrison. Toward the unity of pathological and exertional fatigue: A predictive processing model. *Cognitive, Affective, & Behavioral Neuroscience*, 22:215–228, 2022.

H T McGovern, Alexander De Foe, Hannah Biddell, Pantelis Leptourgos, Philip Corlett, Kavindu Bandara, and Brendan T Hutchinson. Learned uncertainty: The free energy principle in anxiety. *Frontiers in Psychology*, 13 (Article 943785):1–12, 2022.

Hugo Bottemanne and Karl J Friston. An active inference account of protective behaviours during the COVID-19 pandemic. *Cognitive, Affective, & Behavioral Neuroscience*, 21(6):1117–1129, 2021.

Tobias Kube, Liron Rozenkrantz, Winfried Rief, and Arthur Barsky. Understanding persistent physical symptoms: Conceptual integration of psychological expectation models and predictive processing accounts. *Clinical Psychology Review*, 76(Article 101829):1–12, 2020.

Domenico Maisto, Laura Barca, Omer Van den Bergh, and Giovanni Pezzulo. Perception and misperception of bodily symptoms from an active inference perspective: Modelling the case of panic disorder. *Psychological Review*, 128(4): 690–710, 2021.

Thomas Parr, Emma Holmes, Karl J Friston, and Giovanni Pezzulo. Cognitive effort and active inference. *Neuropsychologia*, 184(Article 108562):1–19, 2023.

Alexandre Zénon, Oleg Solopchuk, and Giovanni Pezzulo. An information-theoretic perspective on the costs of cognition. *Neuropsychologia*, 123:5–18, 2019.

Ronald D Adelman, Lyubov L Tmanova, Diana Delgado, Sarah Dion, and Mark S Lachs. Caregiver burden: A clinical review. *Journal of the American Medical Association*, 311(10):1052–1060, 2014.

Marina Bastawrous. Caregiver burden—A critical discussion. *International Journal of Nursing Studies*, 50(3):431–441, 2013.

Zhu Liu, Catrina Heffernan, and Jie Tan. Caregiver burden: A concept analysis. *International Journal of Nursing Sciences*, 7(4):438–445, 2020.