Resource-rational psychopathology

Bilal A. Bari[1] and Samuel J. Gershman[2,3]

[1]Department of Psychiatry, Massachusetts General Hospital, Boston, MA 02114
[2]Department of Psychology and Center for Brain Science, Harvard University, Cambridge, MA 02138
[3]Center for Brains, Minds, and Machines, MIT, Cambridge, MA 02139
Correspondence: Bilal A. Bari, bbari@mgb.org

Funding Sources: This work was supported by a bipolar seed grant from the Harvard Brain Initiative (S.J.G.) and by grant R25MH094612 from the National Institute of Mental Health, Bethesda, MD (B.A.B.).

Conflict of Interest: None

Acknowledgments: We thank Sandeep M. Nayak for providing feedback on an earlier version of the manuscript.

Abstract

Psychopathology is vast and diverse. Across distinct disease states, individuals exhibit symptoms that appear counter to the standard view of rationality (expected utility maximization). We argue that some aspects of psychopathology can be described as *resource-rational*, reflecting a rational trade-off between reward and cognitive resources. We review work on two theories of this kind: rational inattention, where a capacity limit applies to perceptual channels, and policy compression, where the capacity limit applies to action channels. We show how these theories can parsimoniously explain many forms of psychopathology, including affective, primary psychotic, and neurodevelopmental disorders, as well as many effects of psychoactive medications on these disorders. While there are important disorder-specific differences, and the theories are by no means universal, we argue that resource rationality offers a useful new perspective on psychopathology. By emphasizing the role of cognitive resource constraints, this approach offers a more inclusive picture of rationality. Some aspects of psychopathology may reflect rational trade-offs rather than the breakdown of rationality.

Keywords: resource rationality, perception, decision making, information theory, psychopathology

Introduction

We lack a basic understanding and language to explain psychopathology, and without this understanding, we are limited in our ability to diagnose, treat, and prognosticate, among others. Here, we argue that nascent work at the intersection of cognitive science, economics, and information theory has the potential to provide the necessary explanatory framework. We begin with the premise that biological agents are inherently resource limited. *Resource rationality* formalizes the notion that people are doing the best they can, subject to natural information-processing constraints. This resource-rational perspective was developed to explain how people can perform optimally in some domains, and deviate from optimality in other domains (Lewis et al., 2014; Gershman et al., 2015; Griffiths et al., 2015; Lieder and Griffiths, 2020; Bhui et al., 2021; Gershman, 2021). We seek to extend this perspective to gain insight into symptoms that may be shared across states of psychopathology.

The nervous system evolved in the face of myriad constraints, including computational costs (Bossaerts et al., 2019), interference costs (Musslick et al., 2016), metabolic costs (Gailliot and Baumeister, 2007), and others (Shenhav et al., 2017). We focus here on channel capacity, an upper bound on how much information can be transmitted across brain regions (Attneave, 1954; Miller, 1956). We expound two theories of capacity constraints, one applied to perception and the other to action. We will focus primarily on dopamine, as this is the neurotransmitter system with the greatest support for our theories. Psychopathology is far more complex than a single neurotransmitter system and we leave out other relevant systems and brain structures not because they are unimportant, but because the link between them and our theories is more tenuous. We consider a deeper focus on neurobiology outside the scope of this perspective, which we aim to keep at a more theoretical level.

Before getting into details, it is worth stepping back to appreciate the larger conceptual pivot that resource rationality invites us to make. The concept of psychopathology was traditionally based on a division into "pathological" and "non-pathological" minds, but this division has been under strain from both empirical and sociological directions. Empirically, it has become increasingly recognized that many—perhaps all—mental disorders are points on a continuum; there is often no clean dividing line between pathological and non-pathological. Sociologically, the continuum view has led to a "neurodiversity" movement which aims to reframe pathological states as differences rather than deficits. As we will explain below, the resource rationality framework suggests a formalization of the continuum view, where individual differences in cognitive capacity lead to different optimal solutions. *All* of these solutions are optimal, yet they may lead to highly divergent phenotypes. The population may cluster around certain solutions, but these solutions do not necessarily reflect a normatively privileged status. Accepting this proposition opens the door to a computationally informed destigmatization of psychopathology.

Resource rationality does not abandon the notion that some states occupy extremes that require medical treatment. By analogy, a person with missing limbs may be doing the best they can with their available physical resources, but this does not mean that they couldn't do *better* if supplied with prosthetic limbs. Similarly, resource rationality does not guarantee any particular absolute performance level; it only guarantees that an individual will attain a performance level that is superior to the set made available by their supply of cognitive resources. Psychiatric treatment may enable an individual to attain higher absolute performance levels. Resource rationality thus reconciles the continuous nature of psychopathology with the ostensible benefits of treatment.

Rational inattention: capacity limit applied to perception

Given the limitations of biological sensors as well as the statistical challenges of perception, the brain relies on prior, contextual information to constrain what it perceives. Broadly speaking, rational inattention asserts that agents rationally allocate their limited attentional resources (Sims, 2003; Woodford, 2009; Maćkowiak and Wiederholt, 2009; Mackowiak et al., 2018). An equivalent

interpretation, as we will see, is that agents can pay a cognitive cost to obtain a more veridical representation of the world, so long as it does not exceed channel capacity.

Consider an agent inferring a latent variable, such as the time interval between two events. Because sensory signals are imprecise (e.g., time-keeping is noisy), the agent cannot be certain about the underlying latent variable. Bayes' rule states that the agent should combine its sensory evidence with its prior beliefs (e.g., the typical distribution of time intervals) to compute a posterior probability distribution over the values of the latent variable. This is the standard setup in Bayesian models of perception. Importantly, sensory precision is traditionally taken to be outside the control of the agent—an exogenous factor. Rational inattention models generalize this setup to endogenize sensory precision, treating it as a function of attentional control. In other words, sensory precision is modeled as a kind of "cognitive action" that the agent can take, subject to a cognitive cost.

To formalize this idea, we need to first be more precise about what we mean by attention. Following prior work, we conceptualize attention in terms of mutual information (Itti and Baldi, 2009; Feldman and Friston, 2010). Mutual information expresses how much our uncertainty about the latent variable is reduced (on average) after observing data. Intuitively, attending to a signal means extracting information from it—i.e., reducing uncertainty. This information extraction process can be viewed as a kind of communication channel mapping inputs (signals) to outputs (percepts). Like all physical channels it is subject to a capacity limit (an upper bound on mutual information).

Agents will earn more reward on average when their sensory precision is higher. We will refer to the relationship between sensory precision and reward for a given task as the *attentional incentive*. Thus, an agent should increase sensory precision when the attentional incentive is higher. A second factor determining precision is the *attentional cost* incurred by increasing precision, which implicitly depends on the capacity limit. Evidence for these predictions, along with a more technical exposition, is covered further in Gershman and Burke (2022). Figure 1A summarizes the predictions.

Building a bridge to neurobiological mechanisms, Mikhael et al. (2021) developed a rational inattention account of tonic dopamine. Under this account, tonic dopamine subsumes the average reward theory of tonic dopamine, where it encodes the context (state)-specific average reward rate (Niv et al., 2007; Beierholm et al., 2013; Hamid et al., 2016). Specifically, rational inattention does not propose a different role for how dopamine encodes reward than what has been posited previously. Tonic dopamine, by reporting average reward, is hypothesized to set the baseline for learning the value of specific actions within a given state, and can give rise to phenomena like asymmetric learning and exploration/exploitation, ideas which we will elaborate below. Because the rational inattention framework couples average reward to sensory precision, it predicts that changes in tonic dopamine levels should control the allocation of attention, consistent with many pharmacological and physiological findings. The reward-attention coupling also provides an integration of the average reward theory with the active inference theory developed by Friston et al. (2012), according to which tonic dopamine controls the precision ("salience") of external and internal cues (see also Shi et al., 2013, for a related theory applied to time perception).

Policy compression: capacity limit applied to action

All actions, from the mundane to the significant, require memory. These memories are stored in the brain as policies, or mappings from states to actions, where states are defined as the representation of information needed to predict reward (Sutton and Barto, 2018). As an example, imagine being tasked with purchasing groceries for the family. The state representation includes the items available for purchase as well as the individual preferences of family members, and the actions include either purchasing or not purchasing an item. Intuitively, you can satisfy each individual's preferences, but at the cost of a mentally-demanding trip to the store. If you choose instead to ignore certain preferences, you can reduce cognitive demand, by reducing the number of items that must be remembered, at the cost of reducing the overall satisfaction of the family.

Policy compression formalizes this intuition, by conceptualizing the mapping from states to

actions as a communication channel—just like we posited for perception—and postulating that this communication channel has a limited capacity (Parush et al., 2011; Gershman, 2020; Lai and Gershman, 2021). Under policy compression, agents must optimize the trade-off between reward and policy complexity, which we define as the mutual information between states and actions. Because policy complexity is a lower bound on the number of bits needed to store a policy in memory, more complex policies necessitate more bits. If the policy complexity exceeds capacity, then agents must "compress" the policy in order to transmit it across brain regions. Policies with high complexity require greater memory and can lead to greater reward. In contrast, policies with low complexity require less memory to implement, but are generally suboptimal. At the extreme, if the policy is the same in every state, then the policy complexity is minimized (mutual information is 0).

The optimal capacity-limited policy has a number of interesting features. First, it takes the form of the ubiquitous softmax function, in which an "inverse temperature" parameter governs the stochasticity in the policy. When capacity is high, policies become more deterministic (via a larger inverse temperature parameter) and concentrate on the action with maximal value. When capacity is low, policies become less state-dependent (via a smaller inverse temperature parameter). More specifically, the inverse temperature is lower (i.e., choices become more random) when varying the policy complexity has a greater effect on reward, which occurs at low values of policy complexity. Second, the optimal policy includes a perseveration term. When capacity is large, the inverse temperature term is large and actions are largely driven by the values of the underlying states. When the capacity is small, the inverse temperature term decreases, and the perseveration term can dominate the policy. Third, more complex policies result in slower response times, because the brain must inspect more bits to find the coded state (Hick, 1952; Lai and Gershman, 2021; Bari and Gershman, 2023). These regularities are summarized in Figure 1B.

Psychiatric phenomena

Mania

Rational inattention provides a rich language for describing numerous symptoms of mania, which we propose is best understood as an individual's belief that their precision has increased, without an increase in true precision. In other words, mania may be the result of precision *miscalibration*, where precision is overestimated (Mikhael et al., 2021).

Clinically, mania can be a distinctly euphoric state (Cassidy et al., 1998b), one that patients are often unwilling to request or accept treatment for (Baldessarini et al., 2008). It is not uncommon for patients with bipolar disorder, a disorder characterized by oscillations between mania and depression, to self-discontinue medications, either because it makes them feel 'depressed' (e.g., relative to prior mania/hypomania) or in the hopes that they may experience a manic state (Devulapalli et al., 2010; Crowe et al., 2011). Anecdotally, some patients who accept treatment are doing so not for the mania itself to be treated, but because personal experience has taught them that their mania can develop into a mania with psychosis. Rational inattention provides a clue for the intoxicating effects of mania: a subjective increase in precision is associated with an increase in average reward. If mania causes patients to experience the world as highly rewarding, it's understandable why they would desire to remain in that state.

This increase in the estimate of average reward manifests as an asymmetry in how agents estimate value functions. In mania, because agents come to expect reward (i.e., the prior over rewards is shifted higher), they exhibit a persistent 'optimism' in which their value functions are shifted higher, following either positive or negative feedback. In other words, relative to an agent with a veridical estimate of average reward, optimistic agents come to expect reward even when they objectively should not. In mania, this is consistent with clinical intuition and the empirical literature (Alloy et al., 2016; Kwan et al., 2020). Note that rational inattention does not predict faster

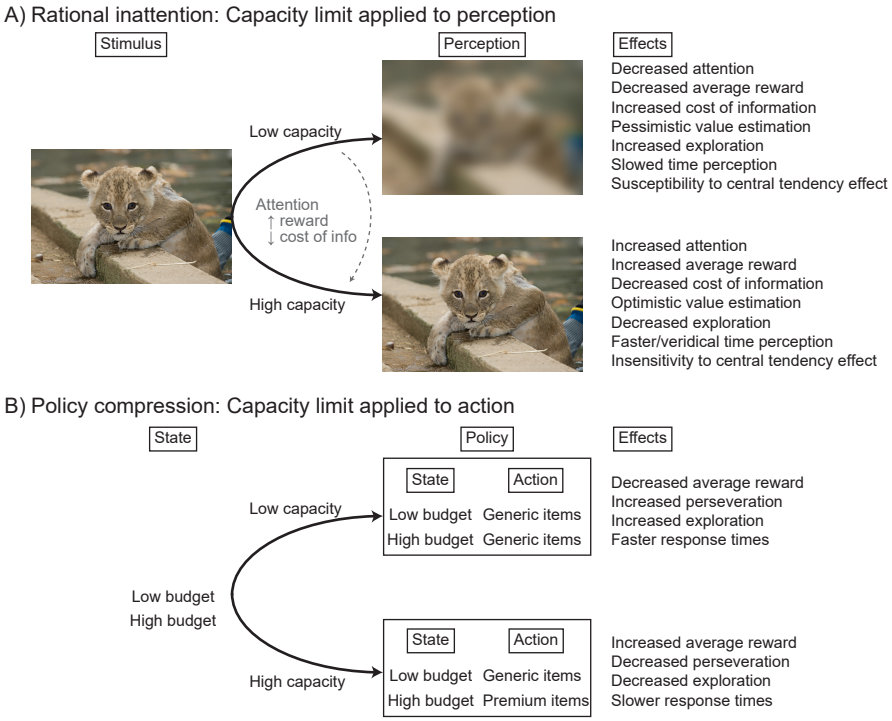

Figure 1: Perception and action as communication channels.
A: Rational inattention describes perception as a communication channel, subject to a capacity limit, or upper bound on amount of information that can be transmitted across sensory channels. In this example, we highlight a stimulus being encoded by the brain either under low capacity conditions or high capacity conditions. The brain is able to increase the capacity of encoding by devoting greater attention - the cognitive process of reducing uncertainty about a stimulus. Under rational attention, the factors that increase attention are increased reward (in our framework, the attention incentive) or decreased cost of information. The consequences of encoding at differing capacities is highlighted to the right. Photo obtained from Smithsonian's National Zoo and Conservation Biology Institute open access images.
B: Policy compression describes action selection as a communication channel, subject to a capacity limit. In this example, an agent is tasked with purchasing groceries under two possible states: low budget, or high budget. With a low budget, the optimal policy is to purchase generic items. With a high budget, the optimal policy is to purchase premium items. Under conditions of low capacity, the agent is state-insensitive, and purchases generic items regardless of the state. Under conditions of high capacity, the agent is highly state-sensitive, and exhibits the optimal policy. The consequences of these capacity limits is highlighted to the right.

learning (e.g., the trial-to-trial change in expectation) from positive feedback under an optimistic prior, only that the value function is initialized optimistically. Learning is otherwise consistent with the predictions of Bayesian inference, in which more surprising observations - those farther from the prior - are learned faster.

High estimated precision implies less dependence on a stored internal prior. Patients in a manic state display the expected hallmarks—they are highly attentive, ever-present, keenly aware of their environments, and outwardly directed. In the context of interval timing, high estimated precision also predicts a faster internal clock (see Mikhael et al., 2021). In mania, patients exhibit myriad symptoms consistent with a faster internal clock. They are are classically psychomotor agitated: they appear to be moving in fast motion, restless and always on the move, with rapid speech that can be difficult to interrupt (Cassidy et al., 1998a). A faster clock also results in faster thought, consistent with subjective experience and occasionally at a pace so rapid as to be aversive. The subjective sense of time is sped up (Bschor et al., 2004) and can become so grossly miscalibrated that patients will sense that tens of minutes have elapsed after only a minute or two.

In the context of reward learning, overestimated precision induces heightened sensitivity to noise, which may appear as a form of "exploration" (choice randomness). Clinically, this manifests in what the Diagnostic and Statistical Manual of Mental Disorders refers to as an "increase in goal-directed activity" (American Psychiatric Association, 2022). Patients in a manic state are famous for starting (but not necessarily completing) dozens of new projects, hobbies, books, television shows, and so on (Dailey and Saadabadi, 2018). In general, the new activities are consistent with what the patient feels is worth pursuing, and not just activity for its own sake. This is generally consistent with the notion of value-based random exploration: activities tend to be those of higher value.

Distractibility, a key diagnostic criterion for mania, is another facet of heightened noise sensitivity due to precision miscalibration. On this account, distraction arises when task-irrelevant distractors are misinterpreted as task-relevant cues.

One aspect worth emphasizing is that all of the above phenomena, with the exception of distractibility, are consistent with a *true* increase in precision and therefore a true increase in capacity. How can we discern a precision overestimation account of mania from a true increase in precision? The former predicts degraded perception while the latter predicts improvement. Consistent with precision overestimation, the literature supports the idea that perception is degraded in mania (Kohler et al., 2011; O'Bryan et al., 2014). As a separate prediction, the precision overestimation hypothesis predicts an increase in random exploration, while true precision increase predicts a decrease. There is some evidence to suggest increased exploratory behavior in mania (Ryu et al., 2017), although there is more work to be done.

Depression

Under rational inattention, if mania can be described by an increase in precision (either real or perceived), then depression in many ways can be viewed as its opposite. With a decrease in precision comes an increase in reliance on an internal prior. Patients with depression frequently speak of a subjective 'grey'-ness of experience, with a sensation that they cannot perceive or experience the world as they did when they were well. If, in depression, the sensory precision is reduced, then the posterior will be dominated by the prior, and subjective experience will necessarily be less rich and less modulated by perception of the outside world. As a result, patients appear inattentive to the outside world (Keller et al., 2019), with their focus directed inward.

Just as increased sensory precision speeds up the internal clock, decreased precision slows it down (Bschor et al., 2004). This too explains the general slowness observed in depression, with overt psychomotor slowing manifesting as sluggish gait and movement, turning mundane tasks such as dressing and showering into time-consuming chores. Speech itself slows down (Koops et al., 2023) and patients report a sense of slowed thinking.

A reduction in precision signals to an optimal agent to expect reduced reward. This may explain the subjective intolerability of depressive states. It comes as no surprise that depression is a risk factor for suicidal thinking (Franklin et al., 2017), though there is some circularity here as suicidal thoughts are a diagnostic criteria for depression. A reduced expectation of reward also shifts the balance of learning towards pessimism. This may lead to the sense of hopelessness that is pervasive in depression (Abramson et al., 1989; Cusin et al., 2010). Reduced precision also renders behavior less responsive to feedback, consistent with what has been observed in depression (Steele et al., 2007). From the perspective of reinforcement learning theory, this is consistent with a reduced learning rate, which has been observed in depression (Brown et al., 2021), though other results have been equivocal (Chen et al., 2015).

Patients may be convinced that they lack agency to meaningfully affect their lives. Here we predict that depression may also reduce the attentional incentive, even in circumstances where patients have clear agency. In other words, patients will perceive a lack of controllability, even if this is at odds with reality (Miller and Seligman, 1975). There is a long and rich literature on learned helplessness in depression (Maier and Seligman, 1976, 2016), and rational inattention provides another perspective: if the attentional incentive is decreased or erased, rational agents should not allocate attention to the task at hand, as attention is only worth the cost if outcomes can be improved. This in turn can manifest as reduced motivation—after all, why engage if the result won't change? Further, clinical experience suggests that to the extent that patients are motivated to interact with the world, they need to frequently be reminded of activities they find highly rewarding (i.e., encouraging patients to exploit). This is consistent with the idea that exploitation should be reduced under conditions of low capacity, consistent with what has been observed with reinforcement learning modeling in depression (Blanco et al., 2013).

Patients with dense depression are known to have cognitive deficits in numerous domains (Hack et al., 2023). Under policy compression, reduced capacity for actions is associated with a decrease in working memory and an increase in perseverative behavior. These deficits manifest clinically, as patients can have a difficult time retaining basic information and can be perseverative in their thoughts and their behaviors (Trick et al., 2016). These deficits have been identified in the lab, as depressed patient show pronounced perseverative errors in set shifting tasks (Martin et al., 1991; Channon, 1996; Ilonen et al., 2000; McGirr et al., 2012) and have working memory deficits (Channon et al., 1993; Burt et al., 1995; Christopher and MacDonald, 2005; Rose and Ebmeier, 2006).

Unlike mania, we make no strong claims about whether depression is due to a true decrease in precision vs. a decrease in estimated precision, as both of these mechanisms will give rise to all the symptoms we discussed. One conceptual difference is that a true decrease in precision results in a decrease in capacity, whereas precision underestimation only decreases perceived capacity, rather than actual capacity. In the latter case, cognitive abilities should be intact, in some sense, with deficits rendered by patients perceptions of their own abilities. In line with this thought, there is some evidence that the cognitive deficits in depression are mediated by variables such as effort (Moritz et al., 2017).

Stimulants

Prescribed stimulants fall into two major classes, amphetamines and methylphenidate. Amphetamines function, in part, by inducing the release of dopamine (Schiffer et al., 2006) whereas methylphenidate functions as a stronger inhibitor of the dopamine transporter (John and Jones, 2007), a protein which reuptakes dopamine into presynaptic terminals. Despite differences in mechanism (with some overlap), these drugs effectively increase synaptic dopamine concentration (Kuczenski and Segal, 1997). Prior work has shown that stimulants decrease the energy the brain uses to perform cognitively demanding tasks, analogous to a reduction in the attentional cost parameter under rational inattention (Volkow et al., 2008). Individuals take these medications to feel more attentive and focused. Time perception is sped up (Lake and Meck, 2013), in line with predictions, and with a faster clock comes faster movements, faster speech, and faster thinking. The mental states induced

by stimulants are distinctly pleasurable and contribute in no small part to potential for misuse. Overall, there is remarkable similarity between the effects of stimulants and mania, but they are by no means identical.

Antipsychotics

Under rational inattention, if increasing dopamine can recapitulate the effects of increased capacity, then decreasing dopamine should do the opposite. Clinically, antipsychotics serve this function, a subset of which functions, in part, by blocking the D2 receptor (McCutcheon et al., 2023). Further, recent work has argued that clinically-efficacious antipsychotics, include those with minimal direct dopaminergic effects, function by modulating D1 receptor-expressing neurons in the striatum (Yun et al., 2023). Rational inattention on its own does not explain how antipsychotics reduce the hallucinations, delusions, and disorganized thought for which they are indicated, but we believe it does help explain the intolerability of this class of medications (Valenstein et al., 2004; Ascher-Svanum et al., 2010). In fact, the largest trial of antipsychotics to date chose 'discontinuation of treatment for any cause' as a primary outcome (Lieberman et al., 2005), which substantiates their intolerability. Patients who have tried numerous antipsychotics describe a sense of feeling subjectively slowed, cognitively dulled, with limited attention, and a sense that they are perceiving the world through a fog. All of these effects are consistent with a reduced capacity. Antipsychotics must be dosed carefully in psychotic illnesses, as they can exacerbate the cognitive symptoms that limit the ability of patients to function (Kelley et al., 1999; Kasper and Resinger, 2003), an entity once known as neuroleptic-induced deficit syndrome (Lader, 1993; Schooler, 1994)

Attention-deficit/hyperactivity disorder

Inattention is a cardinal feature of attention-deficit/hyperactivity disorder (ADHD), and may naturally be explained under rational inattention as reduced sensory precision. This explains the failure to 'give close attention,' 'difficulty sustaining attention,' and 'often easily distracted by extraneous stimuli' that are diagnostic symptoms. The increased reliance on the internal prior can contribute to the sense that their minds seem elsewhere. Distractibility and careless mistakes can be thought of as an increase in random exploration (Hauser et al., 2014) which, for reasons stated previously, is increased under more stringent capacity limits. The dislike of mentally effortful tasks is also well-explained: if one cannot provide the attention necessary to complete a task, due to reduced capacity, then it is rational to avoid those tasks. If stimulants work as proposed above, then the rationale for their use in ADHD is clear.

Schizophrenia

Schizophrenia and other psychotic illnesses are primarily characterized by their positive symptoms: hallucinations, delusions, and disorganized thought. In contrast, the negative symptoms more frequently restrict the ability of patients to fulfill typical societal roles and responsibilities (e.g., maintaining friendships, managing household tasks), in part due to our inability to adequately treat them (Aleman et al., 2017). These symptoms include amotivation, asociality, blunted affect, and general cognitive impairments (Correll and Schooler, 2020). From the perspective of policy compression, a number of these symptoms can be explained by a reduction in channel capacity. Indeed, patients with chronic schizophrenia exhibit reduced capacity (Gershman and Lai, 2021). This can be linked to blunted affect (decreased expressivity of emotions) and alogia (the reduction in quantity of words spoken). Reduced capacity can also explain working memory deficits, which have been robustly demonstrated in schizophrenia (Forbes et al., 2009; Collins et al., 2014). We note that reduced capacity in chronic schizophrenia is confounded by chronic antipsychotic use, which may contribute independently to changes in capacity.

Parkinson's disease

Parkinson's disease is characterized by widespread degeneration of the dopaminergic system (as well as other neuromodulatory systems). From the perspective of policy compression, individuals with Parkinson's disease have reduced capacity compared to age-matched controls (Bari and Gershman, 2023). This reduced capacity can explain a number of cognitive symptoms seen in Parkinson's disease, including: memory problems, language difficulties (word-finding difficulty, naming/misnaming, comprehending complex sentence structure, dysarthria), and general problem solving and executive functioning difficulties (Dubois and Pillon, 1996; Verbaan et al., 2007). Providing patients with dopaminergic agents increases capacity for actions, allowing subjects with Parkinson's disease to entertain more complex policies (Bari and Gershman, 2023). This is consistent with clinical observation that dopaminergic therapy can relieve a number of cognitive symptoms (although not to the same extent as motor symptoms; Dubois and Pillon, 1996; Robbins and Cools, 2014). Counterintuitively, in our analysis, dopaminergic therapy *slows* participants down, as measured by response rates, which is opposite of the general effect of these treatments in relieving bradykinesia. This is consistent with policy compression, which predicts that more complex policies require greater time to decode (i.e., map from the the compressed representation to overt actions; Hick, 1952; Lai and Gershman, 2021).

Rational inattention provides a complementary perspective. Reduced capacity explains reduced attention and slower speed of thinking. It also explains the observation that, in subjects with Parkinson's disease, dopaminergic therapy restores sensitivity to feedback (Frank et al., 2004; Rutledge et al., 2009). Another manifestation is the stronger central tendency effect in Parkinson's disease. In the context of interval timing, this effect describes a tendency of subjects to overreproduce short intervals and underreproduce long intervals in timing reproduction tasks (Malapani et al., 1998, 2002; Shi et al., 2013). Under rational inattention, this is consistent with a strong migration towards the prior induced by conditions of low attention / low tonic dopamine (Mikhael et al., 2021; Mikhael and Gershman, 2022). Consistent with this account, providing subjects with dopaminergic medication reduces the magnitude of the central tendency effect.

Neurodevelopmental disorders: Specific learning disorders

Several specific learning disorders manifest as difficulties in processing specific sources of information. Examples include difficulties in processing language, written information, numerical / mathematical information, and social information. One consequence is a decrease in reliance on these sources, which over time atrophies the brain's ability to use them. Under rational inattention, if these processing difficulties arise from aberrant precision, then it is rational to decrease reliance on them and focus on higher precision sources of information. This has the deleterious consequence of diminishing the brain's ability to use this information, which can create difficulties with functioning later in life. This highlights the need to design curricula that force the brain out of the rational but deleterious underreliance on this aberrant information. Instead, if individuals can be trained to use this low-precision information, precision may increase with experience, especially during valuable critical periods early in life. We turn the interested reader to Jones et al. (2023) for a thoughtful perspective.

Perseveration

Under policy compression, perseveration emerges as the optimal policy under low capacity. If low capacity is common to numerous psychiatric conditions, then we would expect perseveration to arise as a transdiagnostic symptom. Indeed, perseveration is observed in numerous psychiatric conditions, some of which we have detailed above (Serpell et al., 2009). As an extreme, in delirium, it is not uncommon for patients to repeat answers to the first question asked, even if the answer is nonsensical. Patients can act out more complex policies, like those seen in addiction (Lane et al.,

2007; Woicik et al., 2011), which are nevertheless resistant to change. It is seen in conditions ranging from schizophrenia (Crider, 1997), to depression (Martin et al., 1991), to frontal lobe pathology and other neurodegenerative conditions (Joseph, 1999; Oosterloo et al., 2019), to name a few.

Discussion and limitations

Our resource-rational framework has a remarkable degree of overlap with prior work in learning and decision making. First, much of the neurobiology of reward-based decision making is motivated by reinforcement learning theory, which has furnished the field with error-driven learning models (Sutton and Barto, 2018). In these models, reward prediction errors (the difference between actual and expected reward) drive sequential learning. As explained in the Appendix, one form of rational inattention uses this familiar error-driven update rule to estimate the posterior over the parameter of interest; if the parameter is reward, then this rule is the familiar reward prediction error. In deriving policy compression, we assumed the value function was known, an assumption which cannot hold for agents learning in novel environments. We have previously developed process models to iteratively estimate the value function using familiar error-based update rules (Gershman and Lai, 2021). We withhold a deeper discussion of process models as this remains an active area of research.

Second, policy compression provides insights into habits, repetitive behaviors that are famously insensitive to outcome devaluation or contingency degradation, and which can interfere with goal-directed behavior (Dickinson, 1985; Wood and Rünger, 2016; Miller et al., 2019). Similarly, low complexity policies are perseverative since they are dominated by the marginal action distribution, which is not dependent on rewards (Gershman, 2020). Policy compression does not, however, explain the shift from "goal-directed" behavior to habitual behavior that occurs with training (Balleine and O'Doherty, 2010). One hypothesis is that the shift to habitual behavior (i.e., policies of low complexity) may free up finite capacity, which we assume is fixed, so it is not all allocated to one task. Third, rational inattention subsumes an influential account relating tonic dopamine to average reward availability in a given context and response vigor (Niv et al., 2007). In other words, our derivation of rational inattention makes the same predictions as the average reward theory of dopamine, and extends it to precision.

One conceptual leap we have made is to propose that the affective symptoms in mania and depression arise not from deficits in reward processing, but from aberrations in attention. We made this leap based on parsimony, as aberrations in attention allow us to explain not just affective symptoms, but numerous symptoms related to psychomotor state, learning, and decision making. An influential account of anhedonia, a common symptom in depression characterized by the inability to experience pleasure, holds that it may arise as a consequence of impaired reward sensitivity (reduction in the perception of reward magnitude; Huys et al., 2013). The authors note that reward insensitivity, under certain assumptions, is equivalent to over-exploration. This latter view is close to what we propose with rational inattention, which suggests that over-exploration, as a consequence of precision underestimation, may masquerade as reward insensitivity.

We have argued that distractibility arises in mania, due to precision overestimation, and ADHD, due to reduced precision. How does rational inattention predict distractibility in both of these circumstances? At first pass, it would appear that only agents with reduced precision should be affected, since they do not have the sensory precision to focus on the task at hand. It is important to recognize that although distractibility arises in both mania and ADHD, they are observably distinct phenomena on clinical evaluation. What they share in common is an inability to follow a task through to completion. In mania, there is a sense that patients are intensely interested in their environments, focused not just on the task at hand but also task-irrelevant information. In ADHD, there is a sense of disinterest or ambivalence in the task at hand. In mania, therefore, distractibility arises from amplification of signal *and* noise, and patients with mania assign undue importance to inappropriate samples that impinge on their senses. In ADHD, distractibility arises from reduced attention to the signal. Phrased this way, rational inattention provides transdiagnostic insight into

distractibility, and makes it clear that it arises from distinct computational aberrations. Note that rational inattention does not predict distractibility if precision is increased but not miscalibrated (true and estimated precision are equal). It is only in the setting of precision overestimation that we predict distractibility.

Although rational inattention and policy compression have broad implications for psychiatry, we do not want to leave readers with the impression that capacity limits offer complete explanations of psychopathology. How could they describe such a wide range of phenomena? In brief, they don't, at least not in isolation. Our general belief is that the psychiatric phenomena we describe should not be seen as lying along a single dimension (capacity), but that disease states will require several dimensions to sufficiently define the relevant symptoms. For example, we propose that both depression and ADHD result from low capacity (under rational inattention), yet mood symptoms are nowhere to be found in the ADHD diagnostic criteria. We view this as an opportunity for computational psychiatry to identify the relevant symptom dimensions.

In mania, for example, we provide some insight into elation, but mood effects also include irritability or anger[1] (Cassidy et al., 1998b). Patients with depression may not exhibit any symptoms of slowing, nor any measurable cognitive deficits (Hack et al., 2023). Even more extreme, symptoms of mania may coexist with symptoms of depression, as in mixed states, in which racing thoughts coexist with psychomotor slowing, or more generally any combination of mood, speed of thought, and psychomotor state (Kraepelin, 1921; Marneros, 2001). The predictions of rational inattention hold for *optimal* agents, in which reward, precision, and capacity are linked. Perhaps, in mixed states and other disease states, this assumption of optimality is violated. For example, an agent with an optimistic prior and reduced sensory precision may be characterized by what Emil Kraepelin called 'manic stupor,' characterized by elevated mood and psychomotor slowing.

The notion that reward, precision, and capacity are linked for optimal agents has consequences for the relevant causal deficit. We highlighted particular deficits (e.g., precision miscalibration in mania) given our intuition for the relevant causal deficit. However, since the theory links multiple variables together, an aberration in any one could lead to the same symptoms. Neuroscience will play a valuable role in determining the relevant causal deficit. Rational inattention has recently gained traction in neuroscience (Mikhael et al., 2021; Grujic et al., 2022; Wu et al., 2022), and we are hopeful for deeper insights in the future.

Both mania and depression can become severe enough that symptoms of psychosis emerge. Since, under rational inattention, mania and depression exist on opposite sides of the spectrum, this observation suggests that psychotic symptoms may emerge via an independent process. Likewise, mania or depression can be complicated by comorbid anxiety. In ADHD, our theories do not yet provide insight into the hyperactivity symptoms that are the more frequently observed consequences of the disorder. In addition, our theory predicts slowed time perception in ADHD, which is at odds with empirical data demonstrating time perception is more rapid and normalizes with treatment (Smith et al., 2002; Ptacek et al., 2019). This highlights either a limitation of our theory, a limitation of our understanding of ADHD, or - most likely - both.

In schizophrenia, policy compression explains several negative symptoms, but it is does not exhaustively explain all negative symptoms, including asociality and amotivation. It also fails to describe anhedonia, although this may not be a failure of the theory, as there is a body of work advocating for intact hedonic drive in schizophrenia (Burbridge and Barch, 2007; Kring and Moran, 2008; Cohen and Minor, 2010; Dowd and Barch, 2010; Yee et al., 2010; Llerena et al., 2012). Neither policy compression nor rational inattention provides insight into the positive symptoms of psychosis, nor how antipsychotics function to reduce these symptoms.

Amotivation is an interesting case study. Amotivation is also a feature of Parkinson's disease, and is relieved by dopaminergic agonists, at least early in the disease (Pagonabarraga et al., 2015), which we argued increases capacity for more complex policies. Similarly, patients with chronic schizophrenia, who have reduced capacity, also typically suffer from amotivation. Thus, there is

[1]Note, however, that such effects have been argued to represent a distinct subtype of mania (Cassidy et al., 1998a).

suggestive, but incomplete, evidence for a link between capacity and amotivation.

Our discussion of stimulants and antipsychotics is likewise incomplete. Although stimulants and mania have some overlap—stimulant intoxication can manifest as mania—clinical experience makes it obvious that antipsychotics do not phenocopy depression. Although dopamine blockade via antipsychotics is a mainstay of treatment for mania, stimulants are not generally recommended for depression. This is in large part because the neurobiology is far more complicated than we have laid out here, with psychopathology sculpted by the relevant brain structures, neuromodulatory systems, cell types, receptor subtypes and densities, and myriad other details. We have couched capacity limits in the language of dopamine, but we recognize the complexity of disease. Depression, for example, is more commonly thought to involve dysfunction of the serotonergic system (Coppen, 1967; Meltzer, 1990; Owens and Nemeroff, 1994; Harmer et al., 2017), with recent work highlighting a role for extra-neuronal mechanisms (e.g., inflammation; Raison et al., 2006). Stimulants, as another example, involve the release of multiple neuromodulators, not just dopamine (our focus here). In short, a more complete picture will require a theory that encompasses multiple neuromodulatory systems.

Conclusion

Information-theoretic resource rationality provides a rich transdiagnostic language for describing psychopathology. We summarize our perspective in Table 1. Rationality, on this view, does not provide a single solution in phenotype space, but rather a Pareto frontier of optimal solutions. A wide range of psychopathology may be thought of not as suboptimal simply because it results in poorer task performance, but rather optimal performance under an illness-induced capacity limit. Our hope is that these frameworks provide rich ground for development of new theories, behavioral tasks, and for uncovering the neurobiological loci of mental illness.

Appendix: Technical details

This appendix summarizes theoretical results from past papers (Lai and Gershman, 2021; Mikhael et al., 2021; Gershman and Burke, 2022). We refer readers to those papers for further details.

Rational inattention

Suppose an agent is inferring a parameter μ about the world (this can be expected reward, a temporal interval, an object category, etc.) and it observes a sample x. Bayes' rule prescribes a normative solution, which states that agents can combine what they observe, $P(x|\mu)$, with prior information, $P(\mu)$, to generate a posterior estimate of the parameter, $P(\mu|x)$:

$$P(\mu|x) \propto P(x|\mu)P(\mu).$$

For analytical tractability, we will assume the sample distribution is Gaussian, $x \sim \mathcal{N}(\mu, \lambda^{-1})$, with mean μ and precision (inverse variance) λ. If we assume the prior is also Gaussian, $\mu \sim \mathcal{N}(\mu_0, \lambda_0^{-1})$, then the posterior is also Gaussian with mean $\hat{\mu}$:

$$\hat{\mu} = \mu_0 + \frac{\lambda}{\lambda + \lambda_0}(x - \mu_0)$$

This equation takes the form of an error-driven update rule, where $x - \mu_0$ is the error and $\frac{\lambda}{\lambda+\lambda_0}$ is the learning rate, determined by the relative precision between the likelihood and prior. Note that if the agent is inferring reward, then $x - \mu_0$ corresponds to the reward prediction error in

Phenomenon	Mechanistic Hypothesis	Symptoms Explained	Symptoms Not Explained
Mania	Rational inattention Precision miscalibration (overestimation)	Euphoria Optimistic value estimates Increased attention/focus on outside world Psychomotor agitation Rapid speech Rapid passage of time Distractibility	Dysphoria/irritability Mixed states Emergence of psychosis
Depression	Rational inattention Precision miscalibration (underestimation) or precision decrease	Decreased mood Inattentive to outside world/inwardly drawn Subjective "greyness" of experience Psychomotor slowing Slowed speech Slower passage of time Pessimistic value estimates Blunted response to feedback	Psychomotor agitation Irritability Mixed states Emergence of psychosis Depression without cognitive changes
	Rational inattention Decreased attentional incentive	Decreased perception of control (learned helplessness)	
	Policy compression Reduced capacity	Working memory deficits Perseveration	
Stimulants	Rational inattention Reduction in attentional cost	Euphoria Increased attention/focus on outside world Psychomotor agitation Rapid speech Rapid passage of time	Psychosis
Antipsychotics	Rational inattention Reduced capacity	Cognitive dulling / slowing Limited attention Difficulty perceiving world Exacerbation of negative symptoms	Antipsychotic effects
ADHD	Rational inattention Precision decrease	Failure to give close attention Difficulty sustaining attention Mind seems elsewhere Careless mistakes (exploration) Dislike of mentally effortful tasks Distractibility	Hyperactivity symptoms Faster perception of time
Schizophrenia	Policy compression Reduced capacity	Blunted affect Alogia Working memory deficits	Positive symptoms Amotivation Asociality
Parkinson's disease	Rational inattention Reduced capacity	Central tendency effect Reduced attention Slower speed of thinking	Motor symptoms Sleep disorder Psychosis
	Policy compression Reduced capacity	Memory problems Language difficulties Word-finding difficulty Naming errors Difficulty with complex sentence structure Dysarthria General problem solving and executive functioning difficulties Reduced response time w/ dopamine therapy	
Specific learning disorders	Rational inattention Imprecise sensory input	Underreliance on imprecise sensory input and overreliance on other sensory inputs	Practice improving precision of imprecise signals Reduced motivation to engage with imprecise signals (Matthew effect)
Perseveration	Policy compression Reduced capacity	Tendency to repeat actions across diagnostic entities	Hierarchical nature of perseveration (e.g., repetition of actions in some disorders, repetition of action sets in others)

Table 1: Summary of psychiatric phenomena explained by rational inattention and policy complexity.

reinforcement learning theory. Greater relative precision of the prior shifts the posterior estimate closer to the prior (what the agent assumed), whereas greater relative precision of the likelihood shifts the posterior estimate closer to the likelihood (what the agent perceived). Under rational inattention, agents control the likelihood precision λ. When, then, does it make sense to modulate precision?

First, we must formalize what we mean by attention and how it increases information. The information transmission rate across sensory channels is

$$I(\mu; x) = H(\mu) - H(\mu|x),$$

where $I(\mu; x)$ is the mutual information between parameter μ and signal x, $H(\mu)$ is the entropy of the prior, and $H(\mu|x)$ is the entropy of the posterior. Intuitively, high mutual information means that observing the sample reduces uncertainty about the parameter estimate. Low mutual information means the prior and posterior distributions are similar and therefore observing the sample contributes little to uncertainty reduction. Mutual information formalizes what we mean by attention. For our Gaussian generative model, the mutual information is given by:

$$I(\mu; x) = \frac{1}{2} \log \left(1 + \frac{\lambda}{\lambda_0} \right).$$

Shannon's noisy-channel coding theorem states that the minimum number of bits needed to communicate μ without error across a noisy channel is on average equal to $I(\mu; x)$. A corollary, therefore, is that errorless communication is impossible if the agent's capacity is less than $I(\mu; x)$, in which case there is a trade-off between the cost of attention and the cost of error. This trade-off is analyzed by rate distortion theory (Berger, 1971), which is equivalent to rational inattention (Denti et al., 2020).

Let us assume agents receive reward inversely proportional to the squared error $\epsilon = (\mu - \hat{\mu})^2$, which should motivate agents to reduce their error. Intuitively, if λ is small, μ will migrate to the prior and cause a large error between the latent source and its estimate. As λ grows large, this error is reduced. For simplicity, let us assume that the reward agents receive, $u(\epsilon)$, is a monotonically decreasing and differentiable function of this error. Expanding the first-order Taylor series approximation around $\epsilon = 0$ yields $u(\epsilon) \approx u(0) - \theta\epsilon$, where $\theta > 0$ is the negative slope of $u(\epsilon)$ at $\epsilon = 0$. We can interpret θ as an attentional incentive parameter to capture the idea that agents should be motivated to pay attention when reward is contingent on error. The expected reward is therefore $U = \mathbb{E}[u(\epsilon)] \approx u(0) - \theta\mathbb{E}[\epsilon]$.

We can now write down the constrained optimization problem faced by agents

$$\lambda^* = \underset{\lambda}{\operatorname{argmax}}\, U - \kappa I(\mu; m)$$

where κ is the Lagrange multiplier. κ can be interpreted as the attentional cost, to formalize the notion that attention is effortful. κ can be thought of as an "exchange rate:" one unit of reward can be "bought" for κ units of information. Written this way, κ implicitly represents the capacity limit, with a large κ representing low capacity (many bits for one reward) and small κ representing high capacity (few bits for one reward).

This equation formalizes the tradeoff between reward (U) and information: agents can increase λ to increase reward U by reducing their error ϵ, but doing so increases the information rate $I(\mu; m)$, which the agent must keep at or below capacity. Solving this constrained optimization problem yields:

$$\lambda^* = \max \left(0, \frac{2\theta}{\kappa} - \lambda_0 \right).$$

The optimal precision increases as 1) the attentional incentive, θ, increases, 2) the attentional cost, κ decreases, and 3) the prior precision, λ_0, decreases.

For the Gaussian generative model we described above, the optimal expected reward is:

$$U^* \approx u(0) - \frac{\theta}{\lambda^* + \lambda_0} = u(0) - \frac{\kappa}{2}.$$

Mikhael et al. (2021) used this equivalence to posit a rational inattention account of tonic dopamine. The authors propose a model in which tonic dopamine encodes average reward, U, and by the equivalence demonstrated above, also encodes posterior precision, $\lambda^* + \lambda_0$ and the information-reward exchange rate, κ (and implicitly, the capacity).

Policy compression

We model an agent that visits states, s, and takes actions, a, to earn reward. States are defined as the representation of information needed for reward prediction. Each state is visited with probability $P(s)$ and an action is chosen according to a policy $\pi(a|s)$, a probabilistic mapping from states to actions. We conceptualize the policy as a communication channel mapping states to actions. The minimum number of bits to achieve error-free communication of the state identity is given by the mutual information between states and actions:

$$I(S;A) = \sum_s P(s) \sum_a \pi(a \mid s) \log \frac{\pi(a \mid s)}{P(a)},$$

where $P(a) = \sum_s P(s)\pi(a|s)$ is the marginal action distribution. We use the term *policy complexity* to refer to $I(S;A)$. Intuitively, a policy with high complexity is highly state-dependent (e.g., each state maps uniquely to an action), whereas low complexity policies are more state independent. Similar to our derivation of rational inattention, we assume our agent is capacity limited, which induces a trade-off between policy complexity and reward. Agents must therefore compress the optimal policy if they lack the channel capacity to achieve error-free communication.

We can therefore define a joint optimization problem where the agent seeks to maximize reward subject to a capacity constraint. We define the optimal policy[2], π^*, as:

$$\pi^* = \underset{\pi}{\operatorname{argmax}}\, \beta V^\pi - I^\pi(S;A),$$

where V^π is the expected reward under policy π:

$$V^\pi = \sum_s P(s) \sum_a \pi(a|s) Q^\pi(s,a).$$

For analytic tractability, we assume that an agent either learns or has direct access to the action-value function $Q(s,a)$ which defines the expected reward after taking action a in state s. Note that we place the Lagrange multiplier, β, on V^π instead of $I^\pi(S;A)$ since it permits a more straightforward connection to reinforcement learning process models, as we will see.

Solving this equation yields the optimal policy, π^*:

$$\pi^*(a|s) \propto \exp[\beta Q(s,a) + \log P^*(a)],$$

[2]To facilitate direct comparison with rational inattention, we have left out Lagrange multipliers which ensure proper normalization (i.e., $\sum_a \pi(a|s) = 1$). See (Parush et al., 2011) for a full derivation.

where $P^*(a) = \sum_s \pi^*(a|s)$ is the optimal marginal action distribution. The optimal policy is the ubiquitous softmax function, used widely in the reinforcement learning literature; the Lagrange multiplier, β, plays the role of the inverse temperature parameter, governing the exploration-exploitation trade-off. Note that our derivation of the optimal policy made no appeal to exploration/exploitation, which instead arose as a natural consequence of resource constraints. The precise value of β is a function of the policy complexity:

$$\beta^{-1} = \frac{dV^\pi}{dI^\pi(S;A)}.$$

At low policy complexity, where $\frac{dV^\pi}{dI^\pi(S;A)}$ is steep, the optimal β is close to 0. In this regime, Q-values have minimal impact on the optimal policy and the marginal action distribution, $P^*(a)$ dominates. In other words, at low policy complexity, state-independent actions dominate, an insight we have previously used to explain perseveration (Gershman, 2020).

References

Abramson, L. Y., Metalsky, G. I., and Alloy, L. B. (1989). Hopelessness depression: A theory-based subtype of depression. *Psychological review*, 96(2):358.

Aleman, A., Lincoln, T. M., Bruggeman, R., Melle, I., Arends, J., Arango, C., and Knegtering, H. (2017). Treatment of negative symptoms: where do we stand, and where do we go? *Schizophrenia research*, 186:55–62.

Alloy, L. B., Olino, T., Freed, R. D., and Nusslock, R. (2016). Role of reward sensitivity and processing in major depressive and bipolar spectrum disorders. *Behavior therapy*, 47(5):600–621.

American Psychiatric Association (2022). Diagnostic and statistical manual of mental disorders, fifth edition, text revision (dsm-5-tr).

Ascher-Svanum, H., Nyhuis, A. W., Stauffer, V., Kinon, B. J., Faries, D. E., Phillips, G. A., Schuh, K., Awad, A. G., Keefe, R., and Naber, D. (2010). Reasons for discontinuation and continuation of antipsychotics in the treatment of schizophrenia from patient and clinician perspectives. *Current medical research and opinion*, 26(10):2403–2410.

Attneave, F. (1954). Some informational aspects of visual perception. *Psychological review*, 61(3):183.

Baldessarini, R. J., Perry, R., and Pike, J. (2008). Factors associated with treatment nonadherence among us bipolar disorder patients. *Human Psychopharmacology: Clinical and Experimental*, 23(2):95–105.

Balleine, B. W. and O'Doherty, J. P. (2010). Human and rodent homologies in action control: corticostriatal determinants of goal-directed and habitual action. *Neuropsychopharmacology*, 35(1):48–69.

Bari, B. A. and Gershman, S. J. (2023). Undermatching is a consequence of policy compression. *Journal of Neuroscience*, 43(3):447–457.

Beierholm, U., Guitart-Masip, M., Economides, M., Chowdhury, R., Düzel, E., Dolan, R., and Dayan, P. (2013). Dopamine modulates reward-related vigor. *Neuropsychopharmacology*, 38(8):1495–1503.

Berger, T. (1971). *Rate distortion Theory: A Mathematical Basis for Data Compression*. NJ: Prentice Hall.

Bhui, R., Lai, L., and Gershman, S. J. (2021). Resource-rational decision making. *Current Opinion in Behavioral Sciences*, 41:15–21.

Blanco, N. J., Otto, A. R., Maddox, W. T., Beevers, C. G., and Love, B. C. (2013). The influence of depression symptoms on exploratory decision-making. *Cognition*, 129(3):563–568.

Bossaerts, P., Yadav, N., and Murawski, C. (2019). Uncertainty and computational complexity. *Philosophical Transactions of the Royal Society B*, 374(1766):20180138.

Brown, V. M., Zhu, L., Solway, A., Wang, J. M., McCurry, K. L., King-Casas, B., and Chiu, P. H. (2021). Reinforcement learning disruptions in individuals with depression and sensitivity to symptom change following cognitive behavioral therapy. *JAMA psychiatry*, 78(10):1113–1122.

Bschor, T., Ising, M., Bauer, M., Lewitzka, U., Skerstupeit, M., Müller-Oerlinghausen, B., and Baethge, C. (2004). Time experience and time judgment in major depression, mania and healthy subjects. a controlled study of 93 subjects. *Acta Psychiatrica Scandinavica*, 109(3):222–229.

Burbridge, J. A. and Barch, D. M. (2007). Anhedonia and the experience of emotion in individuals with schizophrenia. *Journal of abnormal psychology*, 116(1):30.

Burt, D. B., Zembar, M. J., and Niederehe, G. (1995). Depression and memory impairment: a meta-analysis of the association, its pattern, and specificity. *Psychological bulletin*, 117(2):285.

Cassidy, F., Forest, K., Murry, E., and Carroll, B. J. (1998a). A factor analysis of the signs and symptoms of mania. *Archives of general psychiatry*, 55(1):27–32.

Cassidy, F., Murry, E., Forest, K., and Carroll, B. J. (1998b). Signs and symptoms of mania in pure and mixed episodes. *Journal of affective disorders*, 50(2-3):187–201.

Channon, S. (1996). Executive dysfunction in depression: the wisconsin card sorting test. *Journal of affective disorders*, 39(2):107–114.

Channon, S., Baker, J. E., and Robertson, M. M. (1993). Working memory in clinical depression: an experimental study. *Psychological medicine*, 23(1):87–91.

Chen, C., Takahashi, T., Nakagawa, S., Inoue, T., and Kusumi, I. (2015). Reinforcement learning in depression: A review of computational research. *Neuroscience & Biobehavioral Reviews*, 55:247–267.

Christopher, G. and MacDonald, J. (2005). The impact of clinical depression on working memory. *Cognitive neuropsychiatry*, 10(5):379–399.

Cohen, A. S. and Minor, K. S. (2010). Emotional experience in patients with schizophrenia revisited: meta-analysis of laboratory studies. *Schizophrenia bulletin*, 36(1):143–150.

Collins, A. G., Brown, J. K., Gold, J. M., Waltz, J. A., and Frank, M. J. (2014). Working memory contributions to reinforcement learning impairments in schizophrenia. *Journal of Neuroscience*, 34:13747–13756.

Coppen, A. (1967). The biochemistry of affective disorders. *The British Journal of Psychiatry*, 113(504):1237–1264.

Correll, C. U. and Schooler, N. R. (2020). Negative symptoms in schizophrenia: a review and clinical guide for recognition, assessment, and treatment. *Neuropsychiatric disease and treatment*, pages 519–534.

Crider, A. (1997). Perseveration in schizophrenia. *Schizophrenia bulletin*, 23(1):63–74.

Crowe, M., Wilson, L., and Inder, M. (2011). Patients' reports of the factors influencing medication adherence in bipolar disorder–an integrative review of the literature. *International journal of nursing studies*, 48(7):894–903.

Cusin, C., Yang, H., Yeung, A., and Fava, M. (2010). Rating scales for depression. *Handbook of clinical rating scales and assessment in psychiatry and mental health*, pages 7–35.

Dailey, M. W. and Saadabadi, A. (2018). *Mania*. StatPearls Publishing.

Denti, T., Marinacci, M., and Montrucchio, L. (2020). A note on rational inattention and rate distortion theory. *Decisions in Economics and Finance*, 43:75–89.

Devulapalli, K. K., Ignacio, R. V., Weiden, P., Cassidy, K. A., Williams, T. D., Safavi, R., Blow, F. C., and Sajatovic, M. (2010). Why do persons with bipolar disorder stop their medication? *Psychopharmacology bulletin*, 43(3):5.

Dickinson, A. (1985). Actions and habits: the development of behavioural autonomy. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, 308(1135):67–78.

Dowd, E. C. and Barch, D. M. (2010). Anhedonia and emotional experience in schizophrenia: neural and behavioral indicators. *Biological psychiatry*, 67(10):902–911.

Dubois, B. and Pillon, B. (1996). Cognitive deficits in parkinson's disease. *Journal of neurology*, 244:2–8.

Feldman, H. and Friston, K. J. (2010). Attention, uncertainty, and free-energy. *Frontiers in human neuroscience*, 4:215.

Forbes, N., Carrick, L., McIntosh, A., and Lawrie, S. (2009). Working memory in schizophrenia: a meta-analysis. *Psychological medicine*, 39(6):889–905.

Frank, M. J., Seeberger, L. C., and O'reilly, R. C. (2004). By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science*, 306(5703):1940–1943.

Franklin, J. C., Ribeiro, J. D., Fox, K. R., Bentley, K. H., Kleiman, E. M., Huang, X., Musacchio, K. M., Jaroszewski, A. C., Chang, B. P., and Nock, M. K. (2017). Risk factors for suicidal thoughts and behaviors: A meta-analysis of 50 years of research. *Psychological bulletin*, 143(2):187.

Friston, K. J., Shiner, T., FitzGerald, T., Galea, J. M., Adams, R., Brown, H., Dolan, R. J., Moran, R., Stephan, K. E., and Bestmann, S. (2012). Dopamine, affordance and active inference. *PLoS Computational Biology*, 8:e1002327.

Gailliot, M. T. and Baumeister, R. F. (2007). The physiology of willpower: Linking blood glucose to self-control. *Personality and social psychology review*, 11(4):303–327.

Gershman, S. (2021). *What makes us smart: The computational logic of human cognition*. Princeton University Press.

Gershman, S. J. (2020). Origin of perseveration in the trade-off between reward and complexity. *Cognition*, 204:104394.

Gershman, S. J. and Burke, T. (2022). Mental control of uncertainty. *Cognitive, Affective, & Behavioral Neuroscience*, pages 1–11.

Gershman, S. J., Horvitz, E. J., and Tenenbaum, J. B. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, 349(6245):273–278.

Gershman, S. J. and Lai, L. (2021). The reward-complexity trade-off in schizophrenia. *Computational Psychiatry*, 5(1).

Griffiths, T. L., Lieder, F., and Goodman, N. D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in cognitive science*, 7(2):217–229.

Grujic, N., Brus, J., Burdakov, D., and Polania, R. (2022). Rational inattention in mice. *Science advances*, 8(9):eabj8935.

Hack, L. M., Tozzi, L., Zenteno, S., Olmsted, A. M., Hilton, R., Jubeir, J., Korgaonkar, M. S., Schatzberg, A. F., Yesavage, J. A., O'Hara, R., et al. (2023). A cognitive biotype of depression and symptoms, behavior measures, neural circuits, and differential treatment outcomes: A prespecified secondary analysis of a randomized clinical trial. *JAMA Network Open*, 6(6):e2318411–e2318411.

Hamid, A. A., Pettibone, J. R., Mabrouk, O. S., Hetrick, V. L., Schmidt, R., Vander Weele, C. M., Kennedy, R. T., Aragona, B. J., and Berke, J. D. (2016). Mesolimbic dopamine signals the value of work. *Nature neuroscience*, 19(1):117–126.

Harmer, C. J., Duman, R. S., and Cowen, P. J. (2017). How do antidepressants work? new perspectives for refining future treatment approaches. *The Lancet Psychiatry*, 4(5):409–418.

Hauser, T. U., Iannaccone, R., Ball, J., Mathys, C., Brandeis, D., Walitza, S., and Brem, S. (2014). Role of the medial prefrontal cortex in impaired decision making in juvenile attention-deficit/hyperactivity disorder. *JAMA psychiatry*, 71(10):1165–1173.

Hick, W. E. (1952). On the rate of gain of information. *Quarterly Journal of experimental psychology*, 4(1):11–26.

Huys, Q. J., Pizzagalli, D. A., Bogdan, R., and Dayan, P. (2013). Mapping anhedonia onto reinforcement learning: a behavioural meta-analysis. *Biology of mood & anxiety disorders*, 3(1):1–16.

Ilonen, T., Leinonen, K.-m., Wallenius, E., Karlsson, H., Taiminen, T., Salokangas, R. K., Lauerma, H., and Tuimala, P. (2000). Impaired wisconsin card sorting test performance in first-episode severe depression. *Nordic Journal of Psychiatry*, 54(4):275–280.

Itti, L. and Baldi, P. (2009). Bayesian surprise attracts human attention. *Vision research*, 49(10):1295–1306.

John, C. E. and Jones, S. R. (2007). Voltammetric characterization of the effect of monoamine uptake inhibitors and releasers on dopamine and serotonin uptake in mouse caudate-putamen and substantia nigra slices. *Neuropharmacology*, 52(8):1596–1605.

Jones, S. D., Jones, M., Koldewyn, K., and Westermann, G. (2023). Rational inattention: A new theory of neurodivergent information seeking. *PsyArXiv*.

Joseph, R. (1999). Frontal lobe psychopathology: mania, depression, confabulation, catatonia, perseveration, obsessive compulsions, and schizophrenia. *Psychiatry*, 62(2):138–172.

Kasper, S. and Resinger, E. (2003). Cognitive effects and antipsychotic treatment. *Psychoneuroendocrinology*, 28:27–38.

Keller, A. S., Leikauf, J. E., Holt-Gosselin, B., Staveland, B. R., and Williams, L. M. (2019). Paying attention to attention in depression. *Translational psychiatry*, 9(1):279.

Kelley, M. E., van Kammen, D. P., and Allen, D. N. (1999). Empirical validation of primary negative symptoms: independence from effects of medication and psychosis. *American Journal of Psychiatry*, 156(3):406–411.

Kohler, C. G., Hoffman, L. J., Eastman, L. B., Healey, K., and Moberg, P. J. (2011). Facial emotion perception in depression and bipolar disorder: a quantitative review. *Psychiatry research*, 188(3):303–309.

Koops, S., Brederoo, S. G., de Boer, J. N., Nadema, F. G., Voppel, A. E., and Sommer, I. E. (2023). Speech as a biomarker for depression. *CNS & Neurological Disorders-Drug Targets (Formerly Current Drug Targets-CNS & Neurological Disorders)*, 22(2):152–160.

Kraepelin, E. (1921). *Manic-depressive insanity and paranoia*. E. & S. Livingstone.

Kring, A. M. and Moran, E. K. (2008). Emotional response deficits in schizophrenia: insights from affective science. *Schizophrenia bulletin*, 34(5):819–834.

Kuczenski, R. and Segal, D. S. (1997). Effects of methylphenidate on extracellular dopamine, serotonin, and norepinephrine: comparison with amphetamine. *Journal of neurochemistry*, 68(5):2032–2037.

Kwan, J. W., Bauer, I. E., Hautzinger, M., and Meyer, T. D. (2020). Reward sensitivity and the course of bipolar disorder: a survival analysis in a treatment seeking sample. *Journal of Affective Disorders*, 261:126–130.

Lader, M. (1993). Neuroleptic-induced deficit syndrome: old problem, new challenge. *Journal of Psychopharmacology*.

Lai, L. and Gershman, S. J. (2021). Policy compression: An information bottleneck in action selection. In *Psychology of Learning and Motivation*, volume 74, pages 195–232. Elsevier.

Lake, J. I. and Meck, W. H. (2013). Differential effects of amphetamine and haloperidol on temporal reproduction: dopaminergic regulation of attention and clock speed. *Neuropsychologia*, 51(2):284–292.

Lane, S. D., Cherek, D. R., Tcheremissine, O. V., Steinberg, J. L., and Sharon, J. L. (2007). Response perseveration and adaptation in heavy marijuana-smoking adolescents. *Addictive behaviors*, 32(5):977–990.

Lewis, R. L., Howes, A., and Singh, S. (2014). Computational rationality: Linking mechanism and behavior through bounded utility maximization. *Topics in cognitive science*, 6(2):279–311.

Lieberman, J. A., Stroup, T. S., McEvoy, J. P., Swartz, M. S., Rosenheck, R. A., Perkins, D. O., Keefe, R. S., Davis, S. M., Davis, C. E., Lebowitz, B. D., et al. (2005). Effectiveness of antipsychotic drugs in patients with chronic schizophrenia. *New England journal of medicine*, 353(12):1209–1223.

Lieder, F. and Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and brain sciences*, 43:e1.

Llerena, K., Strauss, G. P., and Cohen, A. S. (2012). Looking at the other side of the coin: a meta-analysis of self-reported emotional arousal in people with schizophrenia. *Schizophrenia research*, 142(1-3):65–70.

Mackowiak, B., Matejka, F., and Wiederholt, M. (2018). Rational inattention: A disciplined behavioral model. *CEPR Discussion Papers*, 13243.

Maćkowiak, B. and Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review*, 99(3):769–803.

Maier, S. F. and Seligman, M. E. (1976). Learned helplessness: theory and evidence. *Journal of experimental psychology: general*, 105(1):3.

Maier, S. F. and Seligman, M. E. (2016). Learned helplessness at fifty: Insights from neuroscience. *Psychological review*, 123(4):349.

Malapani, C., Deweer, B., and Gibbon, J. (2002). Separating storage from retrieval dysfunction of temporal memory in parkinson's disease. *Journal of cognitive neuroscience*, 14(2):311–322.

Malapani, C., Rakitin, B., Levy, R., Meck, W. H., Deweer, B., Dubois, B., and Gibbon, J. (1998). Coupled temporal memories in parkinson's disease: a dopamine-related dysfunction. *Journal of cognitive neuroscience*, 10(3):316–331.

Marneros, A. (2001). Origin and development of concepts of bipolar mixed states. *Journal of affective disorders*, 67(1-3):229–240.

Martin, D. J., Oren, Z., and Boone, K. (1991). Major depressives' and dysthymics' performance on the wisconsin card sorting test. *Journal of clinical psychology*, 47(5):684–690.

McCutcheon, R. A., Harrison, P. J., Howes, O. D., McGuire, P. K., Taylor, D. M., and Pillinger, T. (2023). Data-driven taxonomy for antipsychotic medication: a new classification system. *Biological Psychiatry*.

McGirr, A., Dombrovski, A. Y., Butters, M. A., Clark, L., and Szanto, K. (2012). Deterministic learning and attempted suicide among older depressed individuals: cognitive assessment using the wisconsin card sorting task. *Journal of psychiatric research*, 46(2):226–232.

Meltzer, H. Y. (1990). Role of serotonin in depression. *Annals of the New York Academy of Sciences*.

Mikhael, J. G. and Gershman, S. J. (2022). Impulsivity and risk-seeking as bayesian inference under dopaminergic control. *Neuropsychopharmacology*, 47:465–476.

Mikhael, J. G., Lai, L., and Gershman, S. J. (2021). Rational inattention and tonic dopamine. *PLoS computational biology*, 17(3):e1008659.

Miller, G. A. (1956). The magic number seven plus or minus two: Some limits on our capacity for processing information. *Psychological review*, 63:91–97.

Miller, K. J., Shenhav, A., and Ludvig, E. A. (2019). Habits without values. *Psychological review*, 126(2):292.

Miller, W. R. and Seligman, M. E. (1975). Depression and learned helplessness in man. *Journal of abnormal psychology*, 84(3):228.

Moritz, S., Stöckert, K., Hauschildt, M., Lill, H., Jelinek, L., Beblo, T., Diedrich, S., and Arlt, S. (2017). Are we exaggerating neuropsychological impairment in depression? reopening a closed chapter. *Expert Review of Neurotherapeutics*, 17(8):839–846.

Musslick, S., Dey, B., Ozcimder, K., Patwary, M. M. A., Willke, T. L., and Cohen, J. D. (2016). Parallel processing capability versus efficiency of representation in neural networks. *Network*, 8(7).

Niv, Y., Daw, N. D., Joel, D., and Dayan, P. (2007). Tonic dopamine: opportunity costs and the control of response vigor. *Psychopharmacology*, 191:507–520.

O'Bryan, R. A., Brenner, C. A., Hetrick, W. P., and O'Donnell, B. F. (2014). Disturbances of visual motion perception in bipolar disorder. *Bipolar disorders*, 16(4):354–365.

Oosterloo, M., Craufurd, D., Nijsten, H., and van Duijn, E. (2019). Obsessive-compulsive and perseverative behaviors in huntington's disease. *Journal of Huntington's disease*, 8(1):1–7.

Owens, M. J. and Nemeroff, C. B. (1994). Role of serotonin in the pathophysiology of depression: focus on the serotonin transporter. *Clinical chemistry*, 40(2):288–295.

Pagonabarraga, J., Kulisevsky, J., Strafella, A. P., and Krack, P. (2015). Apathy in parkinson's disease: clinical features, neural substrates, diagnosis, and treatment. *The Lancet Neurology*, 14(5):518–531.

Parush, N., Tishby, N., and Bergman, H. (2011). Dopaminergic balance between reward maximization and policy complexity. *Frontiers in Systems Neuroscience*, 5:22.

Ptacek, R., Weissenberger, S., Braaten, E., Klicperova-Baker, M., Goetz, M., Raboch, J., Vnukova, M., and Stefano, G. B. (2019). Clinical implications of the perception of time in attention deficit hyperactivity disorder (adhd): A review. *Medical science monitor: international medical journal of experimental and clinical research*, 25:3918.

Raison, C. L., Capuron, L., and Miller, A. H. (2006). Cytokines sing the blues: inflammation and the pathogenesis of depression. *Trends in immunology*, 27(1):24–31.

Robbins, T. W. and Cools, R. (2014). Cognitive deficits in parkinson's disease: a cognitive neuroscience perspective. *Movement Disorders*, 29(5):597–607.

Rose, E. and Ebmeier, K. (2006). Pattern of impaired working memory during major depression. *Journal of affective disorders*, 90(2-3):149–161.

Rutledge, R. B., Lazzaro, S. C., Lau, B., Myers, C. E., Gluck, M. A., and Glimcher, P. W. (2009). Dopaminergic drugs modulate learning rates and perseveration in parkinson's patients in a dynamic foraging task. *Journal of Neuroscience*, 29(48):15104–15114.

Ryu, V., Ha, R. Y., Lee, S. J., Ha, K., and Cho, H.-S. (2017). Behavioral and electrophysiological alterations for reinforcement learning in manic and euthymic patients with bipolar disorder. *CNS neuroscience & therapeutics*, 23(3):248–256.

Schiffer, W. K., Volkow, N., Fowler, J., Alexoff, D., Logan, J., and Dewey, S. (2006). Therapeutic doses of amphetamine or methylphenidate differentially increase synaptic and extracellular dopamine. *Synapse*, 59(4):243–251.

Schooler, N. R. (1994). Deficit symptoms in schizophrenia: negative symptoms versus neuroleptic-induced deficits. *Acta Psychiatrica Scandinavica*, 89:21–26.

Serpell, L., Waller, G., Fearon, P., and Meyer, C. (2009). The roles of persistence and perseveration in psychopathology. *Behavior therapy*, 40(3):260–271.

Shenhav, A., Musslick, S., Lieder, F., Kool, W., Griffiths, T. L., Cohen, J. D., and Botvinick, M. M. (2017). Toward a rational and mechanistic account of mental effort. *Annual review of neuroscience*, 40:99–124.

Shi, Z., Church, R. M., and Meck, W. H. (2013). Bayesian optimization of time perception. *Trends in cognitive sciences*, 17(11):556–564.

Sims, C. A. (2003). Implications of rational inattention. *Journal of monetary Economics*, 50(3):665–690.

Smith, A., Taylor, E., Warner Rogers, J., Newman, S., and Rubia, K. (2002). Evidence for a pure time perception deficit in children with adhd. *Journal of child psychology and psychiatry*, 43(4):529–542.

Steele, J., Kumar, P., and Ebmeier, K. P. (2007). Blunted response to feedback information in depressive illness. *Brain*, 130(9):2367–2374.

Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.

Trick, L., Watkins, E., Windeatt, S., and Dickens, C. (2016). The association of perseverative negative thinking with depression, anxiety and emotional distress in people with long term conditions: A systematic review. *Journal of psychosomatic research*, 91:89–101.

Valenstein, M., Blow, F. C., Copeland, L. A., McCarthy, J. F., Zeber, J. E., Gillon, L., Bingham, C. R., and Stavenger, T. (2004). Poor antipsychotic adherence among patients with schizophrenia: medication and patient factors. *Schizophrenia bulletin*, 30(2):255–264.

Verbaan, D., Marinus, J., Visser, M., van Rooden, S. M., Stiggelbout, A. M., Middelkoop, H. A., and van Hilten, J. J. (2007). Cognitive impairment in parkinson's disease. *Journal of Neurology, Neurosurgery & Psychiatry*, 78(11):1182–1187.

Volkow, N. D., Fowler, J. S., Wang, G.-J., Telang, F., Logan, J., Wong, C., Ma, J., Pradhan, K., Benveniste, H., and Swanson, J. M. (2008). Methylphenidate decreased the amount of glucose needed by the brain to perform a cognitive task. *PLoS One*, 3(4):e2017.

Woicik, P. A., Urban, C., Alia-Klein, N., Henry, A., Maloney, T., Telang, F., Wang, G.-J., Volkow, N. D., and Goldstein, R. Z. (2011). A pattern of perseveration in cocaine addiction may reveal neurocognitive processes implicit in the wisconsin card sorting test. *Neuropsychologia*, 49(7):1660–1669.

Wood, W. and Rünger, D. (2016). Psychology of habit. *Annual review of psychology*, 67:289–314.

Woodford, M. (2009). Information-constrained state-dependent pricing. *Journal of Monetary Economics*, 56:S100–S124.

Wu, Q., Nakazato, S., Yang, B., and Shimokawa, T. (2022). A neural basis of rational inattention models: consistency of cognitive cost with the mutual information criterion. *NeuroReport*, 33(15):649.

Yee, C. M., Mathis, K. I., Sun, J. C., Sholty, G. L., Lang, P. J., Bachman, P., Williams, T. J., Bearden, C. E., Cannon, T. D., Green, M. F., et al. (2010). Integrity of emotional and motivational states during the prodromal, first-episode, and chronic phases of schizophrenia. *Journal of Abnormal Psychology*, 119(1):71.

Yun, S., Yang, B., Anair, J. D., Martin, M. M., Fleps, S. W., Pamukcu, A., Yeh, N.-H., Contractor, A., Kennedy, A., and Parker, J. G. (2023). Antipsychotic drug efficacy correlates with the modulation of d1 rather than d2 receptor-expressing striatal projection neurons. *Nature Publishing Group US New York*.