

Junk, Numerosity, and the Demands of Epistemic Consequentialism

Forthcoming in Erkenntnis

Michal Masny
UC Berkeley
mmasny@berkeley.edu

Abstract: Epistemic consequentialism has been challenged on the grounds that it is overly demanding. According to the Epistemic Junk Problem, this view implies that we are often required to believe junk propositions such as ‘the Great Bear Lake is the largest lake entirely in Canada’ and long disjunctions of things we already believe. According to the Numerosity Problem, this view implies that we are frequently required to have an enormous number of beliefs. This paper puts forward a novel version of epistemic consequentialism which avoids these twin demandingness problems. The key is to recognise, first, that the final epistemic value of a true belief depends at least in part on the duration for which it is retained by the agent and, second, that our cognitive makeup places important constraints on which beliefs are retained and for how long.

Word count: 8,840

1. Introduction

Are we epistemically required to believe ‘junk’ propositions such as ‘Great Bear Lake is the largest lake entirely in Canada’, ‘the number of leaves on the tree by my window is 21,794’, or long disjunctions of things we already believe? Many people think that there is no such requirement. In fact, numerous philosophers accept an even stronger claim that we ought not to have such beliefs. As Gilbert Harman (1986, p. 12) puts it, “one should not clutter one’s mind with trivialities”.¹

Epistemic consequentialism seems unable to accommodate this sentiment. At least some versions of this theory tell us to maximise the number of true beliefs that we

¹ See also Goldman (1986), Christensen (1994), Williamson (1998), Feldman (2000), Wallace (2001), Sainsbury (2002), DePaul (2004), MacFarlane (2004), White (2005), Field (2009), Douven (2010), Wedgwood (2012), and Friedman (2018; 2019).

have. But a typical agent in typical circumstances can come to have all sorts of junk beliefs, for instance, by directing their attention to a particular thing in their surroundings, by making a simple inference, or by entertaining available evidence. Thus, epistemic consequentialism seems to require that we come to have such beliefs. This strikes many people as objectionably demanding. Call this *the Epistemic Junk Problem*.

There is a further, related worry. Given how easy it is to form new true beliefs, epistemic consequentialism appears to require that we have a paralyzingly large number of them. Here, the issue is not with the content of our attitudes, but with their sheer quantity. This also seems too demanding. Call this *the Numerosity Problem*.²

In this paper, I develop a novel version of epistemic consequentialism which promises to avoid these twin demandingness problems. My discussion is structured as follows. In Section 2, I outline key background assumptions and the structure of epistemic consequentialism, sharpen the twin demandingness problems, and address some rival responses.

In Section 3, I propose a novel solution to these problems. My first key claim is that junk beliefs and enormous sets of beliefs tend to have little final epistemic value. To defend this claim, I argue that the final epistemic value of a true belief covaries with the duration for which that belief is retained and draw on the findings of cognitive science concerning remembering and forgetting to show that junk beliefs and enormous sets of beliefs are typically retained only briefly. My second key claim is that junk beliefs and enormous sets of beliefs have great instrumental epistemic disvalue because they inhibit the retention of other beliefs and are epistemic dead ends. These two claims inform a novel version of epistemic consequentialism, *Time-Sensitive Epistemic Consequentialism*, that can accommodate the intuition that we are not required to have junk beliefs or an absurdly large number of beliefs.

Finally, in Section 4, I discuss the broader significance and plausibility of the proposed view. In particular, I examine the demands it places on atypical subjects in atypical circumstances, address the challenge that it renders beliefs replaceable, and suggest that it issues directives which are analogous to those implied by our best version of moral consequentialism.

² For recent discussions of these and similar problems, see e.g. Goldman (1999), Alston (2005), Grimm (2009), Friedman (2019), and Singer (2019).

2. Preliminaries

My inquiry concerns epistemic consequentialism and demandingness. A number of preliminary remarks will help us get clearer on this subject matter.

2.1. Epistemic consequentialism

Arguably, the fundamental question of normative epistemology is “What should I believe?”³ Epistemic consequentialism is a family of views which hold that one should have beliefs which promote epistemic value in an appropriate way. There are many ways to spell this out, but I think that it is helpful to think of these views as having an axiological component and a deontic component.

At the axiological level, we are concerned with epistemic value. Two main questions arise here. First, which entities are the bearers of final epistemic value? Many authors assume that attitudes play this role, and the usual suspects include true belief, knowledge, and understanding. For the purposes of my discussion, I will adopt the most popular of these views, namely that true beliefs have final epistemic value (and false beliefs have final epistemic disvalue, and nothing else has final epistemic value or disvalue—henceforth omitted).⁴

Second, what is the structure of epistemic value? One thing we would like to know here is whether value-bearers differ in terms of how much value they have. Many proponents and critics of epistemic consequentialism proceed on the assumption that all true beliefs are equally valuable.⁵ As we will see shortly, it is precisely this commitment which gives rise to the implication that having more true beliefs is always better, and thus to the Epistemic Junk Problem and the Numerosity Problem. In Section 3, I will argue that there are weighty reasons to reject it, which are, in fact, quite independent of the twin demandingness problems.

At the deontic level, we are concerned with promoting epistemic value. This aspect of epistemic consequentialism will not play an important role in the ensuing discussion, but nonetheless merits an outline. There are three key questions here. First, which entities are the objects of evaluation: attitudes, actions, dispositions, or something else? Second, how is epistemic value to be promoted: should we be maximising, satisficing, in the short run or in the long run, etc.? Third, which deontic

³ See, e.g., Berker (2013a) and Sylvan (2020) for this characterization.

⁴ See, for example, Sosa (2001), Goldman (2002), and Alston (2005).

⁵ Proponents of this view include Kvanvig (2003), Lynch (2004), and Horwich (2006).

properties are assigned to objects of evaluation: (epistemically) required or permissible or impermissible, rational or irrational, justified or unjustified?⁶

While my discussion will focus on the structure of epistemic value, it will be helpful to refer to a complete theory of epistemic consequentialism. Let's understand *Simple Consequentialism* as the view that true belief has final epistemic value, that all true beliefs are equally valuable, and that we are epistemically required to maximise epistemic value in the long run. This is a natural departure point for examining the prospects of epistemic consequentialism. Of course, this very view may not be the most popular option among contemporary epistemic consequentialists.⁷ However, my aim in this paper is not to criticize or defend Simple Consequentialism in particular, but rather to learn from its shortcomings and develop a superior alternative.

2.2. Epistemic junk

Paradigmatic *junk propositions* include trivia like 'the Great Bear Lake is the largest lake entirely in Canada' or 'the building across the street has 79 windows', and logical consequences of things we already believe. What these propositions have in common is that our subjective and objective interests are not served by having beliefs about them. That is, we do not find these propositions interesting and they are not conducive to our survival or successful completion of our projects. I will refer to a belief toward a junk proposition as a *junk belief*. Propositions and attitudes which are not junk are *important*.⁸

As noted earlier, for a typical agent in typical circumstances, there are all sorts of junk propositions which are believable. I will assume that a proposition is *believable* just in case one can form a belief about it, perhaps by directing their attention to a particular thing in their surroundings, by making a simple inference, or by considering easily

⁶ For somewhat different characterisations of the structure of epistemic consequentialism, see Berker (2013a), Dunn & Ahlstrom-Vij (2018), and Sylvan (2020a).

⁷ For example, some authors formulate epistemic consequentialism as a synchronic view (concerned with maximising epistemic value at a time rather than in the long run) in an attempt to avoid the so-called 'epistemic trade-offs' objection. But it's not clear that the synchronic view fares any better in this respect (see Berker 2013a) or, indeed, that epistemic trade-offs are always or even typically objectionable (see Singer 2018). Singer (2019) is among the few to endorse Simple Consequentialism in print. But there are also reasons to doubt that the views of some other self-proclaimed epistemic consequentialists are genuinely consequentialist. See fn. 12 for more on this.

⁸ The label 'junk' is due to Friedman (2018). Following Harman (1986), some philosophers use terms like 'clutter' and 'trivial propositions'. This is potentially misleading because 'trivial' is often understood as 'obvious', whereas not all junk propositions are obvious and not all obvious propositions are junk.

available evidence. Many logical consequences, especially disjunctions, of things which we already believe are junk and believable in this sense, and so are trivia and many things available to us through perception.⁹

It should be straightforward to see why Simple Consequentialism runs into the Numerosity Problem and the Epistemic Junk Problem. This view implies that *having more true beliefs is always better*. But the richness of our environments ensures that there are always new true beliefs that we could bring on board. And many or even most of one's jointly believable true propositions at any time are junk. So, we are required to have an enormous number of true beliefs, and we are required to have junk beliefs.

2.3. The pervasiveness of the twin demandingness problems

To appreciate the epistemological importance of the Numerosity Problem and the Epistemic Junk Problem, it is important to recognize that these issues are not exclusive to Simple Consequentialism.

For example, a popular version of *evidentialism* due to Thomas Kelly (2003) holds that one should believe all and only propositions which are supported by the totality of one's evidence.¹⁰ For a typical agent in typical circumstances, there are very many propositions which are supported by the totality of their evidence. And many of those are junk propositions, such as logical consequences of what they already believe or things available to them in perception. Thus, this view implies that we are required to believe all or many of these propositions.

Another form of evidentialism might seem more promising. Richard Feldman (2000) insists that the evidentialist injunction to respect the evidence is that one is merely permitted to believe all and only propositions supported by the evidence. In a similar fashion, *process reliabilism* popularized by Alvin Goldman (1979; 1986) says roughly that a belief is *permissible* if it is the product of a reliable process. Because these views do not issue requirements, they avoid the worry of excessive demandingness.

However, in addition to facing the charge of being *ad hoc*, these alternative views confront a problem which is closely related to the issues of demandingness. For a typical agent in typical circumstances, there are very many beliefs which are supported by the totality of their evidence or could be formed through reliable processes, and many of those are junk. According to reliabilism and Feldman's evidentialism, it is permissible to have all of these beliefs which is clearly incompatible

⁹ The background assumption here is that we are operating with a fine-grained notion of belief on which one does not automatically believe $p \ \& \ q$ when one believes both p and q .

¹⁰ Kelly (2003, p. 625) writes: "If, despite my utter lack of interest in the question of whether Bertrand Russell was left-handed, I stumble upon strong evidence that he was, then I have strong epistemic reasons to believe that Bertrand Russell was left-handed".

with Harman's popular injunction to avoid cluttering our brains. Thus, while these views are not excessively demanding, they are *excessively permissive* in the same range of cases.¹¹ This is hardly a consolation.

All in all, it looks like one cannot avoid the twin demandingness problems and their cognates just by switching corners of the epistemological ring.¹²

There are several strategies that epistemic consequentialists might adopt in an attempt to resolve the twin demandingness problems, but I believe that most of them are not going to work.

To begin with, epistemic consequentialists might deny that the twin demandingness problems are real problems. In a recent paper, Daniel Singer (2019, p. 272) observes that pursuing trivialities is a very inefficient way of acquiring true beliefs. For example, instead of counting leaves on a nearby tree or memorizing a phone book, beings with limited time and attention span like us who wish to maximise epistemic value would be better served by attending an online class on a topic they find interesting. Thus, the thought goes, Simple Consequentialism can accommodate the intuition that we are not epistemically required to have junk beliefs.

This response strikes me as deficient on two counts. First, while it can take hours to acquire certain junk beliefs, such as a true belief about the number of leaves on a nearby tree, there are innumerable other propositions that we could come to believe instantly: for example, by considering logical consequences of things we already believe or by directing our attention to our immediate surroundings. Second, holding considerations of acquisition efficiency fixed, it is clearly better to have some number of important beliefs rather than the same number of junk beliefs, and the Simple View cannot explain that. So, we are yet to get to the heart of the issue.

Alternatively, epistemic consequentialists might admit that these problems are real and seek to revise certain commitments of Simple Consequentialism. However, there are many dead ends here. For instance, it would not suffice to take knowledge rather than true belief as the bearer of final epistemic value. This is because, for a typical agent in typical circumstances, there are numerous knowable propositions at any time, many of which are junk. However, intuitively, we are generally not required to have such amounts or kinds of knowledge. Likewise, adopting a satisficing approach

¹¹ See Friedman (2018; 2019) for a recent discussion of excessive permissiveness.

¹² Many authors, such as Firth (1981) and Berker (2013a; 2013b), classify process reliabilism as a form of epistemic consequentialism (akin to rule-utilitarianism), in part because 'reliability' is often cashed out in terms of 'truth-conduciveness'. However, as Sylvan (2020b) and Dunn & Ahlstrom-Vij (2017) argue at length, this is a mistake: reliabilism lacks certain distinctive features of consequentialism and cannot be given plausible consequentialist foundations.

rather than a maximising approach would not suffice to avoid the Epistemic Junk Problem. For any plausible threshold for satisficing, we could imagine a case in which, to clear this threshold, the agent must believe many junk propositions.

These considerations lead me to believe that we have to re-examine a particular aspect of Simple Consequentialism: the structure of epistemic value. Specifically, it seems to me that the Epistemic Junk Problem and the Numerosity Problem arise because Simple Consequentialism accepts *the Simple View* that all true beliefs are equally valuable.

A number of philosophers have recognized that this view might be too simple and argued, for instance, that a belief's importance for the agent directly affects how much value that belief bears. Call this *the Importance View*.

There are at least two versions of this view. According to the 'threshold' version, junk beliefs have no final epistemic value and important beliefs have equal value. According to the 'increasing' version, junk beliefs have no final epistemic value, and the final epistemic value of a true belief is an increasing function of its importance for the agent. Both versions have their proponents in the literature.¹³

The Importance View certainly vindicates the intuition that we are not epistemically required to form junk beliefs: since they bear no value, there are no reasons to have them. It therefore avoids the Epistemic Junk Problem.

However, this line of response strikes me as unsatisfactory, on three counts. First, the Importance View has some very odd implications, at least if importance is to be understood subjectively. To mention just one: it would be an enormous epistemic tragedy if everyone suddenly lost interest in Leo Messi. After all, the numerous true beliefs that millions of football fans around the world have about his career and personal life would, on this view, completely lose their value.

Second, as Marian David (2001) and Stephen Grimm (2009) discuss at length, it is difficult to see how the Importance View could make sense of the full spectrum of our epistemic appraisals of junk beliefs. For one thing, even if we have no epistemic reason to have *true* junk beliefs, we seem to have an epistemic reason not to have *false* junk beliefs. But if true junk beliefs have no final value, then considerations of symmetry suggest that false junk beliefs have no final disvalue either, and there appears to be no non-instrumental epistemic reason not to have false junk beliefs.

¹³ These are two variants of what Grimm (2009) calls 'the restricted view'. Goldman (1999; 2002) and Alston (2005) seem to embrace the 'threshold' version. Ahlstrom-Vij and Grimm (2013) and Hu (2016) appear to endorse the 'increasing' version.

For another, we also want to be able to assess whether our beliefs are justified or well-formed, and junk beliefs are no exception in this regard. With this in mind, consider the following example from Grimm:

Suppose that on a lazy whim you decide to scan your desktop for motes of dust. After a bit of distracted counting, you then conclude that the desktop is harboring eighteen motes. Given the sloppiness of your method, however, we can suppose that this answer really amounted to little more than a guess: you might very easily have concluded, for example, that there were rather more motes or rather fewer. (2009, p. 250)

Your belief is clearly unjustified in this case. But a proponent of the Importance View will have trouble explaining this. After all, within a consequentialist framework, beliefs are meant to inherit their status from their relationship to value, but since we are dealing with a junk proposition here, there is nothing of final epistemic value to derive this status from.

Finally, note that postulating that importance matters is not enough to address the Numerosity Problem. This is because we often find ourselves in situations in which there are a vast number of things that we have an interest in believing, and so we would be required to believe all such propositions.

Of course, proponents of the Importance View might have things to say in response to these three objections and offer additional arguments in favour of their view.¹⁴ But I think that the presented considerations suffice to motivate a search for a consequentialist solution of a different kind.

3. The solution

3.1. Overview

Baldly stated, my proposed solution to the twin demandingness problems comprises the following simple argument:

1. Junk beliefs and enormous sets of beliefs have little final epistemic value.

¹⁴ For example, as one reviewer suggests, proponents of the Importance View might argue that our cognitive faculties are at least roughly attuned to epistemic value, and so our tendency to remember important beliefs and forget junk beliefs provides evidence that only important beliefs have epistemic value. In response, let me just note that the first premise brings on board certain controversial meta-epistemological commitments, such as that having cognitive faculties which are attuned to epistemic value is an adaptive trait, that I cannot hope to adequately address here.

2. Junk beliefs and enormous sets of beliefs have great instrumental epistemic disvalue.
3. Therefore, we are not epistemically required to have junk beliefs or enormous sets of beliefs.

In this section, I will clarify and defend both premises of this argument. As should be clear from the preceding, neither of these premises is completely novel as a standalone claim. But we are yet to see a compelling ground for accepting the first claim, and this is precisely what I seek to provide.

Of course, for the argument to clear the standard of logical validity, we must also bring on board the consequentialist framework outlined earlier, i.e. the claim that we are epistemically required to maximise epistemic value in the long run. And we have to stipulate that having junk beliefs and enormous sets of beliefs is not the only option available to us (which is an accurate description of our epistemic environment, though we will take a second look at it in Section 4).

3.2. Time Sensitive-Consequentialism

According to the Simple View, all true beliefs have equal final value. I believe that this view is flawed. Epistemic value has an important temporal dimension: the duration for which a true belief is retained affects how much value it bears. This claim finds support in two sets of considerations which are, in fact, completely independent of the twin demandingness problems.

The first set comprises our considered judgments about particular cases. Here is a quick example. Take your favourite scientific or philosophical proposition. Make sure that it is exciting and important. Here is one: the universe is expanding at an accelerating rate. Suppose that this is true and that true beliefs are the bearers of final epistemic value. Now, consider

MANY YEARS: Salma comes to believe that the universe is expanding at an accelerating rate and retains this belief for many years.

Intuitively, in this case, Salma's belief is of great final value. But now compare

BRIEF MOMENT: Salma comes to believe that the universe is expanding at an accelerating rate. Unfortunately, the doorbell rings shortly after and Salma is distracted by a chatty neighbour. As a result, Salma completely forgets that the universe is expanding at an accelerating rate and never comes to have that belief again.

Salma's true belief in BRIEF MOMENT clearly has little or no final value. This is so even though it scores very high in terms of objective and subjective importance. In

fact, we could stipulate that this belief scores very highly in all plausible dimensions of evaluation other than duration. This suggests that the duration for which a belief is retained affects its value.

The thought that duration matters is further supported by cases which involve compensation between epistemic value and disvalue. Consider

LAPSE: Salma forms a true belief that the universe is expanding at an accelerating rate and holds it for a very long time, say, ten years. At the end of this period, however, Salma comes to believe that the universe is *not* expanding at an accelerating rate instead. She holds this false belief for a fairly short time, say, one hour.

I find it implausible that the disvalue of the brief period of having a false belief that not- p could nullify the positive value of the many years of having a true belief that p . But this is among the implications of the view that the final epistemic value of true beliefs (and the disvalue of false beliefs) does not vary with the duration for which they are held. Surely, it is better that Salma had the true belief for ten years and then the false belief for one hour than if she had neither. Likewise, it would be implausible to say that the value of the brief episode of having a true belief that p compensates for the disvalue of the many years of falsely believing that not- p .

Cases like these easily multiply. Instead of providing further examples, let me briefly note that we could easily swap true belief for knowledge, understanding, or any other candidate value-bearing attitude. Thus, the presented considerations support the more general claim that epistemic value has a temporal dimension regardless of what the value-bearing attitudes are. However, it will ease the exposition of the argument of this paper to focus just on true beliefs.¹⁵

The second set of considerations that support the claim that the final value of a true belief varies with the duration for which it is retained has to do with the fact that some other, widely-recognized value bearers have this property. In particular, most philosophers think that it is better to have *prudential goods* for a longer amount of time rather than a shorter amount of time. For example, it is better to experience an hour

¹⁵ It will also be helpful to explicitly assume a particular view about the nature of belief: *representationalism*. On this view, to have a belief is to stand in a particular relation to a mental representation. See, for example, Fodor (1968), Millikan (1993), and Burge (2010). Its most prominent rival is *dispositionalism* which holds that believing a proposition is just a matter of having a certain set of dispositions. See, for example, Ryle (1949), Stalnaker (1984), and Schwitzgebel (2013). At least some versions of dispositionalism imply that we have infinitely many beliefs at any time. Any such view seems unable to serve as a foundation for a theory of epistemic consequentialism because, in any given case, the epistemic value of the outcome will be infinite.

of pleasure rather than just a minute. Likewise, a longer friendship seems more valuable for a person than a shorter one, other things equal. This lends some plausibility to the thought that it is also better to have *epistemic goods*, such as true beliefs, for a longer rather than a shorter period.¹⁶

These two arguments, one case-based and one symmetry-based, shall suffice for now. The remainder of the paper puts the thesis that duration matters to a further test by examining whether it can be used to address the Epistemic Junk Problem and the Numerosity Problem. But we must first sharpen this thesis.

To that end, let's get clearer on *how* exactly duration matters. There are two main candidate views. The first—to my mind, more natural—view is that the final epistemic value of a true belief is an increasing function of the duration for which it is retained. In other words, the longer you have a true belief, the more valuable it is. Analogously for the disvalue of false beliefs. Call this *the Increasing View*.¹⁷ The second candidate view is that true beliefs which are not retained for long enough are not valuable at all, whereas all true beliefs which are retained for long enough are equally valuable, other things equal. Analogously for false beliefs. Call this *the Threshold View*.¹⁸

I think that there are weighty, case-based reasons to embrace the Increasing View. First, the Threshold View gives the wrong verdict about the following comparison:

INTERMITTENT BELIEVING: Salma comes to truly believe that the universe is expanding at an accelerating rate, retains that belief for some time and then forgets it completely. After several weeks, she independently comes to

¹⁶ What about aesthetic value? I'm not sure. According to one view, the primary bearers of such value are aesthetic experiences, and artefacts such as paintings are valuable only in the sense that they have the potential to afford such experiences (Beardsley 1958; Budd 1995; Dickie 1998). If correct, this view would lend further support to my symmetry argument. After all, other things equal, a longer aesthetic experience seems better than a shorter one. However, according to another view, the primary bearers of aesthetic value are certain properties of artefacts such as paintings (Sibley 1983; Zangwill 1984; Goldman 1995). In this context, one could also claim that, other things equal, an older painting is more aesthetically valuable than a newer one, but this intuition is admittedly less prevalent. Thus, pending further inquiry, I think that comparisons with aesthetic value neither confirm nor undermine my thesis that epistemic value has a temporal dimension.

¹⁷ This is analogous to Bentham's quantitative hedonism, according to which the non-instrumental value of an episode of pleasure depends in part on its duration (and in part on its intensity).

¹⁸ Of course, both views should be understood in 'other things equal' terms. We want to leave open the possibility that there are other features which also affect the attitude's final value.

have the same belief again. This cycle repeats a few times over the course of two years.

CONTINUOUS BELIEVING: Salma comes to truly believe that the universe is expanding at an accelerating rate and retains that belief for two years.

Intuitively, continuous believing is more valuable than intermittent believing. The Increasing View delivers this verdict, but the Threshold View is unable to do so. For this latter view implies either that intermittent believing is more valuable (if it is to be interpreted as having many true beliefs, each for long enough) or else that intermittent and continuous believing are equally valuable (if intermittent believing is interpreted as having one true belief for long enough).

Second, the Threshold View appears to get the negative side of things wrong. On this view, an attitude must be retained for long enough to have final epistemic disvalue. But this can't be right. Even briefly held false beliefs seem to be disvaluable, and thus subject to epistemic appraisal.

Finally, only the Increasing View is supported by the abductive argument of the remainder of this paper. For reasons that will become apparent shortly, the Threshold View cannot be used to deliver a fully satisfactory answer to the Numerosity Problem and the Epistemic Junk Problem.

We can refer to the version of epistemic consequentialism committed to the Increasing View as *Time-Sensitive Consequentialism*.

3.3. Remembering, forgetting, and the role of belief content and quantity

The initial step of my argument for the first premise is in place: we now understand why and how the duration for which a belief is retained matters with respect to its epistemic value. We will now examine what it is for a belief to be retained, and how the content and number of beliefs affect the duration for which they are retained. To that end, we need to review a set of findings from cognitive science.¹⁹

Let's start by thinking about remembering and forgetting. Memory is not a uniform faculty. It is common to distinguish between *declarative memory*, which is the memory of information and events, and *procedural memory*, which is the memory of skills.²⁰ Declarative memory is in turn divided into *semantic memory* and *episodic memory*. According to a popular view, semantic memory stores records with primarily

¹⁹ For helpful overviews of empirical and philosophical perspectives on memory that I draw upon here, see Bjork (2011), Frise (2015; 2018), Michaelian (2011), Norby (2015), and Squire (2009).

²⁰ This distinction was introduced by Cohen & Squire (1980).

propositional contents, whereas episodic memory stores records of past experiences involving non-propositional contents. For the purposes of this argument, I am going to assume that this view is correct and focus on semantic memory (henceforth just ‘memory’).²¹

What are the records stored in memory? It is not clear. But since remembering a proposition apparently involves believing that proposition, I think that we can safely talk about storing beliefs.²² It is helpful to distinguish between two types of beliefs in this context. First, there are *occurrent beliefs*: when you tell someone that it was raining yesterday, the belief that it was raining yesterday is ‘before your mind’. Second, there are *non-occurrent* beliefs: when you are asleep or occupied with a game of chess, the belief that it was raining earlier is plausibly still stored in your memory. Plausibly, most of the beliefs stored in one’s memory at any given time are non-occurrent.

What is it for a belief to be stored in memory? Beliefs first enter *working memory* which can be understood as a sort of mental workspace. The vast majority of our beliefs, especially those formed through perception, end their journey there. But some of our beliefs go on to be encoded in *long-term memory*.

Forgetting can occur at various stages of this process. First, beliefs that enter working memory might fail to be encoded in long-term memory. Second, beliefs that are encoded in long-term memory might fail to achieve a stable form. Third, beliefs that are encoded in long-term memory in a stable form might be forgotten. In the first case, we forget non-encoded beliefs. In the latter two, we forget encoded beliefs.

Encoded beliefs can be forgotten in two ways. Sometimes beliefs become *unavailable*: i.e., a stored record is permanently eliminated. This can happen as a result of permanent damage to the memory system (e.g., brain injury), interference (e.g., alcoholic blackout), or cognitive decline (e.g., amnesia). It can also happen ‘naturally’, but empirical research suggests that records stored in long-term memory rarely become unavailable. More often, forgetting encoded beliefs is a matter of them becoming *inaccessible*. This involves being unable to retrieve a record in response to

²¹ For the view that the contents of episodic memory are also primarily or exclusively propositional, see, for example, Fernandez (2006) and Byrne (2010).

²² Why think that remembering involves believing? First, some support comes from Moore-paradoxical sentences such as ‘S remembers that it was raining yesterday, but S does not believe that it was raining yesterday’. Second, Moon (2013) argues that (i) if S remembers that *p*, then S can use *p* as a premise in justifying inferences, and (ii) a premise can be used in justifying inferences only if it is believing.

appropriate stimuli. Roughly, if you try to retrieve some record from your memory and fail, then that record counts as inaccessible.²³

Against this background, we can return to the question of which beliefs count as retained. A belief is retained just in case it is (i) occurrent ('before the subject's mind') or (ii) available (stored in long-term memory) and accessible (can be retrieved in response to appropriate stimuli).²⁴

Thus, according to the Increasing View, the final epistemic value of a true belief increases with the duration for which that belief is occurrent or available and accessible. One helpful way to think about this is that true beliefs accumulate final epistemic value for as long as the agent has access to them.^{25, 26}

With this in place, let's now examine how the content and quantity of beliefs affect the processes of remembering and forgetting. Intuitively, we come to remember things which are of interest to us and forget things that are not of interest to us. This thought finds confirmation in the findings of cognitive science. Although the details of specific theories differ, psychologists agree that the processes of encoding records and rendering records inaccessible are in a curious way sensitive to our subjective interests (e.g. what we find interesting) and objective interests (e.g. what is conducive to our survival).²⁷

Let's start with rendering encoded information inaccessible. According to the influential *new theory of disuse*, interests play a key role in this process. Bjork & Bjork summarise the core of this view as follows:

²³ The distinction between availability and accessibility is due to Tulving & Pearlstone (1966). Since McGeoch (1932), it is widely accepted that forgetting of encoded beliefs typically takes the form of inaccessibility. See also Anderson et al. (1994).

²⁴ We need both conditions because, as we have seen, not all occurrent beliefs are committed to long-term memory, and not all available and accessible beliefs are occurrent.

²⁵ There are other possible ways of losing true beliefs. For example, an agent might acquire misleading evidence which would prompt her to form a false belief instead, or acquire high-order evidence that would lead her to suspend judgment on the matter, or lose grip on a complex proposition which is the object of her belief.

²⁶ The presented account can also accommodate the possibility of storing beliefs in one's *extended memory*, such as a notebook or a smartphone (see Clark & Chalmers 1998). A belief would then count as available if it was recorded in one's extended memory and as accessible if it was retrievable from there. It would accumulate epistemic value at times at which it were accessible to the agent.

²⁷ Bjork (2011) provides a helpful survey of these findings.

In general, the theory of disuse ... says that the items in memory that are readily accessible to us are those items that we have been using (retrieving) lately. ... [T]hat will typically be adaptive. The items that have been retrieved frequently in the recent past will tend to be those items most relevant to our current interests, problems, goals, and station in life. On a statistical basis, those same items will be maximally relevant in the future as well. ... So, in general those things that we are likely to need to recall in the near future will be accessible to us, and those things that are irrelevant or interfering or out-of-date will be inaccessible. (1988, pp. 285-6)²⁸

That is, whether a record will remain accessible is largely determined by the frequency of retrieval, which is, in turn, largely determined by our interests. Records which are not interesting or useful will be selected against and inhibited. It follows that junk beliefs will not remain accessible for long, even if they are encoded in long-term memory.

Our interests also play an important role in determining which beliefs are encoded in the first place. Whether some record is encoded depends on the following factors, among others: (i) the extent to which we perceive the information as relevant to our survival,²⁹ (ii) the degree to which we direct our attention to the object of the belief,³⁰ (iii) the intensity of emotional response it elicits in the agent,³¹ and (iv) how well connected it is to the records that are already stored. Junk beliefs score very low in these respects. Propositions such as ‘the Great Bear Lake is the largest lake entirely in Canada’, ‘the building across the street has 79 windows’, or long disjunctions of things we already believe are not exciting to us or sufficiently connected to our wider body of beliefs.

The upshot is that junk beliefs are normally either not encoded in long-term memory or else very quickly rendered inaccessible, and so retained only briefly. There are, of course, exceptions to this rule. We all have some junk beliefs that we happen to have retained for a very long time. But, as we will see, all that matters for the present purposes is that we typically retain such beliefs for a brief time only, and so they tend to have little epistemic value.

²⁸ See also Anderson (2007) for a discussion of the view that the memory system estimates on the basis of past experience which records are likely to be relevant in the future.

²⁹ This mechanism is known as *survival processing*. For landmark studies, see Nairne et al. (2007; 2008), and for a recent meta-analysis, see Scofield et al. (2017).

³⁰ See Chun & Turk-Browne (2007).

³¹ See Brown & Kulik (1977) and Luminet & Curci (2009) for *flashbulb memories*: vivid memories of finding out an important piece of news. See also Hunt (2003) for the role of distinctiveness and surprise, and Kensinger & Schacter (2008) for the role of arousal in encoding.

Similar considerations apply to very large sets of beliefs. Psychologists agree that there is a limit on how many things we can commit to long-term memory at any time.³² If we tried to memorise too much at once, we would, in fact, fail to memorise almost anything—regardless of whether these beliefs would be junk or important. Thus, enormous sets of beliefs are also retained only briefly and have little epistemic value.³³

3.4. The second premise

In the previous section, I presented some reasons for supporting the first premise of the core argument: that junk beliefs and enormous sets of beliefs generally have little final epistemic value. This goes a long way towards vindicating the initially troubling intuition that we are not required to have these attitudes. But to complete my solution to the twin demandingness problems, we also need to identify the reasons we have against forming junk beliefs and enormous sets of beliefs.

One important set of considerations has to do with our cognitive limitations. First, as we have already seen, there are encoding constraints: for example, there is only so much we can commit to memory at once. Second, there is a limit on how many records can remain accessible at any time. On this point, Bjork notes that

As we learn new information, procedures, and skills, we create the potential for competition with related information, skills, and procedures that already exist in memory. (2011, pp. 3-4)

In short, retrieving some information can cause forgetting of other information. This mechanism is known as *retrieval-induced forgetting*.³⁴ Third, there are retrieval constraints. In particular, we have limited time and cognitive resources to evaluate which records are relevant to our present situation. If we never forgot anything, many queries would produce a paralysing amount of records, only some of which would

³² The basic idea here is that records are first encoded in working memory, which has limited capacity and therefore acts as a bottleneck for encoding in long-term memory. See, for example, Cowan (2001) and Forsberg et al. (2021).

³³ It is worth emphasising here that my claim in this section is not (i) that we cannot form or retain junk beliefs or enormous sets of beliefs, but rather (ii) that such beliefs would be very short-lived. If (i) were true, then the conclusion that we are not required to have junk beliefs or enormous sets of beliefs would straightforwardly follow from the ‘ought implies can’ principle. But (i) is not true.

³⁴ See Anderson, Bjork, and Bjork (1994) for a landmark study and Murayama et al. (2014) for a recent meta-analysis.

be directly relevant to the query.³⁵ Thus, having each new belief comes at a direct epistemic cost, as it inhibits retention of other beliefs, and this cost is especially large in the case of enormous sets of beliefs.

In the case of junk beliefs, there are also indirect epistemic costs that we have to recognize. Paradigmatic junk beliefs such as ‘the Great Bear Lake is the largest lake entirely in Canada’, ‘the number of leaves on the tree by my window is 21,794’, or long disjunctions of things we already believe are epistemic dead ends. They do not serve as evidence for or otherwise aid the retention of our other beliefs.

The opposite is true of many important beliefs. The proposition that the universe is expanding at an accelerating rate can play an important role in larger belief systems and theories, open up new avenues for inquiry, and give support to further important propositions which we can come to believe. It can also reinforce some beliefs one already has by improving their accessibility and resilience to misleading evidence. Thus, junk beliefs come at a significant opportunity cost.

In virtue of these direct and indirect effects on our epistemic lives, junk beliefs and enormous sets of beliefs have great instrumental epistemic disvalue. This is the second premise of the core argument.

3.5. Putting the pieces together

Time-Sensitive Consequentialism holds that the final epistemic value of a true belief is an increasing function of the duration for which that belief is retained. As we have seen, typically, junk beliefs and enormous sets of beliefs are retained for a brief moment only, and so have little final epistemic value. Moreover, junk beliefs and enormous sets of beliefs generally have great instrumental epistemic disvalue. This vindicates the original intuition that, in typical circumstances, we are not required to have junk beliefs or an enormous number of beliefs. Instead, maximising epistemic value is best served by having relatively few beliefs, all of which are important. After all, one upshot of the preceding discussion is that a single true belief retained for a month is more valuable than hundreds retained for a few seconds each. Taken together, these considerations show that Time-Sensitive Consequentialism avoids the twin demandingness problems. In the remainder of the paper, I discuss the broader significance of this solution and potential challenges.

³⁵ Cases of hyperthymia are particularly interesting in this context. Parker et al. (2006) extensively studied a subject known as ‘AJ’ who is described as having “an extraordinary capacity to recall specific items from their personal past” (p. 47). One of the apparent downsides of this condition is that AJ would spend vast amounts of time and cognitive resources navigating through the large quantities of retrieved memories.

4. Further issues

4.1. Atypical agents in atypical circumstances

According to Time-Sensitive Consequentialism, the intuitive claim that we are not required to form junk beliefs is true when typical agents in typical circumstances are considered. However, if circumstances are atypical or if the agent is atypical, then one may well be required to have some junk beliefs or to have an enormous number of true beliefs.

To illustrate this, consider three highly atypical kinds of agents. First, take someone who does not have the capacity to store beliefs in long-term memory.³⁶ For this person, all beliefs would be ephemeral. Second, take someone with an extraordinary capacity to retain all beliefs, regardless of their contents. In either case, all beliefs would be retained for an equally long time, and so, according to the Increasing View, they would be equally valuable. Thus, there would be no reason to prioritize important beliefs over junk beliefs, and these atypical agents could be required to have junk beliefs in some circumstances. There would also be no reason to keep the pool of one's beliefs small at any given time, so these atypical agents could sometimes be required to have very large numbers of beliefs.

Third, we might imagine an agent for whom the process of remembering and forgetting works reversely to ours. This unfortunate person regularly retains junk beliefs and forgets important beliefs. Since the Increasing View implies that their junk beliefs typically have more final value than their important beliefs, they would be required to have junk beliefs in a wide range of circumstances.

Are these implications troubling for Time-Sensitive Consequentialism? I don't think so. For one thing, the intuition that *we* are not required to form junk beliefs and the intuition that *we* are not required to have an enormous number of beliefs does not (or, at least, should not) extend to beings who are cognitively highly dissimilar to us. To take the most striking example, there does not seem to be anything counterintuitive about the claim that an omnipotent being would be epistemically required to believe *all* truths, including trivia and logical consequences of all propositions. Likewise, it strikes me as plausible that a being with extremely limited cognitive capacities would not be required to prioritise important beliefs over junk beliefs. The first two agents are, in relevant respects, like that.

What about the third agent who forgets all important things but can remember trivialities? I think that the point about the applicability and reliability of our

³⁶ The well-known case of the patient 'HM' fits this description. The first observations were reported by Scoville and Milner (1957).

intuitions when dealing with atypical agents is relevant here as well. But, beyond that, I think it is worth reflecting more abstractly on what kind of epistemic life we would want an agent with that kind of psychological makeup to have. To me, it seems like a better and more rewarding use of their cognitive capacities to cultivate a set of beliefs about trivialities instead of trying and failing, day after day, to commit important beliefs to memory.

Indeed, I think that it is a major virtue—and not a vice—of Time-Sensitive Consequentialism that what is epistemically required of an agent depends in part on their capacities and external circumstances.

4.2. Replaceability

Another apparent problem with Time-Sensitive Consequentialism is that this view renders beliefs replaceable, in the sense that it does not discriminate between a single belief retained continuously for a long period of time and an appropriately long sequence of beliefs retained for a brief moment each. To make this challenge more vivid, we might imagine the following scenario.

TRIVIA MACHINE: You are hooked up to a fancy TV set that continuously broadcasts short documentary clips about random matters, TikTok-style. Because the amount of information is overwhelming, the corresponding beliefs are neither interesting nor useful. And because you are a typical agent, beliefs formed on the basis of the presented evidence are held only briefly.

We can then compare the TRIVIA MACHINE life with an ordinary epistemic life in which some beliefs are ephemeral while others are retained for a long time. If the Increasing View (to which Time-Sensitive Consequentialism is committed) is correct, and we get the numbers right, it follows that we should be indifferent between these two epistemic lives. This may be counterintuitive. Call this *the Replaceability Problem*.

How troubling is this implication? First, although it may provide little consolation for those moved by this problem, I would like to note that the Increasing View handles such cases much better than the Simple View. On the latter view, we wouldn't be indifferent between these two epistemic lives. Rather, the TRIVIA MACHINE life would be decidedly better because the number of true beliefs formed in that scenario would be much greater.

Note also that the Importance View—which holds that the value of a belief depends directly on its importance—faces the same issue. To be sure, this view avoids the implication that we should be indifferent between a life in the trivia machine and an ordinary life because you would have only junk beliefs in the former case. However, we can easily construct an analogous scenario in which you are bombarded with information that is at least minimally interesting and useful, and so the resulting

beliefs are not junk, but they are nonetheless retained only briefly. We can call this *the Non-Trivia Machine*, for the lack of a better name. According to the Importance View, such a life would also be decidedly better than an ordinary life in which some beliefs are ephemeral while others are retained for a long time.

Second, as the preceding suggests, the Replaceability Problem does not appear to be a challenge to the Increasing View *per se*. Instead of rejecting this view, we should perhaps revisit some other tenets of Time-Sensitive Consequentialism.

For example, we could instead reject the assumption that epistemic value is additive. In TRIVIA MACHINE, the agent is presented with random snippets of information. A typical epistemic life looks very different. Not only it has a narrative of some kind, but also the beliefs are themselves organised into theories and systems. Perhaps it matters that beliefs come together in these ways. Perhaps when they do, they form an organic unity whose value is more than the value of the mere sum of its parts.

Of course, more work must be done to determine whether the additivity of epistemic value should be rejected. But the important point, for now, is that the Replaceability Problem need not be a threat to the Increasing View, which plays the key role in the presented solution to the Epistemic Junk Problem and the Numerosity Problem.

4.3. Our epistemic predicament

In closing, I would like to highlight that Time-Sensitive Consequentialism issues demands which are analogous to those that we get from the best version of moral consequentialism. The following dramatization will help bring this out.

SAILING THROUGH THE STORM: You are the captain of a small ship caught at sea in a violent storm. You spot a large group of people in the water. There are more of them than you could possibly bring on board. Some of them, you can tell, lack the strength to withstand the waves crushing the deck; if pulled out, they would be immediately swept back into the sea. There are also other people already aboard your ship who are threatened by the waves. Attending to either group may impede your ability to help the other.

I think that this scenario provides a useful metaphor for our epistemic predicament. The richness of our environment ensures that there are always new true beliefs that we can bring on board. And our cognitive limitations make it the case that some of these beliefs, as well as others that we already have, are in danger of being quickly forgotten or otherwise lost. Any choice we make in these circumstances comes at a price.

This is a difficult situation to be in, but I believe Time-Sensitive Consequentialism helps us get through it. We only have a weak reason to bring on board attitudes which

would be immediately forgotten or otherwise lost. And we have a weighty reason against forming beliefs which would inhibit the retention of other beliefs or simply fail to contribute to our further epistemic efforts.

This is analogous to what we would expect from moral consequentialism. For what matters in the ethical case is not how many people we pull out of the water, but rather how many survive. Likewise, what should matter in the epistemic case is not how many true beliefs we come to have, but rather how many we manage to preserve. It is a virtue of Time-Sensitive Consequentialism that it captures this insight.

5. Concluding remarks

Epistemic consequentialism has been challenged on the grounds that it implies that we are required to have an enormous number of true beliefs and that we are required to have junk beliefs. But these critiques stem from the failure to appreciate the complexity of issues concerning the structure of epistemic value and the findings of cognitive science concerning remembering and forgetting. Once we recognise that the duration for which beliefs are retained matters with respect to their value, and that our cognitive makeup places constraints on belief retention, the Numerosity Problem and the Epistemic Junk Problem dissolve. The resulting consequentialist view—Time-Sensitive Consequentialism—is not objectionably demanding.

Compliance with Ethical Standards

The author has no conflicts of interests to declare.

References

- Ahlstrom-Vij, K. and Grimm, S. (2013). Getting it right. *Philosophical Studies*, 166(2). 329-347.
- Anderson, M. C., Bjork, R. A., and Bjork, E. L. (1994). Remembering can cause forgetting: retrieval dynamics in long-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25.
- Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York, NY: Oxford University Press.
- Alston, W. (2005). *Beyond Justification: Dimensions of Epistemic Evaluation*. Ithaca, NY: Cornell University Press.
- Beardsley, M. C., (1958). *Aesthetics*. Indianapolis, IN: Hackett.
- Berker, S. (2013a). Epistemic Teleology and the Separateness of Propositions. *Philosophical Review*, 122(3), 337–393.

- Berker, S. (2013b). The Rejection of Epistemic Consequentialism. *Philosophical Issues*, 23, 361-87.
- Bjork, E. L., & Bjork, R. A. (1988). On the adaptive nature of memory. In Gruenberg, Morris, & Sykes (Eds.), *Practical aspects of memory: Current research and issues*. New York: Wiley.
- Bjork, R. A., & Bjork, E. L. (1992). A new theory of disuse and an old theory of stimulus fluctuation. In A. Healy, S. Kosslyn, & R. Shiffrin (Eds.), *From Learning Processes to Cognitive Processes: Essays in Honor of William K. Estes* (pp. 35–67). Hillsdale, NJ: Erlbaum.
- Bjork, R. A. (2011). On the symbiosis of remembering, forgetting, and learning. In A. S. Benjamin (Ed.), *Successful remembering and successful forgetting: A festschrift in honor of Robert A. Bjork*. Psychology Press.
- Brown, R., & Kulik, J. (1977). Flashbulb memories. *Cognition*, 5, 73–99.
- Budd, M. (1995) *Values of Art: Painting, Poetry, and Music*. London: Penguin.
- Burge, T. (2010). *Origins of Objectivity*. Oxford University Press.
- Byrne, A. (2010). Recollection, perception, imagination. *Philosophical Studies*, 148(1):15–26.
- Clark, A. and Chalmers, D. (1998). The Extended Mind. *Analysis*, 58(1), 7-19.
- Chun, M. & Turk-Browne, N. (2007). Interactions between attention and memory. *Current Opinion in Neurobiology*, 17, 177-184.
- Christensen, D. (1994). Conservatism in Epistemology. *Nous*, 28, 69–89.
- Cohen, N. J. and Squire, L. R. (1980). Preserved learning and retention of pattern-analyzing skill in amnesia: dissociation of knowing how and knowing that. *Science*, 210(4466).
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24(1), 87-114.
- David, M. (2001). Truth as the epistemic goal. In M. Steup (Ed.), *Knowledge, truth, and duty* (pp. 151–169). New York: Oxford University Press.
- DePaul, M. R. (2004). Truth Consequentialism: Withholding and Proportioning Belief to the Evidence. *Philosophical Issues*, 14(1), 91–112.
- Dickie, G. (1988). *Evaluating Art*. Philadelphia, PA: Temple University Press.
- Douven, I. (2010). The Pragmatics of Belief. *Journal of Pragmatics*, 42(1), 35–47.
- Dunn, J., & Ahlstrom-Vij, K. (2017). Is Reliabilism a Form of Consequentialism? *American Philosophical Quarterly*, 54(2), 183-194.
- Dunn, J., & Ahlstrom-Vij, K. (2018). Introduction: Epistemic Consequentialism. In J. Dunn & K. Ahlstrom-Vij (Eds.), *Epistemic Consequentialism*. New York: Oxford University Press.
- Feldman, R. (2000). The Ethics of Belief. *Philosophy and Phenomenological Research*, 60, 667–695.
- Fernandez, J. (2006). The intentionality of memory. *Australasian Journal of Philosophy*, 84(1), 39–57.
- Field, H. (2009). What is the Normative Role of Logic? *Aristotelian Society Supplementary Volume*, 83(1), 251–268.

- Firth, R. (1981). Epistemic Merit, Intrinsic and Instrumental. *Proceedings and Addresses of the American Philosophical Association*, 55, 5-23.
- Fodor, J. (1968). *Psychological explanation*. Random House.
- Forsberg A, Guitard D, Cowan N. (2021). Working memory limits severely constrain long-term retention. *Psychonomic Bulletin & Review*. 28 (2): 537-547.
- Friedman, J. (2018). Junk Beliefs and Interest-Driven Epistemology. *Philosophy and Phenomenological Research*, 97(3), 568–583.
- Friedman, J. (2019). Teleological epistemology. *Philosophical Studies*, 176(3), 673–691.
- Frise, M. (2015). Epistemology of Memory. *Internet Encyclopedia of Philosophy*.
- Frise, M. (2018). Forgetting. In K. Michaelian, D. Debus, & D. Perrin (Eds.), *New Directions in the Philosophy of Memory* (pp. 233–240). Routledge. 1-25.
- Goldman, A. (1979). What is Justified Belief? In Pappas, G. (ed.) *Justification and Knowledge*. Dordrecht: D. Riedel.
- Goldman, A. (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Goldman, A. (1999). *Knowledge in a Social World*. New York: Oxford University Press.
- Goldman, A. (2002). *Pathways to Knowledge: Private and Public*. New York: Oxford University Press.
- Goldman, A. (1995). *Aesthetic Value*. Boulder, CO.: Westview Press.
- Grimm, S. (2009). Epistemic normativity. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Epistemic value*. Oxford: Oxford University Press.
- Harman, G. (1986). *Change in View*. Cambridge, MA: MIT Press.
- Horwich, P. (2006). The Value of Truth. *Nous*, 40, 347–60.
- Hu, X. (2016). Why Do True Beliefs Differ In Epistemic Value? *Ratio*, 30, 255–269.
- Hunt, R. (2003). Two contributions of distinctive processing to accurate memory. *Journal of Memory and Language*, 48, 811–825.
- Kelly, T. (2003). Epistemic Rationality as Instrumental Rationality: A Critique. *Philosophy & Phenomenological Research*, 66(3), 612-640.
- Kensinger, E. A., & Schacter, D. L. (2008). Memory and Emotion. In M. Lewis, J. M. Haviland-Jones, & L. F. Barrett (Eds.), *The Handbook of Emotion* (3rd ed.) New York, NY: Guilford Press.
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Luminet, O., & Curci, A. (2009). Introduction. In O. Luminet & A. Curci (Eds.), *Flashbulb memories: New issues and new perspectives* (pp. 1-10). New York, NY: Psychology Press.
- Lynch, M. (2004). *True to Life: Why Truth Matters*. Cambridge, MA: MIT Press.
- MacFarlane, J. (2004). In What Sense (If Any) Is Logic Normative for Thought? *Manuscript*.
- McGeoch, J. A. (1932). Forgetting and the law of disuse. *Psychological Review*, 39(4), 352–370.

- Michaelian, K. (2011). The Epistemology of Forgetting. *Erkenntnis*, 74, 399–424.
- Michaelian, K. & Sutton, J. (2017). Memory. *The Stanford Encyclopedia of Philosophy*, Summer 2017 edition.
- Millikan, R. (1993). *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.
- Moon, A. (2013). Remembering Entails Knowing. *Synthese*, 190, 2717–2729.
- Murayama, K., Miyatsu, T., Buchli, D., & Storm, B. C. (2014). Forgetting as a consequence of retrieval: a meta-analytic review of retrieval-induced forgetting. *Psychological Bulletin*, 140(5), 1383–1409.
- Nairne, J. S., Thompson, S. R., & Pandeirada, J. N. S. (2007). Adaptive memory: Survival processing enhances retention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33(2), 263–273.
- Nairne, J. S., Pandeirada, J. N. S., & Thompson, S. R. (2008). Adaptive Memory: The Comparative Value of Survival Processing. *Psychological Science*, 19(2), 176–180.
- Nørby, S. (2015). Why Forget? On the Adaptive Value of Memory Loss. *Perspectives on Psychological Science*, 10(5), 551–578.
- Parker, E. L., Cahill, L., & McGaugh, J. L. (2006). A case of unusual autobiographical remembering. *Neurocase*, 12, 35–49.
- Ryle, G. (1949). *The Concept of Mind*. Barnes & Noble.
- Sainsbury, M. (2002). What Logic Should We Think With? In A. O’Hear (Ed.), *Royal Institute of Philosophy Lecture Series*. Cambridge: Cambridge University Press.
- Schwitzgebel, E. (2013). A Dispositional Approach to Attitudes: Thinking Outside the Belief Box. in N. Nottelman (ed.) *New Essays on Belief*. Palgrave Macmillan.
- Sibley, F. (1983). General Criteria and Reasons in Aesthetics. In J. Fisher (ed.) *Essays in Aesthetics: Perspectives on the Work of Monroe C. Beardsley* (pp. 3-20). Philadelphia, PA: Temple University Press.
- Singer, D. (2018). How to Be an Epistemic Consequentialist. *The Philosophical Quarterly*, 68(272), 580-602.
- Singer, D. (2019). Demoting promoting objections to epistemic consequentialism. *Philosophical Issues*, 29, 268-280.
- Scofield, J.E., Buchanan, E.M. & Kostic, B. (2018). A meta-analysis of the survival-processing advantage in memory. *Psychonomic Bulletin & Review*, 25, 997–1012.
- Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgery, and Psychiatry*, 20(1), 11–21.
- Sosa, E. (2001). The Place of Truth in Epistemology. In M. R. DePaul & L. Zagzebski (Eds.), *Intellectual Virtue: Perspectives from Ethics and Epistemology* (pp. 155–180). Oxford: Oxford University Press.
- Squire, L. (2009). Memory and brain systems: 1969-2009. *The Journal of Neuroscience*, 29(41), 12711-12716.

- Stalnaker, R. (1984). *Inquiry*. Cambridge, MA: MIT Press.
- Sylvan, K. (2020a). An Epistemic Nonconsequentialism. *The Philosophical Review*, 129(1), 1-51.
- Sylvan, K. (2020b). Reliabilism without Epistemic Consequentialism. *Philosophy and Phenomenological Research*, 100(3), 525-555.
- Tulving, E., & Pearlstone, Z. (1966). Availability versus accessibility of information in memory of words. *Journal of Verbal Learning & Verbal Behavior*, 5(4), 381-391.
- Wallace, R. J. (2001). Normativity, Commitment, and Instrumental Reason. *Philosopher's Imprint*, 1(3), 1-26.
- Wedgwood, R. (2012). Justified inference. *Synthese*, 189(2), 273-295.
- White, R. (2005). Epistemic Permissiveness. *Philosophical Perspectives*, 19(1), 445-459.
- Williamson, T. (1998). Conditionalizing on Knowledge. *British Journal for the Philosophy of Science*, 49, 89-121.
- Zangwill, N. (1984). Dickie and Doughnuts. *Ratio*, 7, 63-80.