Reliabilism: Holistic or Simple?

Jeff Dunn

`jeffreydunn@depauw.edu`

1 Introduction

A standard statement of Reliabilism about justification goes something like this:

Simple (Process) Reliabilism: S's believing that p is justified if and only if that belief was produced by a belief-forming process that is reliable above some specified high threshold.[1]

The thought behind this is that there are many belief-forming processes, some reliable and some not; those that are reliable yield *justified* beliefs.

Though this is how Reliabilism is often formulated today, and is similar to how Alvin Goldman ([1979]) initially formulates the view, a more sophisticated rule-based version of the theory is offered by Goldman in *Epistemology and Cognition* ([1986]). This version of the theory is widely cited and Goldman has made clear that he thinks that the best version of Reliabilism is the more sophisticated version offered in the later work.[2]

The more sophisticated version makes use of a rule framework, and has us focus on the total reliability of an epistemic system that follows those rules, rather than individual processes. Goldman endorses:

---

[1] This view first appears in Alvin Goldman's ([1979]) 'What Is Justified Belief?'. What I've stated in the text is a slight simplification of the view defended there.

[2] For a sampling of those who cite Goldman's ([1986]) view, and explicitly note that it is more sophisticated than Simple Reliabilism, see Casullo ([1988], p. 208), Lammenranta ([1996], p. 118), Chase ([2004], pp. 123-4), Cruz & Pollock ([2004]), and Adler ([2005], p. 446). For Goldman's own views on this, the best source is Goldman ([1988], p. 58). See also Goldman ([2010], [2009a], [2009b], [2008], [1999]), and Goldman & Olsson ([2009]).

ARI: A J-rule system[3] $\mathcal{R}$ is right if and only if $\mathcal{R}$ permits certain (basic) psychological processes, and the instantiation of these processes would result in a truth ratio of beliefs that meets some specified high threshold. ([1986], p. 106)

Notice that the truth ratio in question is for an entire epistemic system, not for a particular process. Thus, ARI doesn't tell us when a *particular* belief is justified. This is taken care of by Goldman's commitment to:

P1: S's believing p at time t is justified if and only if S's believing p at t is permitted by a right system of J-rules. ([1986], p. 59)

ARI and P1 give us something more complex than Simple Reliabilism. Putting them together in the most natural way yields:

Holistic (Process) Reliabilism (HR): S's believing that p is justified if and only if S's believing that p is the result of psychological processes, the instantiation of which lead to a high truth-ratio.[4]

Why prefer this more complex version of Reliabilism? Two reasons suggest themselves. First, the interaction of various processes is often important. For example, the human process of recalling propositions from memory may only be reliable if it is working together with some sort of process that executes what Gilbert Harman ([1986]) has called *Clutter Avoidance*. This process would do something like purge the memory of old propositions not likely to be needed.[5] Although the process executing Clutter Avoidance is not truth-conducive on its own, when working with normal memory recall, it results in a high truth-ratio. Reliabilism, the thought goes, should be sensitive to this.[6] Second, Goldman has expressed doubts over whether individual processes will often have determinate truth-ratios.[7] If they don't, then Simple Reliabilism is inadequate.

---

[3]A J-rule system is a system of rules about justification. Intuitively, one can think of a J-rule system as a recipe that indicates the permitted psychological processes when building an epistemic agent.

[4]I drop reference to time for ease of presentation.

[5]Kirk Michaelian ([2010]) has recently argued convincingly that to the extent that humans execute something like *Clutter Avoidance*, it works very different than Harman's discussion suggests. However, this fact will not be critical to this discussion.

[6]See Goldman ([1986], p. 115).

[7]See Goldman ([1988], p. 63).

In this paper I will consider two main ways of understanding Goldman's more complex version of Reliabilism. I will argue that neither is successful. I will conclude by arguing that Simple Reliabilism is actually fine as it is.

2 Holistic Reliabilism

Holistic Reliabilism as stated above seems to fit Goldman's ARI and P1 well. A belief is justified if it is permitted by the right rules, and the right rules permit certain psychological processes: those, which when instantiated, result in a high truth-ratio. As stated, however, Holistic Reliabilism is ambiguous. It is ambiguous, because it is not clear what it means for a belief to be the result of the permitted psychological processes—that is, the processes the instantiation of which lead to a high truth-ratio.

Consider, for example, one way of precisifying Holistic Reliabilism:

HR1: S's believing that p is justified if and only if S's believing that p is the result of **all** the permitted psychological processes.

We cannot understand Holistic Reliabilism in this way. The full range of permitted psychological processes must at least include processes corresponding to the different sense modalities, and probably more besides. Very few beliefs—perhaps none—are the result of *all* the permitted psychological processes. HR1 would say that only such beliefs—if there are any—are justified. So HR1 can't be right.

Consider, then, a natural alternative to HR1:

HR2: S's believing that p is justified if and only if S's believing that p is the result of **some subset** of all the permitted psychological processes.

But this cannot be correct either. The problem is illustrated by the following example. Suppose that a certain set of permitted processes yields an overall high truth-ratio. Suppose the set includes a process that executes something like Clutter Avoidance and also a process that executes a kind of memory recall. Since the process of memory recall is a member of the set of permitted processes, HR2 says that beliefs produced by such a process are justified. But suppose that this process of memory recall yields a high truth-ratio only when a process like Clutter Avoidance is doing its work. Now, consider an agent who uses that particular memory recall process, but doesn't have a process of Clutter Avoidance. In this case, the memory recall process is not very

reliable at all. But HR2 says that the belief produced by such a process is justified. That seems wrong. A process that is recommended *only* because it works together with other processes to produce reliable results should not count as justification conferring when it operates (unreliably) on its own, apart from those other processes.

One response to this objection is to adopt a very fine-grained way of typing processes. One could claim, for instance, that there are really two types of processes of memory recall being described in the example: one that operates in a psychological system that has a Clutter Avoidance process, and one that operates in a psychological system that has no such process. This does avoid the objection. Using memory recall without Clutter Avoidance is not, then, to use a process that is permitted. However, it leads to further problems.

On one way of understanding this response, the type of memory recall process we are dealing with depends on whether or not that process is located within a system that also has a process of Clutter Avoidance, but the *other* processes within the system are irrelevant to typing the memory recall process. This way of spelling out the response runs into trouble, however, for on this way of going we lose the holistic nature of the proposal. To see this, suppose that we have what we would initially describe as two process types: $P1$ and $P2$. According to this response, what determines whether $P1$'s presence or absence makes a difference to how $P2$ is typed? The natural thought here is that $P1$ is relevant to typing $P2$ just in case $P1$'s presence or absence influences the reliability of $P2$. But if we go this way, then we are admitting that there are subsets of processes within the total set of permitted processes that are independent of each other. If that's right, then holistic assessment is not needed. We don't *need* to evaluate all the processes an epistemic system has together.[8]

A more holistic way of understanding the response to the objection to HR2 is to say that *all* the processes within a system are relevant to typing any process in that system. That is, we say that two processes are of the same type only if they are each contained in psychological systems that themselves contain all and only the same processes. This ensures that memory recall is not permitted if it is not located within a system that has a process of Clutter Avoidance, and it does so in a way that preserves the holistic aspect

---

[8]This is essentially the same issue that arises with respect to the fourth modification of Holistic Reliabilism, HR4, which is discussed below.

of the proposal.

Unfortunately, this leads to intolerable consequences. No matter what form of reliabilism we adopt, we surely want to be able to say that some of an agent's beliefs are justified and some are not. If, however, we type processes in the fine-grained way suggested above, then we will be prohibited from saying this. To see why, consider a particular example. Suppose we want to say of some agent that his belief that q is not justified. According to HR2, this must be because the belief that q is the result of some process P, that is not a member of a set of processes the instantiation of which leads to a high truth-ratio.[9] Notice that on the view being considered, it is impossible for some of the agent's other processes to be instantiated without P. This is because, on this view, part of what it is to be those other processes is to be in a system together with P. This means that for any process the agent has, it can only be instantiated with all and only the other processes the agent has. Since P is not in the special set, it follows that none of the agent's processes are in the special set. Thus, we get the intolerable consequence that if one of the agent's beliefs is unjustified, they all are.

There is another way of responding to the objection to HR2, however. This response does not require us to type processes in the extremely fine-grained way. Rather, the response is to modify Holistic Reliabilism as follows:

HR3: S's believing that p is justified if and only if
(i) S's believing that p is the result of **some subset** of all the permitted psychological processes, and
(ii) S has **all** the permitted psychological processes.

Notice that this, too, gets around the objection to HR2. The problem with HR2, recall, was that the agent's belief was formed by the permitted memory recall process, but without having the process of Clutter Avoidance. Clause (ii) of HR3 prevents this case from occurring. However, HR3 runs into a clear problem. According to HR3, if you lack one of the permitted processes, then you cannot have justified beliefs. But this is surely wrong. Surely I could lack some permitted psychological process and have justified beliefs. For instance, I retain the capacity to have justified beliefs even if I have gone deaf.

As a final attempt to salvage Holistic Reliabilism, one might offer:

---

[9]Strictly speaking, the belief that q could be the result of several processes, but the point is the same either way.

HR4: S's believing that p is justified if and only if S's believing that p is the result of **some special subset of all** the permitted psychological processes.

The idea here is that what matters is not the full set of processes that the agent has, but rather a special set of processes that produced the belief. On the most plausible way of spelling this out, the special set contains mutually *interdependent* processes that work together for reliable results, and yet are *independent* of the rest of the psychological processes and mechanisms at work in the agent.

This version of Holistic Reliabilism appears to avoid the problems with its predecessors.[10] However, it has done so at the cost of collapsing into something indistinguishable from Simple Reliabilism. The reason to move away from Simple Reliabilism was that individual processes (or individual sets of processes) either lack reliability profiles or have reliability profiles that are misleading. The remedy to this problem, according to Holistic Reliabilism, is a more holistic assessment of reliability. HR4, however, suggests that such a holistic view is not needed: one can focus on clusters of interdependent processes, while ignoring what else is going on in the agent. So, if there is a reason move away from Simple Reliabilism, then there is similarly a reason to move away from HR4.

3   A Different Interpretation

Holistic Reliabilism as formulated in the previous section is untenable. However, there is a different way of interpreting Goldman's claims that might seem to do better. This comes from focusing more attention on the rule framework that Goldman proposes. Goldman writes:

> The point is that rules are *interdependent* with respect to their epistemically relevant properties. In particular, they are interdependent with respect to truth-ratio properties. This is especially clear for inferential rules. [. . .] A sound inferential rule will generate additional true beliefs when applied to true input beliefs. But if other rules permit false beliefs to be formed, then even a sound inferential rule may produce innumerable errors.

---

[10]Indeed, it is similar to the first response made to the objection to HR2.

> Truth-ratio propensities, then, only make sense as applied to rule systems, not isolated rules. (Goldman [1986], p. 115)

In this passage, attention is drawn not so much to the interdependence of *processes*, but instead the interdependence of *rules*. The idea seems to be the following:

Holistic Rule Reliabilism: S's believing p is justified if and only if in believing p, S obeys all the J-rules, which when perfectly obeyed by some epistemic system results in a high truth-ratio.

Now, if the right J-rules simply say of a set of processes that they are permitted, then this view will reduce to some version of Holistic Reliabilism discussed in the previous section.[11] It would simply say that S's believing p is justified if and only if in believing p, S uses the permitted processes.

However, if we allow that the J-rules state not only what the permitted processes are, but also which processes must be combined with other processes, then we have a different view. For example, perhaps one J-rule says that the process of memory recall is permitted *only when* used in conjunction with some process of Clutter Avoidance. Consider, then, the agent that uses the process of memory recall, but without using the process of Clutter Avoidance. Although the process of memory recall is one of the processes that is instantiated in a system that obeys all the J-rules, *simply* using the process of memory recall does not obey *all* the J-rules. So, such a belief would be unjustified according to Holistic Rule Reliabilism. Thus, Holistic Rule Reliabilism avoids the problems that plagued Holistic Reliabilism.

Further, Holistic Rule Reliabilism appears to hold on to the holistic features that allegedly motivate a move away from Simple Reliabilism. We do not consider processes one by one on this view. Rather, we consider the truth-ratio that results from a total system of J-rules being followed perfectly by some epistemic system. Since the total set of J-rules can include rules that encode the interdependence between rules, we have an aspect of holism.

Despite this success, Holistic Rule Reliabilism runs into problems elsewhere. Consider some set of J-rules, $\{\mathcal{J}\}$, that when perfectly followed

---

[11]It's worth noting that this appears to be just what Goldman has in mind (see, Goldman [1988]). However, since such a view faces the problems detailed in the previous section, it is worth exploring alternatives.

results in a truth-ratio well above the threshold specified. Now, consider adding to this set the following rule:

Rogue Rule: The process *Temp* may occasionally be used.

Temp is the process of believing that the temperature is $-T$ when one's thermometer reads 'T'. The process *Temp* is very unreliable. However, if it is used rarely enough (as the rule states), it won't do enough damage to the overall truth-ratio of a system that follows the set of rules $\{\mathcal{J} \cup \text{Rogue Rule}\}$. In other words, a system that perfectly follows this new set of rules will still yield a truth-ratio over the threshold. But then Holistic Rule Reliabilism says that a belief formed by the process *Temp* is justified. It is clear, however, that such a belief is not justified. The fact that following the set of rules $\{\mathcal{J} \cup \text{Rogue Rule}\}$ yields a high truth-ratio does not show that every process approved by those rules is one that adds to the reliability of the belief-forming system. The very holistic nature of Holistic Rule Reliabilism guarantees that this is possible.

In a short section of his ([1988]), Goldman considers this problem. He writes: "It is not enough that a rule system itself have a sufficiently high truth-ratio [...]. It must also be required that none of its *subsets* should have an *in*sufficiently high truth-ratio." (p. 64). Here is the procedure that Goldman is recommending. First, we are to separate a set of potential J-rules into its atomic rules. Note that this first step might be considerably difficult in that one will have to come up with a unique way of counting rules. But suppose this can be done. For simplicity, suppose we are considering a set of J-rules that has three atomic rules: J1, J2, and J3. According to Goldman, we look at all the epistemic systems that follow any subset of these three rules, and record the truth-ratios of each of these systems. If any system that follows any subset of these rules has too low a truth-ratio, the J-rule system is not a right rule system.

Unfortunately, this solution fails. The main reason it fails can be most easily seen if we think about a particular set of J-rules. Consider, for example, the set of J-rules that we think of as an intuitively correct set. Some of these rules will say which processes are permitted; others will say which processes are prohibited; still others will say which processes should be used with others. Suppose one such J-rule in the correct set states: "The following process is permitted: forming beliefs about the presence of physical objects based on normal visual perception in adequate light." Surely something like this rule (with necessary modifications) will be in the correct set of J-rules.

But notice that the set that contains just this J-rule is a subset of the set of correct J-rules. On Goldman's proposal, the truth-ratios for all this rule, then, must be above a given threshold if the correct set really is correct.

Imagine, however, all the epistemic systems that follow this one rule. Some of these epistemic systems follow this one J-rule and in addition many of the other correct J-rules. Such epistemic systems will have high overall truth-ratios. However, there are also epistemic systems that follow this one J-rule, and yet in every other way, use processes that are woefully poor in terms of reliability. Perhaps one such system forms beliefs about physical objects as a result of normal perception in adequate light, but also forms beliefs on the basis of counter-inductive principles, believes not-p in response to a memory that p, and so on. The truth-ratio for this epistemic system will be very low, despite the fact that it follows a subset of the correct set of J-rules. The general point here is that the truth-ratio will be very low for most epistemic systems that follow only a small subset of the correct J-rule system, because a small subset specifies so little about the behavior of an epistemic system. So, it seems, no J-rule system will satisfy the extra condition Goldman proposes.

One might attempt to avoid this result by arguing that the atomic J-rules govern not single processes, but rather clusters of processes. If these clusters are big enough, they may specify enough about the behavior of the epistemic systems following the rules in that cluster so as to guarantee that the truth-ratio for such systems is suitably high. But this proposal faces a dilemma: if the clusters include enough rules, then we can always sneak in something like Rogue Rule; if they do not include enough rules, the original problem arises again.

This suggests that perhaps I am misinterpreting Goldman. Perhaps he only means to claim that each subset of the rules must have a high truth-ratio *with respect to the beliefs produced by the process governed by the rule in question.* This would avoid the problem. However, such a view is indistinguishable from Simple Reliabilism. What we would be doing is evaluating individual processes in terms of their truth-ratios independent of the other processes at work in an epistemic system. The holistic aspect of the proposal would be lost once again.

An alternative solution, still in the spirit of Goldman's solution, is to try to prohibit rules like the Rogue Rule as follows:

A right set of J-rules $\{\mathcal{J}\}$ is such that:

1. when perfectly followed yield a truth-ratio n, greater than some specified high threshold, and

2. include no rule r such that the set $\{\mathcal{J} - r\}$ when perfectly followed, yield a truth-ratio greater than n.

This doesn't require that *every* subset of a rule system yield a truth-ratio over a certain threshold, nor that *every* subset of a rule system yield a truth-ratio at all. It only requires that certain still-quite-large subsets of rule systems (all the subsets of $\{\mathcal{J}\}$ containing one less member than $\{\mathcal{J}\}$) have truth-ratios over a certain threshold. Thus, it retains *some* amount of holistic evaluation.

I think this view about J-rules, paired with Holistic Rule Reliabilism, is the most defensible version of Goldman's ideas about reliability assessment. Despite this, it still seems to fail to offer a real alternative to Simple Reliabilism. After all, to adopt this view is to admit that a holistic assessment of an entire epistemic system is not strictly what we are interested in. According to this view, the fact that following the set of rules $\{\mathcal{J}\}$ yields a high truth-ratio does *not* show that a belief formed by a process that violates none of the rules in $\{\mathcal{J}\}$ is justified. Instead, the reliability of processes, *on their own*, is held to be relevant to the justification of a belief produced by that process. In fact, if this were not so, the view would not be able to say that a belief formed by the Rogue Rule is unjustified.

Goldman's insight is that sometimes several sub-processes can combine in interesting ways to yield a process that has reliability properties not had by each of the sub-processes. However, this does not motivate moving away from Goldman's ([1979]) initial formulation of Reliabilism to a totally holistic assessment of reliability. We can capture Goldman's important insight with Simple Reliabilism, and just be clear that the processes in question can sometimes be composed of sub-processes that are not necessarily reliable in isolation. This does point attention to the Generality Problem[12]—the problem of how exactly to specify the relevant processes—but this is a problem that Reliabilism has long been known to have.

References

[1] Adler, J. [2005]: 'Reliabilist Justification (or Knowledge) as a Good Truth-Ratio', *Pacific Philosophical Quarterly*, **86**, pp. 445458.

_____

[12]See Conee & Feldman [1998].

[2] Casullo, A. [1988]: 'Revisability, Reliabilism, and A Priori Knowledge', *Philosophy and Phenomenological Research*, **49**, pp. 187-213.

[3] Chase, J. [2004]: 'Indicator Reliabilism', *Philosophy and Phenomenological Research*, **69**, pp. 115-137.

[4] Conee, E. & Feldman, R. [1998]: 'The Generality Problem for Reliabilism', *Philosophical Studies*, **89**, pp. 1-29.

[5] Cruz, J. & Pollock, J. [2004]: 'The Chimerical Appeal of Epistemic Externalism', in R. Schantz (ed.), *The Externalist Challenge*, Berlin: Walter de Gruyter, pp. 125-142.

[6] Goldman, A. [1979]: 'What Is Justified Belief?', in George Pappas (ed.), *Justification and Knowledge*, Dordrecht: D. Reidel, pp. 1-23.

[7] Goldman, A. [1986]: *Epistemology and Cognition*, Cambridge, MA: Harvard University Press.

[8] Goldman, A. [1988]: 'Strong and Weak Justification', *Philosophical Perspectives, Vol. 2, Epistemology*, pp. 51-69.

[9] Goldman, A. [2008]: 'Immediate Justification and Process Reliabilism', in Quentin Smith (ed.), *Epistemology: New Essays*, New York: Oxford University Press, pp. 63-82.

[10] Goldman, A. [1999]: 'Internalism Exposed', *The Journal of Philosophy*, **96** pp. 271-93.

[11] Goldman, A. [2009a]: 'Internalism, Externalism, and the Architecture of Justification', *Journal of Philosophy*, **106**, pp.

[12] Goldman, A. [2009b]: 'Reliabilism', *The Stanford Encyclopedia of Philosophy (Fall 2009 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2009/entries/reliabilism/>.

[13] Goldman, A. [2010]: 'Epistemic Relativism and Reasonable Disagreement', in R.Feldman and T. Warfield (eds.), *Disagreement*, New York: Oxford University Press.

[14] Goldman, A. & Olsson, E. [2009]: 'Reliabilism and the Value of Knowledge', in D. Pritchard et al. (eds.), *Epistemic Value*, Oxford: Oxford University Press, pp. 19-41.

[15] Harman, G. [1986]: *Change in View*, Cambridge, MA: MIT Press.

[16] Lammenranta, M. [1996]: 'Reliabilism and Circularity', *Philosophy and Phenomenological Research*, **56**, pp. 111-124.

[17] Michaelian, K. [2010]: 'The Epistemology of Forgetting', *Erkenntnis*, **74**, pp. 399-424.