King's Research Portal

Document Version
Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Roush, S. (2017). The Difference between Knowledge and Understanding. In R. Borges, C. de Almeida, & P. D. Klein (Eds.), *Explaining Knowledge: New Essays on the Gettier Problem* Oxford University Press; Oxford.
<https://global.oup.com/academic/product/explaining-knowledge-9780198724551?cc=gb&lang=en#>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

The Difference between Knowledge and Understanding - final

Sherrilyn Roush

King's College London

Short Abstract:

Characterizing Gettier cases as failures of probabilistic *relevance matching*, and knowledge as tracking, we can explain the nature and value of knowledge simultaneously. Avoidance of gettierization has the independent value of improving understanding of why p is true, and relevance matching gives a theory of understanding as distinct from knowledge.

Medium Abstract:

I characterize Gettier cases as failures of understanding, and give a theory of what it is to understand why proposition p is true. This view is based on the concept of probabilistic relevance matching, having one's dispositions to believe p mirror the probabilistic relations that p has to all other matters. Based in probability, the view yields a clear relationship, and also distinction, between the concept of understanding and the concept of knowledge defined in terms of probabilistic tracking. With these tools we are able to see that gettierization avoidance has a value independent of the value of knowledge, viz. understanding, but that it is also in the nature of tracking-type knowledge to discourage gettierization quite specifically. The concept of understanding here captures several key features of this phenomenon, such as breadth, depth, and appreciation of connections, and allows us to see understanding as simulation.

Long Abstract:

In the aftermath of Gettier's examples, knowledge came to be thought of as what you would have if in addition to a true belief and your favorite epistemic goody, such as justifiedness, you also were ungettiered, and the theory of knowledge was frequently equated, especially by its detractors, with the project of pinning down that extra bit. It would follow that knowledge contributes something distinctive that makes it indispensable in our pantheon

of epistemic concepts only if avoiding gettierization has a value that can be explained without presupposing the value of knowledge.

Tracking-type knowledge has a value that no other logically possible conditions on true belief does. As an Evolutionarily Stable Strategy it preserves appropriate belief states through time and changing circumstances. If we characterize gettierization through the concept of relevance matching, then we see that avoiding gettierization has a value independent of that of knowledge, namely, understanding, and that it is unnecessary to add a clause to the tracking conditions to make them suppress gettierization directly, though fallibly. The bright line of value is between gettierization avoidance and understanding on the one hand and knowledge on the other, and so should be the bright line defining concepts.

The concept of relevance matching is key to a definition of what it is to understand *why* p is true, as opposed merely to knowing *that* p is true. Perfect tracking implies perfect relevance matching, so knowledge and understanding are intimately connected but understanding also requires that one own states that accomplish the relevance matching rather than achieving it vicariously. The theory of understanding based on relevance matching implies that understanding requires appreciation of not only p but its connections to other matters, and explains how it is possible to know that p is true without understanding why. The view implies that understanding is literally simulation, and is suggestive about understanding other minds.

1. What is the Gettier Problem?

In formulating the problem that bears his name Edmund Gettier stepped back from traditional thinking about knowledge and named something he thought all of the views had in common despite the differences among them that previous discussions had focused on.¹ They all defined knowledge as true belief plus justifiedness, however they went on to describe the latter property. (Gettier 1962) Thus in providing examples where the subjects are intuitively justified but something was missing that it seemed most natural to call “knowledge”, the Gettier Problem opened up space for thinking about what knowledge is that was not restricted to what we would now call internalist justifiedness of belief. Since this problem was identified many alternative conceptions of justification have been defined and continue to flourish, internalist concepts of justification have been refined and better understood, and in the process philosophers

have made explicit dimensions of knowledge itself that were not discussed until this period, for example, defeasibility (Lehrer and Paxson 1969), reliability (Goldman 1967, 1976; Dretske 1971; Nozick 1981), and virtue (Sosa 1991, Zagzebski 1996). The Gettier cases also presented a challenge for fallibilism as such (Zagzebski 1994), and have forced us to think about the nature and role of luck in a term of appraisal of epistemic achievement (Pritchard 2005). For all that one might deplore a scholastic style of conceptual analysis that the problem gave rise to, the provocation it gave also had positive, enlightening consequences, both substantively and methodologically.

Even granting this is valuable work, though, one could regard it as a distant side effect of a project that is itself pointless: finding a gettier-proof definition of knowledge. The most interesting arguments in this direction are not that the project is doomed to failure, which seems obvious, but that it would have no value if it succeeded. One argument to this effect, made by Mark Kaplan in the 1980s, takes it as a requirement of epistemological theorizing that it advance or clarify the conduct of inquiry. (Kaplan 1985) The search for an anti-gettierization clause to close the gap between justified true belief and knowledge does not do this, he argues, because in inquiry once we check that we have a justification, broadly speaking², for our belief there is nothing more that we can check to make sure that our justified belief is in addition knowledge, in the sense of gettier-proof. Thus, whatever the successful anti-gettierization clause gave us would be of no use. This criticism is not as satisfying as one might like, I think, since it presumes we know ahead of time that whatever property the anti-gettierization clause identifies is something we already know enough to take into account when we go about checking that we have good evidence. This is less open than I would like to be to the possibilities that we do not fully understand the gettierization phenomenon, and that our current practices are incomplete.

An argument that takes the line a step further, that whatever the anti-gettierization clause identifies is not of epistemic value, comes more recently from Jon Kvanvig. (Kvanvig 2003) He takes it that if theorizing does not show its worth in practice or in any other way, it can and should offer explanations of its subject matter. In particular, he expects a theory of the nature of knowledge to explain why knowledge is valuable, more valuable than mere true belief. However, he argues, no account of why knowledge is valuable can also serve as a theory of its nature. The going accounts of what it is about knowledge that makes it valuable may succeed, but they do not naturally address the

Gettier Problem, so a definition – a proper statement of the nature of knowledge – will require the addition of further clauses to rule out the problematic Gettier cases. The prospects are dim for finding clauses that succeed, but even if we succeeded they would always be ad hoc and gerrymandered patches on a sleek property already identified as making knowledge more valuable than mere true belief.

It is not only that such modifications make for an ugly definition, but that it is unclear how fulfillment of any of the further conditions proposed would add value to true belief plus subjective justification or intellectual virtuosity, or one's favorite alternative criterion beyond true belief. And if we said that their value is that they get us finally to knowledge, then we would be presupposing the value of knowledge in order to explain the value of one of its components, that very component that made us call the property knowledge rather than the more limited phenomenon of justified, or virtuous, or ... true belief. If the gettier-proof phenomenon has no added value over a gettier-prone property that straightforwardly adds value to mere true belief, then it is hard to see why the study of the gettier-proof phenomenon – that is, the defining of knowledge as that project had come to be pursued – has any value either. One might think an epistemologist should be embarrassed to take the position that knowledge is not valuable, but Kvanvig points out in his defense that clear value has been attributed to phenomena that are (or are taken by some to be) components of knowledge – subjective justification, intellectual virtuosity, etc. – and this does some justice to the everyday intuition that knowledge has value.

A theorist of knowledge thus has two challenges, one to say why the property X that she takes to be required beyond true belief adds value to mere true belief, the other – the secondary value problem – why the further property that provides gettier-proofing adds value to true belief plus X, and she must do the second without presupposing that knowledge is valuable. The suspicion that a gettier-proof definition of knowledge will always be gerrymandered is what makes these distinct challenges: if property X had done the job of gettier-proofing, the second problem would be solved by the solution to the first. In my favorite definition of knowledge the X property that goes beyond true belief is not subjective justification or intellectual virtuosity but tracking, so my first task in this paper will be to explain how tracking adds value to mere true belief.

Tracking also does a lot to gettier-proof our beliefs, but as with all fallibilist criteria Gettier cases will be possible. However, instead of launching a search for further

conditions to narrow that gap, I will investigate gettierization taken by itself, without concern for the theory of knowledge. This is not because knowledge lacks value, or is not to be understood through the value of its constituents, but because, as I will argue, avoidance of gettierization has a nature and value of its own that does not rest on that of knowledge. The value of knowledge defined as requiring tracking also does not rest on it. Following these lines we should conclude that the mistake has been to try to force the concepts of knowledge and of gettierization avoidance into the same box, by assuming that knowledge must take the form of true belief + X + an anti-gettierization clause.

If we want to carve our concepts along the bright lines in the field of values of the phenomena, we have another option than those that have been explored thus far. We can take knowledge to be defined by tracking, and take gettierization avoidance to define a different concept. In developing my account of what gettierization is, we will see that this other concept is most naturally taken to be *understanding*. Both concepts, of knowledge and of understanding, will be defined by sleek properties, and their value will be understood through the value of all of their components. Gettierization avoidance is not per se a part of the value of knowledge, so it will not be a separable part of the definition of knowledge, though the tracking property that defines knowledge will naturally, fallibly, target what is gettierization-specific.

The one bit of violence this account will do to widespread intuitions is to allow some Gettier cases to be counted as knowledge. However, this is superior to Kvanvig's approach in that he neither gives an account of the nature nor of the value of knowledge in terms of its constituents, nor offers anything to explain the intuition that something epistemically valuable is missing in Gettier cases. My carving of the concept of knowledge, and my account of gettierization as failure of a property that defines understanding, achieve both. Something is missing in Gettier cases, but it is not knowledge; it is understanding.

2. Knowledge is more valuable than mere true belief.

If knowledge is tracking, then it has a high value that mere true belief does not. To explain that value, I will give an outline of how it is derived, following an account given earlier. (Roush 2010) We start with a game that a subject plays with Nature.

Nature ↓ You →	$b(p)$	$-b(p)$
p	(0, 10)	(0, -20)
$\neg p$	(0, -7)	(0, 5)

In any given round of the game Nature has two options, to make proposition p true or to make p false. In a given round the subject's options are to believe p or not believe p . To have a game there must be payoffs associated with the four possible outcomes. Nature's payoffs are in the first position in each ordered pair, and those of the subject, say you, are in the second spot. Nature is indifferent to the consequences of the play so her payoffs are zero in every possible outcome. The particular numbers corresponding to your payoffs do not matter, but some of their relationships do. I will assume that we are concerned with p for which when p is true it is better for you to believe it than not to believe it, and when p is false it is better not to believe it than to believe it. These payoff assumptions are sufficient for derivation of the value of tracking-type knowledge that I describe here, and they are also necessary.

There are a great many p that satisfy these conditions, but what eventually comes out of this derivation may lead one to suspect that the result was rigged by these payoff assumptions. However, this is definitely not so, for several reasons. First, the challenge to explain the value of knowledge is that of explaining why it is more valuable than mere true belief. No one has so far expected us to explain why knowledge of p is valuable for p where true belief is not valuable, or where avoiding false belief has no value, probably because we would not normally expect knowledge to be valuable in such cases either. So, the payoff assumptions about which states are more valuable than which are not begging the explanatory question, and they are not a qualification on the answer to the value question.

Second, the payoff assumptions do not trivialize the result I will describe by being only trivially different from the result. This is because the payoffs concern belief states, and

the result will be about strategies. Strategies are dispositions or regularities in your responses to the plays of the other player. So, they are conditional, and modal, and general; they are rules. States are outcomes in a single round of the game. In a given round, you may land in a comparatively good state or a comparatively bad state. In our game the upper left and lower right outcomes are best for you among your possibilities, because they have the best payoffs. $10 > -20$ and $5 > -7$. The lower left and upper right possibilities are comparatively bad. These good and bad outcomes in a given round are not assumed to be any more than co-incidences of two things, p 's truth or falsity, and belief or non-belief in p , on a given occasion.

Strategies are different because a strategy is a general rule for responding to all the possible plays of the other player, so it is a set of conditionals, and it will have consequences for where you land how frequently in all or a large fraction of trials of the game. As for strategies in all game theory, so here, no assumption is made that they involve conscious or deliberate actions. A strategy need only be a regularity of response, here of belief response. Pancreatic cells have strategies in the intended sense because they respond to the presence of sugar by producing insulin and to the absence of sugar by not producing insulin. However, differently from some game theory, here I do not assume in the definition of the game that the player of belief states has knowledge of which play Nature has made. The game is not that you will be told whether p is true or false and decide whether to believe it. Knowledge will exist, when it does, at the level of strategies, which we will evaluate in terms of their game-theoretic properties. Thus, representation of the value question as a game does not beg the question of what counts as knowledge or what its value is.

An example of a strategy in the true belief game is the following:

When nature plays p , play $\neg b(p)$.

When nature plays $\neg p$, play $\neg b(p)$.

Probabilistically,

$P(\neg b(p)/p) > s$, for some high s

$P(\neg b(p)/\neg p) > t$, for some high t

This would be a good strategy for someone who wanted to avoid mistaken beliefs at all cost, perhaps Descartes in certain moods. However, it would not be good for a subject

with the payoff structure we assumed above, for whom when p is true it is better that he believe p than not believe p . This strategy would lead him to believe p on (roughly) only $1-s$ of the occasions on which it is true.

Another strategy is given by the following two conditions:

$$\begin{array}{ll} P(-b(p)/-p) > s \leq 1 & \text{variation (sensitivity)} \\ P(b(p)/p) > t \leq 1 & \text{adherence} \end{array}$$

These are the tracking conditions, which are dispositional, and the pair of them form an example of a strategy in the true belief game because they say what your belief state does in response to all of the possible plays of nature, p and not- p : You satisfy them if the probability you do not believe p given that p is false is high, and the probability you do believe p given that p is true is high. That is, in response to nature's play of not- p , you play no belief and in response to nature's play of p you play belief most of the time in most of the ways that p could be false or true ($s = t = \text{high}$). In other work I have developed these conditions as a definition of knowledge; here we see that they form a strategy in the true belief game.

To respond according to the tracking strategy is not a mere choice, and sometimes not a choice at all. It is possible to exhibit the tracking strategy without a decision or deliberate effort, and usually a mere decision to have it does nothing to help you fulfill it. An example of the first occurs in perception, where we get the equipment to fulfill this strategy for free from having normal organs. We do not choose to believe there is a tiger if and only if there is one. There are situations – maybe this is one of them – where, as with pancreatic cells and insulin secretion, we could not choose to believe otherwise if we tried.

For an example of the second, where a mere choice does not get us the status of having the tracking strategy, take any proposition investigated by scientists. For a dramatic case, consider the existence of the Higgs boson. If human beings are to have the tracking dispositions for this p we have to acquire them, which in this case requires building something like a Large Hadron Collider. The harder a thing is to know, the more work will be required to acquire the tracking dispositions. Our task in getting ourselves to fulfill the tracking conditions when these dispositions did not, as in perception, come for free, is to get ourselves dispositions to detect which play Nature has made.

There are an infinite number of logically possible response strategies in the true belief game other than the tracking dispositions, and they correspond to all possible general conditions on the relation between p and your belief in p , and so all possible definitions of knowledge of the form true belief + X that do not have tracking among their requirements. As long as the condition in a theory of knowledge is general it can be written as a strategy. The X of a given theory of knowledge (e.g., good reason to believe, reliable process of belief formation) is extensionally equivalent to some strategy because the X provides an answer for every situation where p is true and p is false as to whether the subject who fulfills X believes p . Does the subject who is virtuous believe p in this situation when p is true? Does the reliable subject believe p when it is false in that situation? For cases where the X property does not determine whether the subject believes p , it will either give a probability of that or saying nothing at all. If the X property says nothing to determine what the subject's belief state is, that is extensionally equivalent to specifying a random response.³

We can now compare how all of these strategies serve us in the true belief game. With thresholds s and t set at 95%, the tracking strategy leads us to land in the winning upper left and lower right squares almost all of the time. I.e., in 95% of the rounds of the game we win. In order to judge other strategies in comparison to this, the true belief game must be converted into a symmetric signaling game with all strategies compared to each other as each plays to win the game with Nature. Suppressing the details, if we call the tracking strategy R then all the other possible strategies (i.e., other theories of knowledge) are members of $-R$, and we get an interesting implication: the subject who is a tracker of p has what is called an *Evolutionarily Stable Strategy* (ESS).

What this means is that the tracking type strictly dominates any type following *any* other conditions beyond true belief, in the struggle for survival and utility. Any given round can be lost by any individual, but this type wins over time. This is significant for the triviality worry one might be having again: far from being trivial it is not just that tracking has an added value over mere true belief, but that its added value is unique: no other logically possible added condition on knowledge – for example, justifiedness, virtue, reliable process, etc. – will swamp or even match the value tracking adds to mere true belief. The ESS property is very strong game-theoretically because its consequences hold regardless of the dynamics of interaction, that is, regardless of how or how often subjects meet, reproduce, exchange information, and choose their next plays in light of

past outcomes. It follows that if we think intuitively that knowledge can be of evolutionary or utilitarian value, then having the ESS property is a unique explanatory advantage of the tracking theory.

What does the added value of an ESS amount to intuitively, for a human being in real life? Suppose you want to go to Larissa. You believe route A will get you there by 12, and suppose this is actually true. Suppose that when p is true it is more valuable to you to be in a state believing p than not believing p and when p is false it is more valuable to you to not believe p than to believe p . To satisfy these conditions it is sufficient that you do want to get to Larissa by 12. If so, then these days you have a choice, whether to use a paper map or real-time GPS.

If you use a paper map then you may have a true belief when you start, and it may turn out that you continue to be right as you wend your way down the road, but there are other possibilities, such as road work and impassable traffic. If you use GPS (and it works) you will be prepared for these possibilities. Even without GPS you will of course learn of these obstructions when you get to them, but if you had used GPS (and it worked) you would have detected them ahead of time and had more time to adjust. You are more likely to get to your appointment in time if you use GPS. The obstruction scenario might not end up being actual – in which case you will turn out to actually have a true belief whichever method you use. But impassable traffic and road construction are possible at the beginning your trip, and remain realistic possibilities as you continue down the road.

What do you have in that GPS that you do not have in a map? You have strong tracking, in the specific sense of a strong and fine-grained disposition to believe p when it is true and not believe p when it is false. This disposition that you have now prepares you for non-actual possibilities where this road will not get you to Larissa by 12, possibilities that could become actual at any time during your journey. If you use a map then you will track a little – if that road did not even geographically link your current location with your destination, then a decent map would not advise you to believe it did – but the tracking is at a much lower level. It prepares you for fewer scenarios alternative to the actual, so your variation and adherence probabilities are conformably lower.

The general upshot is that if you use the GPS then you have a true belief plus a contingency detector, one value of which is preparedness, a value that you possess *now*

because of the GPS and hence tracking you have now. If your utilities concerning belief states are going to remain the same, then a tracking device is what you want to have now. This applies to any contingent proposition whose truth value matters to you, because its truth value may change.

Accounting for the value of knowledge in terms of its being an Evolutionary Stable Strategy may sound all tooth and claw, but there is nothing in the ESS property that restricts the value it brings to consequentialist or even practical terms. Here the true belief was valuable because we wanted to go to Larissa but there need not be any further thing, practical or otherwise *for which* the true belief states serve as a means to an end. You may now want to have true beliefs and not false beliefs about German literature just for the sake of it. As long as you now are someone to whom it is valuable to have such states in the future too, or to whom the continued existence of a type – the knower-of-German-literature type – is valuable, tracking-style knowledge will be more valuable to you now than mere true belief.

The contingencies that tracking makes you good against need not take the form of changes to the truth value of the proposition believed, as it does in the case of the road to Larissa. Even if no contingency will change the truth value of the proposition, as in the case of facts about historical German literature, there are still many contingencies that affect whether you believe the proposition, and do so differently in different circumstances. If you care about maintaining an appropriate belief state with respect to *p*, believing *p* if and only if it is true, then you need to care not only about believing or not believing *p* now when it is true or false respectively, but also that your belief state does not change in response to cues suggesting it is false though it is not or true though it is not. Tracking gives you that.

3. What is so bad about Gettierization?

Valuable though knowledge is, one does not need to care about it in the slightest in order to be concerned about Gettier cases. This is because, as I will argue, avoiding gettierization has a distinctive value that does not depend at all on the value of knowledge.

There is an impressive menu of interesting, intuitive, and fruitful ways of thinking of Gettier cases. I do not claim any of these accounts is wrong, but I will add another to the list, and characterize these cases in terms of probabilistic relevance. To begin with, a

bare bones intuitive way of describing the problem is that despite the subject's belief being justified, what makes the subject believe p is not what makes p true. Consider a case modeled on one of Gettier's original examples. Imagine an office with three workers, Smith, Jones, and Brown. Smith and Jones are friends but neither of them knows Brown. Smith was with Jones when he bought a Ford car, and since then Jones has given Smith a ride to work in it every day. So Smith believes q , that Jones owns a Ford, and has good reason to believe it. Smith also believes p , that someone in the office owns a Ford, whether he thinks about it explicitly or not. After all, that follows logically and trivially. And so, Smith has a justified belief in p , surely.

However, it turns out that a few years back Jones decided to change the status of his car, and now leases it instead of owning it. q is false, though Jones neglected to tell Smith about this change because it did not qualify as an interesting life event, and he is still driving the same car. We know that q false does not imply that p is false. p might be true for a different reason, and so it is in our example. Though Jones does not own a Ford, nevertheless someone in the office owns a Ford because, as it happens, Brown does.

Smith's belief in p is true, and was justified. He had good reason to believe it – we cannot expect him to rule out all logically possible ways of being wrong or we would never count anyone as justified. But a lot of us have the feeling he does not have knowledge, so this and many like it look like counterexamples to the traditional theory of knowledge. One way of describing why is that despite the fact that he is justified he still gets a true belief only accidentally, by double luck. He had bad luck with his justification, and then luckily had perfectly compensating good luck with the truth value of p .

Another way of describing what is wrong, that is present in all of the Gettier cases I have seen, is that what makes p true is not what makes the subject believe it. What makes it true that someone in the office owns a Ford is that Brown owns a Ford. What makes Smith believe that someone in the office owns a Ford are facts about Jones. I will gloss this situation by saying that what is *relevant* to S 's believing p is not *relevant* to a similar degree, to p 's being true, or vice versa. In our case, whether Brown owns a Ford is not relevant to whether Smith believes someone in the office owns a Ford. Only Jones is relevant to that. But whether Brown owns a Ford is relevant to whether in fact someone in the office owns a Ford.

I will express these relevance relations more precisely in probability. To begin:

$$P(b(p)/-q.r) = P(b(p)/-q.-r)$$

$$P(p/-q.r) \neq P(p/-q.-r)$$

The first equation says that the probability that Smith believes p given that Jones does not own a Ford is the same whether Brown owns a Ford or not. Smith is not paying any attention to Brown. The problem is that Brown does make a difference to whether p is true when Jones does not own a Ford. As in the second equation, the probability of p given that Jones does not own a Ford is different depending on whether Brown owns a Ford or not. Now when a Smith so described who believes p when he believes q makes an improbable mistake of believing q when q is false, the difference between his belief dispositions and the dependences between things in the world will be exposed. His belief in p will be either false or Gettiered, depending on whether Brown does not or does own a Ford respectively.⁴

We can get a better feel for the relevance point by writing it out more fully. First, here is what it is for q to be highly positively relevant to whether you believe p :

$$P(b(p)/q) \gg P(b(p)/-q)$$

You are much more likely to believe p given that q is true than you are given that q is false. q 's truth value is something you rely on in believing or not believing p . This can also be written as the ratio of those two terms being much greater than 1:

$$P(b(p)/q) / P(b(p)/-q) \gg 1$$

There is another, distinct, question, whether q is positively relevant to p 's being *true*. This can be represented as

$$P(p/q) \gg P(p/-q)$$

This says that p is much more likely when q is true than when q is false – q makes a large, positive difference to whether p is true – and it can also be written as the ratio of the two terms being much greater than 1.

These two things, the relation between q and p and between q and your belief in p , may be in line with each other or not. That is, those relevance relations can be *matched* or *mismatched*. We have a *relevance match for p on q*

$$P(b(p)/q) | P(b(p)/\neg q) \approx P(p/q) | P(p/\neg q)$$

when the difference that q makes to whether you believe p (the term on the left hand side) is approximately the same as the difference that q makes to whether p is the case (the right hand side). We have a *relevance mismatch for p on q* whenever

$$P(b(p)/q) | P(b(p)/\neg q) \not\approx P(p/q) | P(p/\neg q)$$

That is, when q 's truth value makes significantly more of a difference or less of a difference to whether you believe p – the left-hand side – than it does to whether p is true – the right-hand side.

In our Gettier case we had a relevance mismatch because q – Jones owning a Ford -- makes more difference to whether Smith believes p – someone owns a Ford – than it does to whether p is true. There is also a mismatch on r because r – Brown owning a Ford – makes more difference to whether p is true than it does to whether Smith believes p . So my picture of gettierization is this: when you have such a case you will always have a relevance mismatch on a q the subject is depending on for his belief in p , or a mismatch on some r that p 's truth value depends on, and either of these situations implies the other.

Although this condition will be fulfilled when you have a Gettier case, the converse does not hold. Not all cases of relevance mismatch are actual gettierizations, since for that the mismatch has to be exposed. The subject who tracks q has a grip that is usually good enough to avoid exposure of the mismatch: as long as our Smith believes q when q is true and does not believe q when q is false, he will not be believing p when q is false or not believing p when q is true either, and so will not be Gettiered, because he is a subject who is only using q to determine whether to believe p . So the question what is happening with that relevant proposition he is not paying attention to does not often arise for a subject who fulfills the tracking conditions for p . And this point is not restricted to the tracking conditions but would apply to the other known conditions on knowledge that go beyond true belief. Note that this characterization of gettierization is not restricted to the false-lemma form of the Gettier case just described. There will

always be a q or q's whose truth values the Gettiered subject's believing p or not is responsive to, whether the subject has a belief state with respect to the q's or not. That this analysis does not depend on the subject having a belief state with respect to q is borne out in the relevance match and mismatch equations, where the conditions in the conditional probabilities are all q and -q, and involve no b(q) or -b(q).

To summarize my characterization of Gettier cases, we need three things, 1) that the person bases belief in p on q when q is false, 2) that there is a relevance mismatch on q for the first condition to exploit, and 3) that p is true. 1) and 2) bring out the distinctive element in a Gettier case by focusing on the dependence on q of the subject's belief in p. Improving tracking of p by being more sensitive and adherent to p decreases the probability of the first feature, by either reducing or improving the quality of the subject's reliance on q.

As for the second feature, obviously increasing relevance matching decreases the probability of 2) directly. We can also ask whether improving the tracking of p reduces the probability of 2) by asking whether improved tracking of p improves relevance matching for p on q. Conveniently, the relationship between tracking and relevance matching can be expressed very neatly using total probability:

$$\mathbf{P(b(p)/q)} = \frac{P(b(p)/p)P(q/b(p).p)\mathbf{P(p/q)}}{P(q/p)} + \frac{P(b(p)/-p)P(q/b(p).-p)P(-p/q)}{P(q/-p)}$$

$$\mathbf{P(b(p)/-q)} = \frac{P(b(p)/p)P(-q/b(p).p)\mathbf{P(p/-q)}}{P(-q/p)} + \frac{P(b(p)/-p)P(-q/b(p).-p)P(-p/-q)}{P(-q/-p)}$$

The ratio of the bold terms on the left-hand sides expresses the relevance of q to whether I believe p, i.e., the probability that I believe p given q divided by the probability that I believe p given that q is not the case. The ratio of the bold terms on the right-hand sides expresses the relevance of q to p, i.e., the probability that p is true given that q is true over the probability that p is true given that q is false. For relevance matching we want the relevance level of q to p, the ratio of the bolds on the right, to translate through these other terms into the same relevance level of q to my believing p, the ratio of the bolds on the left; we want the ratio of the left bolds to equal the ratio of the right bolds.

These equations seem to be complicated to evaluate, but some regularities are immediately apparent. For example, the level of relevance matching for p on a given q depends on how well I fulfill the tracking conditions for p:

$$P(b(p)/q) = \frac{\mathbf{P(b(p)/p)}P(q/b(p).p)P(p/q)}{P(q/p)} + \frac{\mathbf{P(b(p)/-p)}P(q/b(p).-p)P(-p/q)}{P(q/-p)}$$

$$P(b(p)/-q) = \frac{\mathbf{P(b(p)/p)}P(-q/b(p).p)P(p/-q)}{P(-q/p)} + \frac{\mathbf{P(b(p)/-p)}P(-q/b(p).-p)P(-p/-q)}{P(-q/-p)}$$

The leftmost bold terms express the level to which I fulfill the adherence condition, and the rightmost bold terms completely determine and are determined by the level to which I fulfil the variation (sensitivity) condition for p.

All of the terms set between the tracking conditions and the terms evaluating relevance matching make the relationship between the two features multi-dimensional, but in the extreme cases the relationship is simple. If I am perfectly sensitive to p – the first tracking condition – then I never believe p when it is false, and $P(b(p)/-p)$ and the whole second summands with them, go to zero in both equations:

$$P(b(p)/q) = \frac{P(b(p)/p)P(q/b(p).p)P(p/q)}{P(q/p)}$$

$$P(b(p)/-q) = \frac{P(b(p)/p)P(-q/b(p).p)P(p/-q)}{P(-q/p)}$$

If in addition I am a perfect adherer – I always believe p when p is true, the second tracking condition – then $P(b(p)/p)$ is 1. That forces

$$\frac{P(b(p)/q)}{P(b(p)/-q)} = \frac{P(p/q)}{P(p/-q)}^5$$

That is, perfect tracking of p implies perfect relevance matching for p on q. The same argument can be done for p and any q, so perfect tracking of p implies perfect relevance matching for p on every q. Thus, if I fulfill the extra condition for knowledge beyond true belief infallibly, then I cannot be gettiered. Of course, any theory with an infallibilist requirement can disqualify Gettier cases, because such a theory allows no room at all

for luck or mismatch. What is most interesting is what happens to a subject who tracks fallibly. The fallibilist tracking theory, with thresholds s and t set at less than 1, disqualifies the Gettier case I went over and many standard cases, though presumably not all possible cases.

However, it is pretty clear a priori that there is no fallibilist theory that disqualifies all such cases (Zagzebski 1994), so the interesting questions are comparative ones of better and worse. Fallible tracking addresses the Gettier problem directly by directly reducing the chances of 2) above. This is because not only does perfect tracking of p imply perfectly relevance matching for p on every q , but also the better you track p the better you relevance match for p on every q . Fallible tracking will not make gettierization impossible, but improving your tracking will always make it less likely. To put that into context, there are fallibilist theories for which this improvement relation does not hold. For example, though infallible justification will eliminate the possibility of gettierization, increasing internalist justification of belief in p does not necessarily reduce the chances of mismatch, i.e. of gettierization, because internalist theories have no condition at all on the relation of the subject's belief to the world other than that the belief actually be true. Since relevance matching is a relation between one's beliefs and the world, it seems there is no way for the conditions internalist theories do have to guarantee improvements of relevance matching. It would be of interest to examine this question for non-tracking externalist theories, for example whether improving the reliability of one's process of belief formation necessarily increases relevance matching.

However, given that fallible tracking allows Gettier cases as possible one might wonder how I can take myself to have addressed Kvanvig's argument that it is not possible to understand both the nature and the value of knowledge simultaneously because of Gettier cases. This is because Kvanvig's argument was not that no criterion could rule out all Gettier cases – we already knew that no fallibilist theory is likely to be able to do that. His point was that when you come up with conditions that seem to explain the value of knowledge they face the Gettier problem, and the fixing that you have to do to address those cases is ad hoc from the point of view of the condition that insured the value of knowledge.

For that claim, what is important in what we have seen so far is that the tracking conditions got no ad hoc tweaking, no adjustments at all, when they went from explaining the value of knowledge to directly reducing the chances of gettierization. The

same conditions did both. The tracking conditions addressed what is distinctive about Gettier cases specifically, rather than just generically decreasing luck: not only does improved tracking of p make it less likely that a subject will believe p when q is false – the first thing we need for a Gettier case – but also better tracking makes it less likely that the subject will be gettiered if he does believe p when q is false – the second, relevance mismatch, aspect – because his believing p or not will have a greater responsiveness to whether or not r is true, for all r relevant to p .

Looking in the other direction, from relevance matching to tracking, we find a few more interesting relationships. For the first claim I made, that perfect tracking implies perfect relevance matching, the converse is also true. Perfect relevance matching for p on every q implies perfect tracking. However, while better tracking of p improves relevance matching for p on every q , increasing your relevance matching for p on *some* q does not necessarily increase your tracking of p .

A diagram can help to explain these claims.

Perfect tracking

Perfect relevance matching

If you track p perfectly, then whether you use an indicator q or not, your belief in p co-varies perfectly with p . So the proposition that you believe p , $b(p)$, is substitutable for p in every relevance relation p has to any q . So, your belief in p relevance matches for p on every q . In the other direction, if your belief in p perfectly relevance matches for p on every q that means that the belief in p stands in the same relation to each q as p stands to that q . So your believing p is substitutable for p in every relation that could affect the truth value of p . That is, you track p .

So far perfection. Human beings are not perfect trackers, and although fulfilling the tracking conditions does not require the use of intermediaries, creatures like us invariably need them. So, imagine that your believing p relates to the truth value of p fallibly via an intermediary, an indicator, q .

Your believing p depends on q strongly as a proxy. Then to what extent you track p depends on how good a proxy q is for p . If it is a perfect proxy for p – indicated by the solidity of the arrow between them – then q is true if and only if p is true, and the fidelity of your tracking of p is limited only by how closely your belief in p depends on that perfect proxy. If q is not a perfect proxy for p – indicated by the arrow between them being dashed – then your tracking has two limitations,

so the relation between p and your belief in p is weaker than in the previous diagram. It is also necessarily weaker than either the relation between $b(p)$ and q or that between p and q .

For relevance matching in the imperfect case, let us call the proposition that your believing p uses as an indicator q_1 , to make room for additional factors.

When q_1 is a better proxy for p , other things equal, then your tracking of p is better. But also, the substitutability between q_1 and p is tighter, and that makes q_1 vary with the other factors in p 's web of relevance in a way closer to exactly the way p does. Because your belief in p is following q_1 , your belief in p automatically becomes more relevance matched to those other q 's as q_1 becomes more substitutable for p in all of p 's relations to other propositions. Better tracking of p improves relevance matching on q for every q .

In the opposite direction, if you improve your relevance matching for p on every q , your belief in p becomes more and more substitutable for p in every relation that p stands in to every q . That is, your believing in p or not becomes more substitutable for p in every circumstance that makes a difference to whether p is true or false. Thus, your belief in p tracks p better. However, you may improve your relevance matching for p on some q , say q_2 , by getting your belief in p into a relation to q_2 that is more similar to p 's relation to q_2 , but achieve this at the expense of your belief's nice relation to another q , q_1 . Maybe you outsource detection of q_1 so that you can get a handle on q_2 , and your third party is not as good as you are at detecting q_1 . In this case you improve your relevance matching to one q but because you reduce it with respect to another q your believing p or not does not necessarily become more substitutable for p 's being true or not in all the relations p stands in to other matters. It does not necessarily increase the chances your belief state is right in all of the circumstances corresponding to the truth or falsity of those q 's, which means you do not necessarily track p any better. Improving relevance matching for p on one q , or any finite number of q , does not imply improved tracking of p .

From all of this we can clearly see that gettierization avoidance – i.e., relevance matching – is related to and not gerrymandered with respect to the requirement that

gives us an account of knowledge – i.e., tracking. That improved tracking implies improved relevance matching and so decreases the chances of gettierization answers to the intuition that knowledge requires gettierization avoidance. That relevance matching and tracking are nevertheless quite distinct properties is the reason we should not think of the project of giving an account of the nature of knowledge as defined by chasing down Gettier cases.

4. The Value of Relevance Matching and the Nature of Understanding

In answer to Kvanvig's challenge, and further to my point that gettierization-proofing is not per se a defining ingredient of knowledge, I will explain why gettierization avoidance is valuable without presupposing that knowledge is valuable. For this it will suffice to explain why relevance matching is valuable even if we do not value having an ESS in the true-belief game.

We know that tracking is the best insurance that the appropriateness of our belief states will survive time and changing circumstances, for propositions about which it matters that we have correct belief states. What is relevance matching intuitively, and what does it give us? I think we can naturally say of the gettierized subject discussed that he does not understand why it is true that someone in the office owns a Ford, and I propose that relevance matching for p on a sufficient set of q 's is the key to understanding why p is true. For current purposes I will assume as a brute fact that understanding is valuable.

Intuitively, understanding why p is true is different from knowing that p is true, and this account will be able to explain how. Whatever the value is of understanding why a proposition p is true, or why it would be true if it were true, it is not merely for possessing or even preserving correct belief states with regard to p 's truth value. Indeed, when we ask why p is true or why it would be true if it were true, we presuppose that it is true or hypothetically imagine that it is, respectively. Although understanding why p might be true could help us determine whether it is, p 's truth value is not per se the concern of the "Why?" question, and this is why the actual and hypothetical "Why?" questions can sensibly be given the very same answer. A second intuitive difference between knowing that p is true and understanding why p is or would

be true is that understanding involves appreciation of more than p alone. Relevance matching will give us a compelling picture of what that appreciation should be like.

The fact that knowledge and understanding are different has a further significance. Combined with my view that the value of gettierization avoidance is improved understanding, it provides some justification for my resistance to the expectation that a theory of what knowledge is must rule out all Gettier cases. We do not need to understand why p is true in order to know that it is. In the Gettier case discussed here we do not want to have to say that the subject should check whether Brown owns a Ford before we give him the knowledge star on his forehead, and we do not have to. Smith can know that p is true despite his defect of not understanding why it is true.

There are a variety of options for stipulating how much relevance matching there must be, and to what set of q a belief in p must be relevance matched, in order to count as understanding why p is true, and in what way. I will start with a very strong definition of what it is to improve one's understanding of why p is true, and then explain some ways it can be qualified, extended, and multiplied. This definition will also serve as the template for comparing levels of understanding in the static case since one state's being an improvement over another is the same as one state being better than another according to the relevant criteria.

As a first pass, *understanding of why p is true improves*⁶ only if there is an increase in relevance matching for p on some q, and no decrease in relevance matching on any other q of such a degree as to outweigh the increase.⁷ We will see below that we need a further condition for a full definition of understanding. We also need to recognize some weaker but significant variations on and partial versions of the condition just stated. For example, there can be a kind of increase in understanding if we improve relevance matching on some q while losing some matching on others. Maybe we stopped caring about the q's of a particular type; our appreciation of these newly matched q's would still have increased our understanding of the things we care about.

Also, the definition of improvement in understanding just given can easily be generalized to the case where we understand better by appreciating more about why p would be true if it were true, whether it is, or we know it is, or not. The current account identifies multiple ways that one can have partial understanding. One might have overall quite good relevance matching for p on all q while the tracking level of p that this

brings one does not meet the threshold chosen for knowledge. One might understand why p would be true although it is actually false because one may relevance match for many q, and even track p, while p is false. Relevance matching and the tracking conditions are dispositional and fallibilist. For this reason one may also relevance match for p on many q's while some of those q's are actually false, and nothing says one does not actually believe they are true. Of course, frequently actually having belief in a given q when that q was not true (or vice versa) would not be consistent with relevance matching on q, because it is not consistent with having the right dispositions to believe p. But these general points agree with an intuitive sense many have that one can have a good deal of understanding of why p is true, or would be true, without all of one's beliefs about the surrounding subject matter being true.

Understanding has the dimensions of breadth and depth. Breadth would correspond to relevance matching to a large number of q's spread from p far out into p's web of relevance. Depth would come from the density of the q's one relevance matches to, from the number of q's one matches to between p and a given q or between one q and another. Some of the kinds of intermediate q's would be causes, correlated matters, and laws and generalizations of which the relation between q and p or between two q's are instances.

It is not sufficient for understanding to have knowledge of a long conjunction of propositions in the subject matter of p, as Elgin put it when she drew attention to the feature of understanding that involves an appreciation of the coherence and connections among those propositions. (Elgin 2009) Relevance matching evidently requires appreciation of the connections, since one's dispositions to believe p are required to mirror p's relevance relations and lack of relations with all other propositions, or an appropriate subset of them.

The partial definition of improvement of understanding, combined with our earlier point that increase in relevance matching for p on all q increases the level of tracking of p, imply that improving one's understanding of why p is true will improve one's tracking, hence knowledge, of p. Combining this with the fact cited earlier that increasing tracking of p increases relevance matching for p on all q we see that at the most abstract level, knowledge and understanding as so far defined are very closely related. Indeed, as we saw above, perfect tracking and perfect relevance matching are logically equivalent. This may seem to make the two phenomena more closely related than intuition allows.

Knowledge and relevance matching, hence understanding, have different values, as we saw earlier, and logical equivalence does not make them the same thing, but there are further pragmatic differences, to be discussed presently, and in the final analysis here the concepts will not be logically equivalent either.

First the pragmatic differences. We live in the short and medium run, and whether we approach epistemic harmony with p by aiming at knowledge or at understanding can lead to dramatically different consequences at these timescales. For contingent reasons, unlike gods we have to use intermediaries in order to track p . At any given point we only have enough resources to address some factors relevant to p and not others.

Sometimes a few intermediaries are enough to track p ; sometimes it takes only one indicator. This is what we should look for if our goal is simply to be able to predict or be in line with the disposition of p 's truth value in the future. Tracking p logically implies a high level of relevance matching to all q – which creates a problem I will address below – but a grasp on one indicator will not make us aware of the other q that p and my belief in p are so similarly related to. It may give us nothing at all to say about other factors relevant to p .

Instead of searching for an indicator that will allow us to efficiently follow the truth value of p , we could choose to spread our resources and investigate multiple q 's, trying to get at their relation to p . Since we would not necessarily increase net relevance matching we would not necessarily improve tracking of p , or understanding strictly speaking as so far defined, and if we did improve tracking it would not be by the most efficient means. For purposes of getting a handle on p 's truth value this would be a roundabout strategy of inquiry. However, in improving relevance matching to some q we would have made the kind of partial progress toward understanding what makes p true discussed above. Moreover, as we went about this we would sometimes, though as I will explain below not necessarily always, be making ourselves aware of more factors relevant to whether p is true, whether they suffice together to be a highly tracking indicator or not. Although I do not think awareness of the factors is necessary for their useability, there are definitely situations where awareness can make them more useable. Whether you take a knowledge approach or an understanding approach to p , your efforts will also take you along the road to the other, but since you will typically have different results in the short run, which approach you choose matters very much.

Knowledge and understanding are even more different than the partial definition of understanding provides for, and different for more than pragmatic reasons. I drew out above how it is that one can have a lot of understanding of why *p* would be true without knowing that *p* is true. One can also know that *p* is the case without understanding why it is, and this is not merely a pragmatic matter, but goes to what understanding is.

One has knowledge without understanding when one uses a good indicator but has no grasp of other relevant matters. If your thermometer is in good working order then you can know that Mr. Washington has a fever merely by having your belief that he has this symptom co-vary with the thermometer reading. That is because the thermometer will give you *tracking* of his body temperature – you will believe that he has that temperature roughly if and only if he does. But this tracking surely does not give you understanding of why he has this symptom; understanding in this case would involve a diagnosis, or an account of the physiology of temperature regulation in mammals. We saw above that it is a logical fact that fulfilling the tracking conditions for *p* implies that your belief in *p* has a high level of relevance matching on all *q*, that is all things related to the fact that Mr. Washington has a fever – the more substitutable your belief in *p* is for *p*, the more it will stand in the same dispositional relation to every *q* as *p* stands to that *q*. So the condition that increase in understanding requires increase in net relevance matching alone would imply that a thermometer can give you understanding of why a patient has a fever, and thus is too weak to provide a full definition.⁸

We might think that the best way to fix this is to make a distinction between indicators of *p* and what makes *p* true, and to require that the subject be in touch with the latter as well as the former. However, this would not address the problem since tracking *p* already makes the subject relevance match for *p* on the *q* that make *p* true because it logically implies that the subject relevance matches on every *q*. The problem is not with which *q* she is required to match, but the nature of the matching. What the subject merely following the thermometer does not have is states corresponding to the *q*'s. If the subject tracks *p* then her belief in *p* will co-vary with the *q*'s as *p* co-varies with the *q*'s, but she will not necessarily own anything that corresponds to the *q*'s.

A natural thought is to make a further requirement that the subject have beliefs about the *q*'s in *p*'s web of relevance, but I think that this would be too strong. Having beliefs about the *q*'s suggests an ability to call up those considerations and talk about them.

Some epistemologists have endorsed a requirement of reflective access in order to be counted as understanding. Thus, Duncan Pritchard:

Understanding clearly is very amenable to an account along internalist lines, in the sense that it is hard to make sense of how an agent could possess understanding and yet lack good reflectively accessible grounds in support of that understanding. Understanding thus cannot be 'opaque' to the subject in the way that knowledge, by externalist lights at least, can sometimes be. (Pritchard 2015)

However, I do not think this jibes with how we use the word "understanding". There are many cases of understanding where our grounds are not reflectively accessible. If I understand why she married Gabriel, that does not imply that I could list all of the factors relevant to her decision. I would probably be able to list some, but for the rest I could say merely that I have a feel for how she thinks and their relationship, and if I know her, Gabriel, and the situation well then I think we would allow that I have more understanding than I am able to explain. Again, if I say I understand why he is traumatized, I do not mean that I could list the elements of the torture he was delivered, and explain how each affected his psyche. An indication that I understand, and even, possibly, that with respect to some factors I understand more deeply than an ability to speak about it would show, is my feeling a shiver when I hear a description of things that were done to him. I agree with Elgin 2009 that it is not understanding if it is not useable by the subject, but states that are not reflectively accessible can enable and otherwise affect our actions, thus be useable, and just as they can co-vary with the belief in *p*, they can co-vary with other beliefs, including beliefs that are reflectively accessed and used explicitly in inferences.

I reject the requirement that the subject who understands have beliefs corresponding to the *q*'s she relevance matches on, but the belief requirement is on the right track in requiring that the subject own something that is disposed to "move" in relation to the subject's belief in *p* in the same way that the *q*'s' truth values are disposed to move in relation to the truth value of *p*. Thus, I will not only require that the subject's belief in *p* move with respect to the *q*'s in the same way that *p* moves with them, i.e., relevance matching, but also that the subject have states corresponding to the *q*'s relevant to *p* that move with the subject's belief in *p* in the same way that the *q*'s move with respect to *p*. The subject will thereby possess in herself a dynamical copy of *p* and its web of relevance. The subject will possess not just a static image like a photograph but a simulation, in the strict sense requiring that the dynamics of the simulation follow the

same laws as the process being simulated. The same simulation quality will obtain when the states in the subject that correspond to the q_i happen to be beliefs, in virtue of the dynamical web of relevance matching that those beliefs achieve. Although the view of understanding proposed here is not in any way restricted to the case of understanding human behavior, it is suggestive as a way of fleshing out the simulationist view of how we read other minds.¹

The contrast between understanding as mere relevance matching, which is too weak, and the subject relevance matching by means of her own states can be seen in the following diagrams:

¹ When the subject's states that carry the relevance matching are beliefs, then the person's understanding may be reflectively accessible as a theory, in the broad sense of that term used when discussing our knowledge of other minds. In the terms of that discussion, when the subject relevance matches via states that are not beliefs she would be seen as simulating. Both count as understanding on my view in virtue of the relevance matching rather than because of the type of state that carries it, a view I develop further in Roush (2016), and both have the character of simulation.

The subject on the bottom not only has a belief in p that is dispositionally related to the q_i 's in the same ways that p is related to those q_i 's but also has states s_i that her belief is related to in the same ways as p is related to the q_i 's. If the relevance matching relations are fulfilled, then the s_i and q_i will be related in a regular way – the horizontal dotted arrows below -- that is a relation in which the s_i 's track the q_i 's, respectively, to some degree, in the sense of tracking defined above.⁹ This would look as follows:

It is tempting to require that the subject's states, s_i , that correspond to the q 's be if not beliefs then at least mental states, but the example of the shiver makes me reluctant to go even that far. A shiver is, or at least includes, a bodily reaction, and this is not merely accidentally related to the understanding, since the torture itself was inflicted on the body. If the states need not be mental they could be located in the foot, though, and how absurd it would be to think we could understand why p is true via a network of foot-dispositions! However, the foot states still have to be dispositionally related to the

subject's belief in p in the right way in order for the set-up to count as understanding, and the belief in p is at least partly a mental state, so it is not as if the mind can be entirely uninvolved. Understanding by foot might be inefficient and tiring given that a foot also has other things to do, but that is a pragmatic matter.

Understanding as simulation, as participation of the understander in a working model of what the understandee undergoes, provides a suggestive picture of what is happening in the phenomenon of compassion fatigue and burnout among mental health workers and emergency room nurses. Understanding by foot might sound hilarious but if caring human beings cannot suppress the surely unconscious impulse to simulate, if no one can understand without simulating, then a strategy of scaling down the simulation and displacing it to a different location in oneself than the phenomenon is in the sufferer could be helpful for those caring for distressed and traumatized people to avoid becoming distressed and traumatized themselves. To account for the fact that it is possible to have knowledge without understanding I think we have to require that the subject own states corresponding to the q 's and not just the p , but it should not be required that those states be always located in a particular part of the understander.

Summary

Relevance matching is a good way to make explicit both what is missing in Gettier cases, and some essential features of understanding why proposition p is true, for any p . It allows us to see that avoiding gettierization is valuable without presupposing the value of knowledge and to trace out some key similarities and differences between knowledge and understanding. It allows us to account for the dimensions of breadth, depth, connectedness, and useability of understanding, and to explain why it is possible to have knowledge that p is true without understanding why it is true, and how it is possible to gain some understanding of why p could be true without gaining knowledge of whether it is. The picture of understanding as simulation that emerges, when in the final definition we add a condition of ownership of a web of states in addition to belief in p , is suggestive as a way of fleshing out what it means to understand other minds.

References

Dretske, F. (1971), "Conclusive Reasons", *Australasian Journal of Philosophy* 49: 1-22.

- Elgin, Catherine Z. (2009), "Is Understanding Factive?" In Adrian Haddock, Alan Millar & Duncan Pritchard eds., *Epistemic Value*. Oxford: Oxford University Press. 322-330.
- Gettier, E. (1963), "Is Justified True Belief Knowledge?" *Analysis*, 23: 121-123.
- Goldman, Alvin (1967), "A Causal Theory of Knowing", *Journal of Philosophy* 64: 357-372.
- (1976), "What is Justified Belief?" in G.S. Pappas ed. *Justification and Knowledge*. Dordrecht: D. Reidel, 1-23.
- Mark Kaplan (1985), "It's not what you know that counts", *Journal of Philosophy* 82 (7):350-363.
- Lehrer, K. and Paxson, T. (1969) "Knowledge: Undefeated Justified True Belief", *Journal of Philosophy* 66: 225-37.
- Kvanvig, Jonathan (2003), *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Nozick, Robert (1981), *Philosophical Explanations*. Cambridge, MA: Harvard University Press, 172-185, 197-217.
- Pritchard, Duncan (2005), *Epistemic Luck*. Oxford: Oxford University Press.
- (2015), "Knowledge, Understanding, and Epistemic Value," section 9 in *Epistemology* ed. A. O'Hear, Cambridge: Cambridge University Press.
- Roush, S. (2005), *Tracking Truth: Knowledge, Evidence, and Science*. Oxford: Oxford University Press.
- (2010), "The Value of Knowledge and the Pursuit of Survival", *Metaphilosophy* 41: 255-278.
- Sosa, Ernest (1991), *Knowledge in Perspective*. Cambridge: Cambridge University Press.
- Zagzebski, L. (1994). "The Inescapability of Gettier Problems", *Philosophical Quarterly* 44: 65-73.
- (1996), *Virtues of the Mind*. Cambridge: Cambridge University Press.

¹ I am grateful to audiences at the Eastern APA 2012, University of British Columbia, UC Berkeley Logic Colloquium, King's College London, and Kyoto University, and to Michael Della Rocca, David Papineau, Bill Brewer, and Jake Wojtowicz in particular for helpful feedback.

² Kaplan's argument is general, including whatever criterion one might like to substitute for justification.

³ Decomposed this way those X conditions that are not tracking will typically be disjunctive, but that just makes them what are called "mixed" strategies.

⁴ I assume that there are no other people in the office. Otherwise Brown's not owning a Ford would not imply the falsity of "Someone in the office owns a Ford".

⁵ $P(b(p)/p) = 1$ implies $P(b(p).p) = P(p)$. From there, algebra gives $P(q/b(p).p)/P(q/p) = 1$, and likewise for the $\neg q$ case. So, $P(q/b(p).p) | P(q/p) = P(\neg q/b(p).p) | P(\neg q/p) = 1$. This implies $P(b(p)/q) = P(p/q)$ and $P(b(p)/\neg q) = P(p/\neg q)$, which together imply $P(b(p)/q) | P(b(p)/\neg q) = P(p/q) | P(p/\neg q)$, i.e., perfect relevance matching for p on q .

⁶ The concept I am defining takes as object why a proposition is, or would be, the case in contrast to many who define a concept of understanding a subject matter, e.g., the Comanche dominance of the Southern Plains. (E.g., Kvanvig 2003, Elgin 2009, Pritchard 2015) The reason usually given for the latter is that understanding that p is not appreciably different from knowing that p . But “I understand that p ” sounds like “I gather that p ”, and understanding why p is, or would be, the case, is definitely much more than either of those.

⁷ Clearly this comparison will require a measure for relevance matching. My preference is the likelihood ratio, but the issue is beyond the scope of this paper.

⁸ The fact that tracking brings relevance matching raises the question how a Gettier case as defined in this paper is possible at all. The reason is that both concepts, knowledge and understanding, are defined fallibly, and the character of the spheres of cases in which we are allowed to fall short is different for the tracking and relevance matching criteria, even if they are, so to speak, the same “size”. This is an issue to which the way of measuring relevance matching will likely matter.

⁹ The relations between the q_i 's and $b(p)$ and between $b(p)$ and the s_i 's will themselves will be tracking relations of some degree in both directions, and I showed in Ch. 5 (2005) that tracking has the property of *transitivity enough*, so each s_i will track its q_i to a derivable minimal level that depends on the levels of the two tracking relations supporting their relation.