Sherrilyn Roush - 2017

Bayesian Re-Calibration: A Generalization

Roush (2009) derived a probabilistic framework for updating one's first-order degree of belief in q in light of evidence about one's own reliability in making q-like claims, thus providing the probabilistic rationality constraint for resolving epistemic self-doubt. In this note the argument is generalized to the case where the evidence about one's reliability or one's degree of belief in q is uncertain, by development of a Jeffrey-style version of the Re-Calibration equation. This allows illustrative applications of the framework to examples where higher-order evidence is of varying qualities. The equation is applied here to the familiar examples of hypoxia and peer disagreement.

There are many situations in which information about our own reliability is available to us, and some where it is at odds with our confidence in first-order propositions. An emergency room doctor may quickly come up with a diagnosis for a patient, and then reflect that he hasn't slept in 36 hours. (Christensen 2010a) An eyewitness for a murder may be confident that the murderer is number 3 in the line-up – how could she forget such a vivid event? – and then remember reading that eyewitnesses are generally quite overconfident, especially when the event was stressful. A pilot may calculate that she has enough fuel to easily make it fifty more miles than her original plan, but on glancing at the altimeter reading of 10,500 feet she remembers the condition of hypoxia, a lack of circulating oxygen that one is susceptible to at altitudes above 10,000 feet, that causes cognitive impairments, and that is rarely noticed by the person herself at its initial onset. (Christensen 2010b)

Such cases present us with the questions: Does rationality require us to revise our original confidence? If so, then by how much? And on what grounds are these answers justified? In the Bayesian Recalibration approach we represent this question what to do with evidence of our (un)reliability explicitly using higher-order probability. Evidence about whether we are reliable judges can come in many shapes and sizes, from many directions, and with a variety of types of things affecting reliability. It may come from testimony – a friend tells you that you've been slipped a hallucinogenic drug – or track record – as when, or if, a weatherman keeps a record of how often it rained when he was confident it would – or knowledge of a mechanism – as when we know our eyes won't be giving us information because we're blindfolded. But all evidence about our reliability has in common that it is information about the relationship between our beliefs and the world, roughly whether our beliefs formed in the way they were when coming to a degree of belief about q tend to be true. In the Bayesian Re-Calibration framework, what the multifarious evidence says about this bottom line is what determines whether we should revise our first-order belief in light of it.

One way to make the concept of a level of reliability of our beliefs precise is through the notion of a calibration curve. We can state the relation of the subject's belief-state concerning q to the way the world is – her reliability – as an objective conditional probability function of her degrees of belief, represented as probabilities (thereby assuming she is coherent):

$PR(q/P(q)=x) = y$ Calibration Curve

The objective probability of q given that the subject believes q to degree x is y. This is a function according to which reliability y[1] varies with the independent variable of confidence, x, with different variables used in order to allow for the possibility that the subject's degree of belief tends not to match the objective probability, and that the level and direction of mismatch can vary with the level of confidence.[2] The curve is specific to proposition q and to the subject whose probability function is P. A subject is calibrated on q, on this definition, if his calibration curve is the line x = y.[3]

Calibration curves are widely studied by empirical psychologists who find that even though the curves vary by individual, there are regularities in the human population. In tests taken in controlled settings, human beings' reliability tends on average to vary systematically and uniformly with confidence, with for example high confidence tending to overconfidence, as in eyewitness testimony, and low confidence tending to underconfidence. The latter is an example illustrating the fact that if we represent reliability via a calibration curve we get a general framework which allows for not only self-doubt but also cases where the news about one's reliability may warrant an increase rather than a decrease in one's confidence about q.

Within these population averages, calibration curves vary with subject matter, sub-group, individual traits, professional skills, and particular circumstances. All kinds of evidence about a subject's belief-forming processes, methods, circumstances, track-record, and competences are relevant to estimating this function. In real life no one could get enough evidence to warrant certainty about an individual's calibration curve for an individual proposition q in a particular set of circumstances, but Bayesianism allows us to use whatever evidence we do have to form a rational degree of belief about what a person's calibration curve is, or what value it has for some argument x, or at least to figure out what it would take to determine an appropriate final degree of belief. The novel but natural step of Bayesian Recalibration is to make use of the fact that one can have such a confidence about one's own calibration curve.

A person having received news about her reliability about q that is at odds with her confidence about q is in the following type of state. She is confident and more or less correct that she believes q to degree x, that is, P(q) = x, but also has an uncomfortably high level of confidence, say $\geq .5$, that she is less or more reliable than x about q at that confidence. That is, she has confidence $\geq .5$ that the objective probability of q when she has x-level of confidence in q is different from x, which we would write P(PR(q/P(q)=x) ≠ x)

---

[1] In the empirical psychology literature on calibration, this quantity is usually referred to as accuracy rather than reliability.

[2] In this it is distinct from the notion of reliability given by Alvin Goldman (1980) and applied to processes for forming beliefs. With that concept a belief was reliably formed if the process that produced it gives true beliefs most of the time. Both calibration and reliability in this sense concern the general relation of one's belief-like states to the truth, but the calibration concept measures match of confidence with objectivity probability rather than flat-out beliefs with truth or falsity, and it can characterize a concerning mismatch in either direction, of over- or under-confidence.

[3] The calibration curve is distinct from the calibration *score*, which is the calibration-related idea most often discussed in philosophy and which has been criticized and largely dismissed, on grounds laid out in Seidenfeld 1985. The calibration score is a root mean squares measure of how far a person's calibration curve is from the x = y line. A single score has a one-many relationship with an infinite number of calibration curves, so a lower calibration score does not uniquely determine how a reduction in the score has been achieved. Consequently one may improve one's calibration score at the expense of informativeness, e.g., having every day the same confidence in rain, a confidence that matches the annual frequency of rain. The use of the calibration curve here is very different. Due to the fact that the subject will be conditionalizing on it, what the person should do with her confidence is uniquely determined, and the improvement is not judged by its effect on her overall calibration score.

$\geq .5$. Let us call that different value y, so that $P(PR(q/P(q)=x) = y) \geq .5$, $y \neq x$. Whether one has underslept and thereby might be less reliable than one is confident, or one's gender and upbringing have made one less confident than one is reliable, whether or not the reason for unreliability is that one tends to mistake evidential support relations and whether one does or does not think a given evidential support relation obtains, make no general difference to what reliability is for purposes of recalibration. Reliability here is simply about how far one tends to get things right when doing the sort of thing one did in coming to be confident in q to level x;[4] it is about the relation between one's confidence and the way things are.

Thus the state one is in when new and unsettling evidence about one's reliability arises looks as follows:

$P(q) = x$

$P(P(q) = x) = .99$ (high)

$P(PR(q/P(q) = x) = y) \geq .5$, $y \neq x$

You actually believe q to degree x, you are confident (say at .99) that you so believe, and you have an uncomfortably high level of confidence that you are not calibrated for q at x, that the objective probability of q when you are x confident of q is y, y different from x. This state escapes incoherence for two reasons. One is that one's confidence either about one's degree of belief or one's reliability is not 1, and unlike some conditional probability formulations of self-doubt (Roush forthcoming), the slightest uncertainty about either of these things is enough to make it coherent to attribute any degree of discrepancy between your believed confidence and your believed reliability.

This is made possible by a second factor, that (un)reliability is expressed here as an objective conditional probability, and coherence alone does not dictate how subjective and objective probabilities must relate. In evaluating my own P I compare it to a different probability function, PR, that gives an objective probability for q given my subjective degree of belief in q, and coherence doesn't dictate the relationship between two different functions. In this case the second function is not an expert function P_E that declares unconditionally what value the subject with maximal knowledge or rationality would give to q, $P_E(q) = x$, as we find in some uses of higher-order probability (Gaifman 1986, Christensen Elga maximally rational agent), but a calibration function, a conditional probability that tells one what objective probability is indicated by one's subjective probability. The calibration function is specific to the subject herself and its use in coming to her final confidence means that her initial confidence in q is not left behind but integrated, as we will see in examples below. Another difference between the expert-function and the calibration-curve approaches is that there are obvious ways to investigate calibration curves empirically. Expert functions could be compiled from opinions of those identified as experts, but it might be hard to recruit enough maximally rational subjects for a statistically significant study so we might tend to be left appealing to intuitions about what seems rational.

Once the defeating information about the relation of a subject's credences to the world is expressed in objective probability it can be represented explicitly as a consideration the subject takes on board in assessing the quality of the degree of belief she takes herself to have in q and resolving the question what her degree of belief should be after she takes the evidence about her reliability on board, thus:

[4] What sort of mistake might have been made is relevant to the estimation of the calibration curve for a particular person and proposition and occasion, since the reliability of one's answer will also vary with method. But the particulars of why you are uncalibrated do not make a difference to how you correct for it or with what justification.

P(q/P(q)=x & PR(q/P(q)=x)=y) = ?

This asks for the degree of belief the subject should have in q on condition that she actually has degree of belief x in q and the objective probability of q given that she has degree of belief x in q is y. This expression is the left-hand side of the Synchronic Reflection Principle (van Fraassen 1984)

$P_t(q/P_t(q)=x) = x$

also known as Self-Respect (Christensen 20), with a further conjunct added to its condition. SR doesn't specify what to do when there is another conjunct and so is not suited to explicitly represent the question of self-doubt, which means that the self-doubting examples above are not counterexamples to it. (Roush 2009) However, some in the past have endorsed variants on an unrestricted version of SR (e.g., Koons 1992) where the value of this expression is x regardless of what other conjunct might be present:

P(q/P(q)=x & r) = x, for r any proposition Unrestricted Self-Respect (USR)[5]

Dutch book arguments that might give support to SR (though see Christensen 1991, Briggs 2009, Roush 2016) do not do the same for USR, leaving us with a need to find other ways of evaluating it when r is the statement of a calibration curve.

It is not incoherent, but it is baldly counterintuitive to suppose that the subject should have degree of belief x when she believes that her so believing is an indicator that the objective probability of q is not x, and a principled argument can also be made to this effect. (Roush 2009) Unpacking the condition P(q)=x & PR(q/P(q)=x) = y, it seems to say that my credence is x and when my credence is x the objective probability is y, inviting us to discharge and infer that the objective probability of q is y. If so,[6] then the expression reduces to:

P(q/PR(q)=y) = ?

which is the left-hand side of a generalization of the Principal Principle

---

[5] Obviously USR needs admissibility conditions to rule out counterexamples like taking r to be q itself. Since the calibration approach being discussed does not depend on USR this can be ignored in this context.

[6] P(q)=x and PR(q/P(q)=x)=y do not by themselves imply PR(q)=y. To discharge that conditional probability requires that PR(P(q)=x) is high, so one must ask under what conditions the credence being x makes the objective probability high that the credence is x. Under what conditions does A imply that A is objectively probable? This is an interesting question but it can be side-stepped in the derivation by appealing to a variant of the Principal Principle (Vranas 2004),

P(q/B & Ch(q/B)=y) = y Conditional Principle (CP)

to which no one has objected. The generalization of the Conditional Principle from chance to any objective type of probability is:

P(q/B & PR(q/B)=y) = y General Conditional Principle (GCP)

This says that the credence in q given that B is true and that the objective probability of q given that B is true is y, is y. Here too there are questions of admissibility, but as with PP it's easy to expect that there exists a useful domain in which CP and GCP are true.

Cal as derived in the text is an instance of GCP:

P(q/(P(q)=x & PR(q/P(q)=x)=y)) = y Cal

$P(q/Ch(q)=y) = y$ Principal Principle (PP)[7]

from chance to any type of objective probability. PP says that your credences in propositions should conform to what you take to be their chances of being true, and, admissibility issues notwithstanding, it is hard to deny that there exists a domain in which the Principal Principle is compelling, and surely one where the generalization to any type of objective probability is too. If so then the answer to the question what the subject's credence in q ought to be in light of her consideration of information about her reliability is:

$P(q/(P(q)=x \ \& \ PR(q/P(q)=x)=y)) = y$ Cal

Cal says that your credence in q given that your credence in q is x and the objective probability of q given that your credence in q is x is y, should be y.

Cal is a synchronic constraint, but if we revise our credences by conditionalization then it implies a diachronic constraint:

$P_{n+1}(q) = P_n(q/(P_n(q)=x \ \& \ PR(q/P_n(q)=x)=y)) = y$ Re-Cal

This calibration approach tells the subject how to respond to information about her cognitive impairment in every case. It uses the information about herself to correct her belief about the world. Intuitively it is a graded generalization of the thought that if you knew of someone (or yourself) that he invariably had false beliefs, then you could gain a true belief by negating everything he said.

Cal and Re-Cal give an explicit characterization of self-doubt and justification of a unique and determinate response to it on the basis of deeper principles that are compelling independently of the current context. Cal follows from only two assumptions, first that probabilistic coherence is a requirement of rationality, and second that rationality requires one's credences to align with what according to one's evidence are the objective probabilities. Re-Cal comes from further assuming that updating our beliefs should occur by conditionalization.

Although self-doubt under the current definition of it is not an incoherent state, Cal implies that rationality always requires a resolution of the doubt that brings matching between the two levels, and tells us that the matching consists in the alignment of subjective and perceived objective probabilities. High confidences in "q", "I have confidence x in q", and "the objective probability of q when I have confidence x in q is low" are not incoherent, but they do violate the Principal Principle. Re-Cal tells us how to get back in line with PP.

Though Re-Cal has us conditioning on second-order evidence, the adjustment it recommends depends on both first- and second-order evidence and does not always favor one level or the other. How much authority the second-order claim about the reliability/calibration curve has depends very much on the quality of one's evidence about it. This can be seen by imagining being uncertain about, e.g., one's calibration curve, i.e., $P(PR(q/P(q) = x) = y) < 1$, and doing a Jeffrey conditionalization version of Re-Cal, developed below. But even in case one has perfect knowledge of one's calibration curve, the role of the first-order evidence in determining one's revised first-order belief is ineliminable. The verdict, the level of confidence, that the first order gave you for q is the index for determining which point on the calibration curve is relevant to potentially correcting your degree of belief. To understand why this is not trivial, recall that the curve can in principle and does often in fact have different magnitudes and

---

[7] This formulation of PP suppresses a conjunct r of the kind that USR above has, meaning that it also requires admissibility conditions to identify the domain in which it is true. Those would be carried over in the use of PP in what follows here.

directions of distortion at different confidences. Your initial degree of belief in q has the role of an indicator of the objective probability of q, so the fact that it has one value or another is crucial.

The way that Re-Cal adjudicates between the first- and second-orders can be illustrated through the case of hypoxia. The information that she is at 10,500 feet combined with her knowledge of the phenomenon of hypoxia raise the objective probability that the pilot's calculations have been unreliable, that they've been of a sort that cannot be trusted to lead to true belief most of the time. Since she is aware of this she is forming an opinion about her calibration curve on its basis, but in this case the information is quite insufficient for confidence about what the curve is. The information is multiply inadequate because she does not learn that she *does* suffer from hypoxia or by how much the probability of that has ticked up, and the information doesn't identify what systematic error might be made, such as whether if she did err it would be by over- or underestimation. Indeed for all the information she has acquired, there may not even be a curve since that requires a regularity in which a particular confidence is a good indicator for a particular objective probability. The unreliability she may have may not be systematic but scattershot, so she doesn't have nearly enough information to know the probability of her conclusion about the fuel given that she has a particular confidence.

Forming a justified conclusion about even one region of a calibration curve requires a lot of information, as one should expect because that curve yields an objective probability for q, and thereby determines a unique particular revision of the first-order belief. Even without that, though, Bayesianism comes already equipped with tools for forming confidences on the basis of uncertain evidence, which work at the second order just as they do at the first order. A Jeffrey version of the Re-Cal conditionalization, a version which takes account of uncertainty about what your new information is by using a weighted average of its possible values, looks like this:

$P_f(q) = P_i(q/(B \& PR(q/B)=y))P_f(B \& PR(q/B)=y) + P_i(q/-(B \& PR(q/B)=y))P_f(-(B \& PR(q/B)=y))$

Jeffrey Re-Cal

where B stands for $P_i(q) = x$, i.e., the subject's initial degree of belief in q is x. In this equation the influence of new second-order evidence depends on how sure it is, in just the same way as evidence influences credence in Jeffrey conditionalization at the first order.

To illustrate how to process uncertain second-order information using the hypoxia case, we instantiate with F, the tank has enough fuel for fifty more miles, and x, the subject's initial confidence in F:

$P_f(F) = P_i(F/(B \& PR(F/B)=x))P_f(B \& PR(F/B)=x) + P_i(F/-(B \& PR(F/B)=x))P_f(-(B \& PR(F/B)=x))$

Jeffrey Re-Cal

where B stands for $P_i(q) = x$. For ease of exposition we will call the terms in this equation by these names:

$P_f(B \& PR(F/B)=x)$	Reliable
$P_f(-(B \& PR(F/B)=x))$	Off
$P_i(F/(B \& PR(F/B)=x))$	If-Reliable
$P_i(F/-(B \& PR(F/B)=x))$	If-Off

Your hypoxia information told you, let us say, to be 50% sure that you were a reliable calculator when coming to confidence x. (That is, let us say that half of the people at 10,500 feet get hypoxia.) This means you should be 50% confident that your answer, x, matches the objective probability and 50% sure that it doesn't. I.e., $P_f(PR(F/B)=x) = .5$ and $P_f(PR(q/B)\neq x) = .5$. Assuming, in all of these calculations, perfect knowledge of B, i.e., *that* your degree of belief is x, yields $P_f(B \& PR(F/B)=x) = .5$ and $P_f(B \& PR(F/B)\neq x) = .5$ That is, you should have half your confidence in the Reliable term and half in Off. The first term in the first summand of Jeffrey Re-Cal, $P_i(F/(B \& PR(F/B)=x))$, is settled by the argument above from the Principal Principle (or its close relative the Generalized Conditional Principle);[8] the confidence you put in If-Reliable should be x, your initial confidence. These together imply that the first summand in Jeffrey Re-Cal is (x)(.5).

To evaluate the second summand you need a value for If-Off, $P_i(q/-(B \& PR(q/B)=x))$. That is, you need not only a confidence that the objective probability of F isn't x when you have x confidence (Off), but the probability it has instead in that case. If you don't know your calibration curve then you don't know, but you can do a weighted average of the objective probabilities you regard as possible. If you think that if you are off then you're equally likely to be anywhere from near right to as far as possible off in either direction, then those probabilities could be anywhere from 0 to 1 (except for the value x), and the average will be .5, which makes If-Off .5. These points imply that the second summand is (.5)(.5) = .25. Thus if you had originally been 99% sure of F, you will now need to be (.99)(.5) + (.5)(.5) = 74.5% sure. Your confidence will be weighted towards your original confidence but incorporate a 50% confidence that it should be some other value, you have no idea which. 74.5% confidence in F might be low enough to justify the pilot staying on course if she has no positive reason for diverting, but that will depend on her utilities.

Holding your initial confidence, x, at .99 as your confidence that you are unreliable increases, i.e., as the Off term grows, the second summand will grow, but because your confidence was so high the first will shrink by more, and the overall revised degree of belief will be even lower. E.g., if Off is .7 instead of the .5 we assumed, then your final confidence will be 65% rather than 74.5%. If one's confidence that one is reliable increases, i.e., if Reliable grows, the second summand will shrink, but because of your high initial confidence the first summand will grow by more to give a revised confidence greater than 74.5%, though still less than the initial .99. If Reliable = .6 instead of .5, then the final confidence is 80% rather than 74.5%. If Reliable = 1 then you are certain you are calibrated, the second summand is zero, and your initial confidence stands.

If the initial confidence in q is lower than .99, then other things equal one's judgments of reliability have less effect. We saw a drop from .99 to .745 when we had the subject take the evidence as making it 50% sure she was unreliable. If she started at .7, then becoming 50% sure that she was unreliable would bring her confidence down by only 10 points to .6. An initial confidence of .5 is a fixed point in that no balance of confidences that one is and isn't reliable changes it away from .5.

All of these cases are on the assumption that If-Off has a uniform distribution. If you think that were you off you would be only near off then If-Off would be more concentrated around your initial confidence in q, and a given value for Off would not move your confidence as much. If-Off is also the place where you

[8] See fn. 6.

would register any information you had about the likely direction of your mistake had you made one. The only case this term makes no difference to is again an initial confidence of .5.

These points have interesting consequences for a case of the much-discussed phenomenon of higher-order evidence called peer disagreement, the situation where you find that a person you regard as an epistemic peer, i.e., with the same general reliability, has a different confidence than you do in the same proposition.[9] For example, it has been noted that if after both calculating your shares of a dinner bill your peer declares his confidence that they are $500 each, then either the restaurant is much fancier than it seems or you can be sure that his answer is way off, and intuitively this means you have little to no reason to take the mere disagreement as implying you should lower your confidence in your answer, despite the fact that you regard him as a peer.

You can apply Jeffrey Re-Cal to your dinner partner's confidence as you do to your own, using his confidence in the place of B, and letting q be the claim that your individual shares are $500. Your confidence in Off for his q and x should be high – the calibration curve depends on the proposition and it is independently obvious that this q is quite improbable regardless of anyone's confidence in it. Your confidence in If-Off for his q should be extremely low for the same reason. Regardless of your friend's general reliability on the mathematics of dinner bills, the objective probability of his answer in this case strongly indicates that in endorsing q in particular he's made a huge mistake, and that dominates. It makes the second summand dominate and its value extremely low and so the confidence it would be rational for you to have in q extremely low, pretty much just as low as it was before you heard his answer to the question. His confidence in q having virtually no effect on the estimated objective probability of his q means that the value of Off for your q and x does not change either, your first summand dominates, and you should remain at your original confidence in your original answer. This accords with our intuitions about the case, and explains it by the fact that no one's high confidence in this q makes much difference to what you should think this q's objective probability is.

Thus, we have found through these two cases that it is possible for Jeffrey Re-Cal to give us informative answers about what to do with evidence about our reliability even when we have little specific evidence about our calibration curve. Neither the first order nor the second is always dominant; both orders always make a contribution to determining the resolution at the first order of conflicts between orders, and their relative contribution depends on the quality of the evidence at each order. Jeffrey Re-Cal gives an explicit explanation of why the facts in the case of hypoxia mean you should revise and the facts of a particular case of peer disagreement mean you should remain steadfast, and it has consequences for every case of evidence about reliability, which could be explored in more detail. It makes available all of the resources of the Bayesian framework so those answers will be principled and based on a few assumptions that are independently plausible even if not universally shared, such as the Principal Principle.

References

Briggs, Rachael (2009), "Distorted Reflection," *Philosophical Review*, Vol. 118, No. 1, 59-85.

[9] The answer for this case also determines an answer for the case where you have equal confidence in different propositions.

Christensen, David (1991), “Clever Bookies and Coherent Beliefs”, *Philosophical Review*, 100: 229–47.

-------- (2007b), “Epistemic Self-Respect”, *Proceedings of the Aristotelian Society*, Vol. CVII, Part 3: 319-337.

------- (2010a), “Higher-Order Evidence”, *Philosophy and Phenomenological Research* 81(1): 185-215.

------- (2010b), “Rational Reflection”, *Philosophical Perspectives*, 24, *Epistemology*, 121-140.

Elga, A. (2013) “The Puzzle of the Unmarked Clock and the New Rational Reflection Principle”, *Philosophical Studies* (2013) 164: 127–139.

Gaifman, Haim (1986), “A Theory of Higher-Order Probabilities”, in *Causation, Chance and Credence*, Brian Skyrms and William Harper eds. Volume 41 of the The University of Western Ontario Series in Philosophy of Science. Springer, 191-219.

Kelly, Thomas (2005), “The Epistemic Significance of Disagreement”. In John Hawthorne and Tamar Szabo Gendler (eds.) *Oxford Studies in Epistemology*, volume 1 (Oxford University Press), pp. 167–196.

------- (2010), “Peer Disagreement and Higher-Order Evidence,” in R. Feldman and T. Warfield, eds., *Disagreement*. New York: Oxford University Press, 111-174.

Koons, Robert (1992), *Paradoxes of Belief and Strategic Rationality*. Oxford: Oxford University Press.

Lewis, David (1980), “A Subjectivist’s Guide to Objective Chance”, in his *Philosophical Papers*, Vol. 2. New York: Oxford University Press, 1986.

Roush, Sherrilyn (2009), “Second-Guessing: A Self-Help Manual”, *Episteme* 6(3): 251-68.

------- (2016), “Knowledge of Our Own Beliefs”, *Philosophy and Phenomenological Research*

------- (forthcoming), “Epistemic Self-Doubt”, *Stanford Encyclopedia of Philosophy* (expected 2017)

Seidenfeld, Teddy (1985), “Calibration, Coherence, and Scoring Rules” *Philosophy of Science* 52: 274-294.

van Fraassen, Bas C. (1984), “Belief and the Will”, *Journal of Philosophy* 81, 235–56.

Vranas, Peter B.M. (2004), “Have your cake and eat it too: The Old Principal Principle Reconciled with the New”, *Philosophy and Phenomenological Research* Vol. LXIX, No. 2: