

Knowing full well: the normativity of beliefs as performances*

ERNEST SOSA

TWO OF PLATO'S BEST-KNOWN DIALOGUES are inquiries about knowledge. The *Theaetetus* inquires into its nature, the *Meno* also into its value. Each dialogue, I will suggest, involves the same more basic question: What sort of normativity is constitutive of our knowledge? A belief that falls short of knowledge is thereby inferior. It is better to know than to get it wrong, of course, and also better than to get it right just by luck. What is involved in such evaluation? An answer to this more basic question enables a solution for both Platonic problems. In this paper we consider mainly this question: What is the epistemic normativity that is constitutive of knowledge?

We shall assume that knowledge requires at a minimum a belief that is true. Our inquiry into the nature of knowledge thus takes a more specific form. Our question is this: What condition must a belief satisfy, in addition to being true, in order to constitute knowledge? The question of the nature of knowledge has been central to epistemology in recent decades, as it was for Plato.

Edmund Gettier showed that more is required for a belief to constitute knowledge, beyond its being competently held: that is to say, competently acquired or sustained.¹ For one thing, a belief can be false despite being competent. If the believer then competently deduces something true from his false belief, this true conclusion cannot thereby amount to knowledge. Yet, if we competently deduce a conclusion from a premise that we competently believe (even after drawing the conclusion), we thereby competently believe that conclusion as well. So a belief can be both true and competently held without amounting to knowledge.

* This paper keynoted the first Midwest Epistemology Workshop, and was also delivered in Taiwan as the first of three Soochow Lectures in Philosophy.

¹ This is the celebrated Gettier problem, with a vast literature.

Post-Gettier, the Platonic problem takes this form: What further condition, added to, or in place of, *being competently held*, must a true belief satisfy in order to constitute knowledge?

On the contemporary scene, the second Platonic problem, that of the *value* of knowledge, has more recently moved to center stage. For Plato this was the problem of how knowledge can be quite generally more valuable than its corresponding true belief, if a merely true belief would be no less useful. Thus, a true belief as to the location of Larissa will guide you there no less efficiently than would the corresponding knowledge. In line with this, we ask: How if at all can knowledge be *as such* always better than the corresponding merely true belief?

In connection with both problems, we will assume that there is some further condition (however simple or complex) that a belief must satisfy in order to constitute knowledge, beyond being a belief and being true. This condition must add normatively positive content, moreover, sufficient to explain how it is that *knowledge*, which must satisfy this further condition, *is as such always better than would be the corresponding merely true belief*. When one ponders a question, for example, there is some respect in which it would always be better to answer knowledgeably than to answer correctly but just by luck.

Let's begin with the *Meno* problem

§1. The value of knowledge

The aim of belief is said to be truth. When you sincerely pose a question to yourself, for example, you want a correct answer. When you reach an answer to your question through adopting a certain belief, the aim of your belief is the truth of the matter. If the aim of a belief is thus truth, then once true that belief would seem to have what really matters epistemically, irrespective of its aetiology.

How then can a truth-reliably produced true belief be better than one that is no less true, regardless of how reliably it may have been produced? Conclusion: knowledge is really no better as such than merely true belief.

«Any argument leading to that conclusion», it may well be replied, «must have its premises examined: perhaps the aim of belief, and of inquiry, is not just truth, but also knowledge. This would explain how and why it is that knowledge (with its required aetiology) is after all better than merely true belief».

What follows will defend this reply by placing it in context, by explaining its content, and by drawing some implications.

§2. A step back

How indeed *is* truth our aim? How should we understand the value we place on it? More explicitly, our aim is presumably to *have* the truth. So, it is the *attained* truth that has corresponding value. How then should we more fully describe our true objective? Is it just the accumulation of believed truths? Compare how we assess accurate shots, those that hit their targets. What is it that people value under this rubric? Is it the accumulation of accurate shots?

Someone casually draws a large circle on the beach right by his feet, aims his gun, and hits the target. Does he thereby attain, at least in some small part, a previously standing objective: namely, that of securing accurate shots? Is that an objective we all share, given how we all share the concept of a *good* shot? Don't we all want good things (other things equal)?

That, I trust we agree, is quite absurd.

Yet, the shot at the beach could be an accurate, good shot nonetheless, as the marksman hits his target in the sand. Although, from one point of view, given the low or even negative value of the target, this accurate shot has little value of its own, yet from another, performance-internal perspective it is graded as quite accurate, a good shot, maybe even an excellent shot if the marksman steps back far enough from the target. Even when the shot is difficult, however, its status does not derive from any standing preference of people for an accumulation of accurate, difficult shots. There is no normative pressure on us to bring about good shots, not even if we grasp perfectly well what it takes to be a good shot, and have this uppermost in our consciousness at the time. There is no normative pressure to bring about even excellent shots, none whatever that I can discern. (N.B.: What we are not normatively pressured to accumulate is *shots, nor even excellent shots, regardless of how excellent they may be*).

Compare now our intellectual shots, our beliefs. A belief may answer a question correctly, but may have very little value nonetheless, if the question is not worth asking. The value of a target will surely bear on the worth of any shot aimed at that target. Arbitrary selection of an area by your feet at the beach yields a silly target. Similarly, suppose you scoop up some sand and proceed laboriously to count the grains. You then take up the question of how many

grains are contained in that quantity of sand. If you reach your objective of answering that question correctly, what is your performance worth? Do you thereby fulfill, at least in some small part, a previously standing objective, that of securing more and more true beliefs? This seems about as implausible as is the corresponding view about the shot at the beach.

In what way, then, *does* the truth of our beliefs have value? One thing that does plausibly have prima facie value is the satisfaction of our curiosity. So, even if the question as to the number of grains is of little worth, if someone gets interested in that question anyhow, then the satisfaction of his curiosity will in an obvious way have value *to him* (and perhaps even, to some small extent, *for him*). This is of course a way for the truth to have value to someone and for someone. After all, if one is curious as to whether *p*, this is just to be curious as to whether it is true that *p*. There are not two instances of curiosity here: (a) as to whether *p*, and (b) as to whether it is true that *p*. So, what we want when we value the truth *in that way* is to have our questions answered, and of course answered correctly.

Sheer curiosity, whatever its basis, thus invests the right answer to a question with some value, though the value might be small and easy to outweigh, as with the question about the grains of sand. Having the answer to that particular question may add so little to the life of the believer, while cluttering his mind, that it is in fact a detriment all things considered, if only through the opportunity cost of misdirected attention.

Similar considerations apply to the shot aimed from a foot away at the sand on the beach. The sheer desire to hit that target, whatever its basis, gives value to the agent's hitting the mark. But it might well be that hitting that mark imports little value for anyone. Spending his time that way may even be a detriment to the agent's life. Nor is it plausible that we humans have generally a standing desire for accurate shots, nor that we place antecedent value on securing such shots. Accuracy will give value to that shot at the sand only dependently on the gunman's whim to hit that target.

Even if that shot at the beach fulfils no human interest antecedent to the gunman's whim, it may still be a better shot, better as a shot, than many with higher overall value. Take a shot at close quarters in self-defense that misses the targeted head of the attacker but hits him in the shoulder and stops the attack. A bad, inaccurate shot, this one, but more valuable than the accurate shot at the beach. (Had it been better as a shot, moreover, a more accurate shot, it might

have constituted a terrible murder, since the attack did not justify shooting to kill).

Are beliefs like shots in that respect? Is a belief a performance that can attain its internal aim while leaving it open whether it has any intrinsic value, and whether it serves or disserves any external aim? Let us explore this view of belief.

§3. Knowledge as a special case

All sorts of things can «perform» well or ill when put to the test. Rational agents can do so, but so can biological organs, designed instruments, and even structures with a function, such as a bridge. A bridge can perform well its function as part of a traffic artery. When a thermostat activates a furnace, it may perform well in keeping the ambient temperature comfortable. When a heart beats, it may perform well in helping the blood circulate. And so on.

A puppet performs well under the control of a puppeteer if its hinges are smooth, not rusty, and well oiled, so that its limbs are smoothly responsive. A bridge might perform well by withstanding a storm. We credit the puppet, as we do the bridge, if its good performance flows appropriately from its state and constitution.

The puppet «performs» (well or ill), as does the bridge, and thus produces performances. But it would be quite a stretch to consider it an «agent». Human beings are different, in any case, if only because we are *rational* agents. Not only are there reasons why we perform as we do. There are also reasons that we *have* for so performing, and *for* which, *motivated* by which, we perform as we do. This is not just a matter of having aims in so performing. After all, the thermostat and the heart do have their aims. But they are motivated by no such aim; no such aim gives them reasons motivated by which they perform as they do.²

Human motivation is on another level, even when the performance is physical, as in athletic or artistic performance.

² True, we could perhaps, just barely, make sense of an extended sort of «motivation» even in those cases, as when a nearby torch fools the thermostat into activating the air conditioner even when the room is already cool. It still in some broad sense has a reason for performing as it does, a «motivating reason». Despite the non-trivial resemblance, nonetheless, this is clearly a metaphorical extension, if only because of the vastly greater complexity involved in human motivation. In any case, a thermostat does not literally have a mind, or any motives.

The archer's shot is a good example. The shot aims to hit the target, and its success can be judged by whether it does so or not, by its accuracy. However accurate it may be, there is a further dimension of evaluation: namely, how skillful a shot it is, how much skill it manifests, how adroit it is. A shot might hit the bull's-eye, however, and might even manifest great skill, while failing utterly, *as a shot*, on a further dimension. Consider a shot diverted by a gust of wind initially, so that it would miss the target altogether but for a second gust that compensates and puts it back on track to hit the bull's-eye. This shot is both accurate and adroit, yet it is not accurate *because* adroit, so as to manifest the archer's skill and competence. It thus fails on a third dimension of evaluation, besides those of accuracy and adroitness: it fails to be *apt*.

Performances generally admit this threefold distinction: accuracy, adroitness, aptness. At least so do performances with an aim (assuming any performance could ever be wholly aimless).

A performance is better than otherwise for not having *failed*, i.e., for not having fallen short of its objective. In line with that, it is *good* if it succeeds, if it reaches its objective. A performance is at least good *as such* for succeeding, even if it is a murderer's shot. The shot itself may still be an excellent shot, despite how deplorable is the broader performance in which it is embedded.

A performance that attains its first-order aim without thereby manifesting any competence of the performer's is a lesser performance. The wind-aided shot scores by luck, without thereby manifesting appropriate competence. It is hence a lesser shot by comparison with one that in hitting the mark manifests the archer's competence.³ A blazing tennis ace is a lesser shot if it is a wild exception from the racket of a hacker, by comparison with one that manifests superb competence by a champion in control. And so on. Take any performance with a first-order aim, such as the archery shot and the tennis serve. That performance then has the induced aim of *attaining* its first-order aim. A performance X attains its aim $\langle p \rangle$, finally, not just through the fact that p , but through the fact that it *brings it about* that p .⁴

³ A shot might manifest an archer's competence without its accuracy doing so. The shot with the two intervening gusts is a case in point. How does that shot manifest the archer's competence? By having at the moment of release an angle, direction, and speed that would take it to the bull's-eye, in relevantly normal conditions.

⁴ Just as its being true that p entails its being true that it is true that p , so one's bringing it about that p may entail that one brings it about that one brings it about that p , assuming such iteration always makes sense.

The case of knowledge is just the special case where the performance is cognitive or doxastic. Belief aims at truth, and is accurate or correct if true. Belief has accordingly the induced aim of *attaining* that objective. Belief aims therefore not just at accuracy (truth), but also at aptness (knowledge). A belief that attains both aims, that of truth and that of knowledge, is for that reason better than one that attains merely the first. That then is a way in which knowledge is as such better than merely true belief.⁵

The account of epistemic normativity as a sort of performance normativity has thus two virtues. It provides an explanation of the nature of knowledge, which amounts to belief that is apt, belief that is an apt epistemic performance, one that manifests the relevant competence of the believer in attaining the truth. And, secondly, it explains also the extra value of knowledge beyond that of merely true belief.

Unfortunately, the account encounters a troubling objection, which we next consider.

§4. The problem of withholding

What's the problem?

The normative judgment that knowledge is as such better than merely true belief is of a piece with the normative judgment that withholding is better than believing when the evidence is insufficient. Since both judgments are epistemically normative, one would expect them to be closely akin. But that is not what one finds on first inspection.

If truth is the first-order aim of our cognitive endeavors, it is not obvious how to assess suspension of judgment with respect to that objective, so it is correspondingly hard to see how we could apply our AAA normative structure (accuracy, adroitness, aptness) of performances to such withholdings. These are after all precisely *non-*performances. How then can they be brought within the sphere of our performance normativity? And if

⁵ Even if performances do not have the automatically induced aims just suggested, we still retain an account of why knowledge is better than merely true belief, since apt performances, in general, are as such better than those that attain success only by luck. So, beliefs provide just a special case of that general truth. This account still depends of course on our view of knowledge as apt belief, belief that manifests the relevant competence of the believer in reaching its aim of truth.

they are not thus assimilable, serious doubt is cast on our claim to have uncovered the most relevant epistemic normativity involved in our intuition that knowledge is as such better than merely true belief.

Let our archer now be a hunter rather than a competitor athlete. Once it is his turn, the competitor must shoot, with no relevant choice. True, he might have avoided the competition altogether, but once in it, no relevant shot selection is allowed. The hunter by contrast needs to pick his shots, with whatever skill and care he can muster. Selecting targets of appropriate value is integral to hunting, and he must also pick his shots so as to secure a reasonable chance of success. The shot of a hunter can therefore be assessed in more respects than that of a competitor/athlete. The hunter's shot can be assessed twice over for what is manifest in it: not only in respect of its execution competence, but also in respect of the competence manifest in the target's selection and in the pick of the shot.

Not taking a shot at a particular target may or may not involve a performance. You might *fail to take that shot* because at the time you are asleep, for example. Alternatively, you might intentionally and even deliberately forbear. If your deliberate forbearing has an aim, moreover, and if the aim is attained, then your forbearing succeeds, and may even be a performance, indeed one that is apt.

Suppose a domain in which an agent puts in performances with an aim, whether athletic, artistic, academic, etc. This yields a derivative aim: *to avoid failure*. You can aim to avoid failure, moreover, without aiming to attain success, at least not ground-level success. When a hunter decides not to take a shot at a certain high-value target, for example, his performance, his forbearing, has its own aim of avoiding failure. To forbear is precisely *not* to aim at first-order success. Nevertheless, forbearing has an aim of its own: namely, avoiding failure.

Take then a hunter's performance of forbearing, which succeeds in avoiding ground-level failure. It does attain *that* aim. Is it thereby apt? Yes, so it is by our account; that is what we have to say. The forbearing *is*, after all, a performance with an aim of its own, and it does attain that aim, in doing which it does manifest a sort of competence.

What if it is a shot that the hunter very obviously *should* have taken? What if he makes a big mistake forbearing?

How do we avoid the unwelcome result that the forbearing is apt despite being one that obviously should not even have occurred? One option is to grant that it is a *narrowly* apt performance, while defining a broader aptness that it lacks. Let us explore this option.

Consider Diana's forced choice between taking a shot and forbearing from doing so. If she opts to take the shot, then her archery skills come into play. If they produce a hit, then her performance, her shot, manifests her narrow competence, and is hence narrowly apt. Compatibly with this, nonetheless, her shot selection might have been incompetent.

That is one way for a narrowly apt shot to be broadly objectionable. The huntress who forbears taking a shot that she obviously *should* take fails in her performance of forbearing. Her forbearing avoids ground-level failure, but is deplorable nonetheless because it fails properly to manage relevant risk.⁶

§5. Varieties of aptness

A performance is apt if its success manifests a competence seated in the agent (in relevantly appropriate conditions). It does not matter how fragile was the continued presence of the competence, or its appropriate conditions, when the agent issued the performance. A performance can thus easily fail to be «meta-apt», because the agent handles risk poorly, either by taking too much or by taking too little. The agent may fail to perceive the risk, when he should be more perceptive; or he may respond to the perceived risk with either foolhardiness or cowardice. He might perform on the ground level although the risk of failure is too high; or he might forbear although it is pusillanimous of him not to plunge ahead.

The aptness of a performance is thus to be distinguished from its meta-aptness. Either one can be present without the other.

An archer/hunter's shot selection and risk taking may be excellent, for example, and in taking a certain shot he may manifest his competence at assessing risk, while the shot itself nevertheless fails, being unsuccessful (inaccurate) and hence inapt. The shot is hence meta-apt without being apt.

⁶ Here I mean to stay within the domain of hunting, where we assess shots as good hunting shots. Impressing a girlfriend, and bonding with a rich uncle, are objectives irrelevant to the assessment of a shot as a good hunting shot, though of course they remain relevant to the assessment of the shot in other ways.

Conversely, the hunter may take excessive risk in shooting at a certain target, given his perceived level of competence (he has been drinking) and the assessed potential for wind (it is stormy). When he shoots, he may still fall just below the level of competence-denying inebriation, however, and the wind may happen to fall calm, so that his shot is (through *that* stroke of luck) quite apt. Here the shot is apt without being meta-apt.

Our shift from the competitor archer to the hunter archer, with his much wider latitude for target or shot selection, imports therefore the following distinction:

A shot is *apt* iff the success it attains, its hitting the target, manifests the agent's first-order competence, his skillful marksmanship.

A shot is *meta-apt* iff it is well-selected: i.e., iff it takes appropriate risk, and its doing so manifests the agent's competence for target and shot selection.

Neither aptness nor meta-aptness is sufficient for the other. They vary independently.

If Diana shoots, her shot might itself be both apt and meta-apt. If she forbears, her forbearing might be meta-apt, though of course it will not be apt on the ground level, since it does not even aim for success on that level. The forbearing might be meta-apt, nevertheless, in being a proper response to the perceived level of risk, a response that manifests her meta-competence.

Sometimes an agent responds properly by performing on the ground level, in which case that *positive performance* is meta-apt; sometimes the proper response is to forbear, so that the *forbearing* is meta-apt.

Arguably, a shot could be both apt and meta-apt while still falling short in that it is not *in virtue of being meta-apt* that it is apt. Thus, a shot might manifest a hunter's risk-assessment competence, and it might issue from his competence as an archer, in conditions appropriate for such shots, while yet the shot might be apt, not through the meta-competence of the archer, but only through a kind of luck. Diana might assess risk aptly and then just toss a coin to decide whether to shoot. But, a fully apt performance is not only apt but also meta-apt. To know something full well requires one's constitutive belief to be fully apt.⁷

⁷ And, indeed, apt, at least in part, *because* it is meta-apt by *manifesting* its meta-aptness.

§6. Full aptness and reflective knowledge

A performance attains thus a special status when it is apt at the ground level and also its aptness is explained through competent risk assessment. Suppose this risk-assessment issues in the performer's knowing that his situation (constitutional and circumstantial) is favorable (where the risk of failure is low enough) for issuing such a performance. If these conditions all obtain, then the performance's aptness might stem from its meta-aptness; that is to say, its aptness might be relevantly explicable through the performer's meta-knowledge that his first-order performance is likely enough to succeed and be apt.

This applies to performances such as a shot that hits its prey. That shot is superior, more admirable and creditable, if it is not only apt, but also meta-apt, and, further, *fully* apt: that is, apt because meta-apt. This happens, for example, when the aptness of Diana's shot stems from her meta-competence in assessing risk properly, so that the shot's aptness manifests her competence for taking apt shots, a competence that essentially includes her ability to assess risk well.

Aptness comes in degrees. One shot is more apt than another, for example, if it manifests a more reliable competence. On one dimension, a shot by a tennis champion may be no better than a similarly paced and placed shot by a hacker. On another dimension, however, the champion's shot manifests his prowess on the court, while the hacker's nearly identical shot is just lucky, and skillful only minimally or not at all. The champion's shot manifests competence, moreover, on two levels. It manifests his sheer athletic ability to hit with good pace and placement, and with impressively good percentage. But it can and normally does manifest also her good shot selection, her ability to attempt shots with a favorable percentage of success. The hacker's shot falls short on both dimensions.

The champion's shots are apt, meta-apt, and *fully* apt (i.e., apt in a way that manifests meta-aptness). For a shot to have the property of being apt is for its success to manifest a competence seated in the agent. This whole arrangement is itself something that the agent might be able to arrange (or not), and not simply by exercising the first-order competence seated in him. The agent might be able to choose when and where to exercise that competence, for one thing, and might manifest more or less competence in such a choice.

The same is true of the archer/hunter's shot. It can be apt in that its success, its accuracy, manifests the agent's competence in relevantly appropriate

conditions (no wind, enough light, distance within proper bounds, and so on). But it can also manifest the agent's meta-competence for target and shot selection. If so, then it is no accident that the shot is made in specific conditions where the archer's competence is up to the task of producing success with a high enough percentage. In other words, the agent's risk perception is then competent enough, and this competence is manifest in his knowledge that the level of risk is appropriate. On one level, how apt the shot is depends on the degree of competence manifest by its success. But, on another level, the full aptness of the shot depends also on the meta-competence manifest by its success. A performance is *fully* apt only if its first-order aptness derives sufficiently from the agent's assessment, albeit implicit, of the chances of such success (and, correlatively, of the risk of failure).

Here the agent is on a meta-level. He must take into account the likelihood that his competence is (and will remain) intact and that the relevant conditions are (and will remain) appropriate, and he must assess how likely it is that his action from such a competence in such conditions will succeed. Suppose he takes his chances of such success to be high enough (and the risk of failure low enough), and he is right, knowledgeably so, the chances being as he takes them to be, and his competence and conditions being relevantly as envisaged. Suppose further that he exercises his competence accordingly, so that the (first-order) aptness of his shot is owed to his meta-competence, is owed sufficiently to his getting it right about the chances of success, and to his getting *this* right as a manifestation of that meta-competence. That shot is then more fully apt and more fully creditable in proportion to how fully all of that falls into place.

We have thus found a further level of *performance-based* normativity. *Epistemic* normativity is, once again, a special case also in this more complex and subtle way. Animal knowledge is first-order apt belief. Reflective knowledge is animal belief aptly meta-endorsed by the subject. We can now see that knowing something full well requires that one have animal and reflective knowledge of it, but also that one know it with full aptness. It requires, that is to say, that the correctness of one's first-order belief manifest not only the animal, first-order competences that reliably enough yield the correctness of the beliefs that they produce. One's first-order belief falls short if it is not appropriately *guided* by one's relevant meta-competence. This meta-competence governs whether or not one should form a belief at all on the question at issue, or should rather withhold. It is only if this meta-competence is operative in one's forming a belief at all on that subject matter that one's belief can reach the epistemic

heights. One's first-order belief is apt in proportion to how reliable is the first-order competence manifest in its success. What is more, it is more fully apt in proportion to how reliable is the meta-competence that its success also manifests. This meta-competence is manifest at a remove, however, because the meta-knowledge *that it is a belief likely enough to be apt on the ground level* is constituted by the fact that the correctness of the corresponding meta-belief itself manifests the subject's relevant meta-competence.

Fully apt performances are in general better as performances than those that succeed without being apt at all, and also than those that are apt without being fully apt. Diana's apt shot that kills its prey is a better shot if apt than if successful only by luck and not through competence. Moreover, it is also a better, more admirable, more creditable shot, if its success flows also from her target-selecting, shot-picking competences⁸. Her shot is more creditable in that case than it is when the right competence *is* manifest in conditions required for a successful first-order performance, *but only by luck external* to any such selection meta-competence on her part.

Epistemic normativity is again just a special case of all that. Apt belief, *animal knowledge*, is better than belief that succeeds in its aim, being true, without being apt. Apt belief aptly noted, *reflective knowledge*, is better than mere apt belief or animal knowledge, especially when the reflective knowledge helps to guide the first-order belief so that it is apt.⁹ In such a case the belief is fully apt, and the subject *knows full well*.

⁸ It might be thought that one needs to know in some detail how the faculty works, if one is properly to be credited for its successful manifestations. But there is surely a kind of «credit» that is in place even for the unreflective subject, if only a sort *similar* to the credit attributable to a thermostat for keeping the room warm.

⁹ In fact proper reflective knowledge will *always* guide or help to guide its corresponding animal belief. Proper reflective knowledge will after all satisfy requirements of coherence, which means not just logical or probabilistic coherence of the respective belief contents, but also the mutual basing relations that can properly reflect such coherence among the contents. Cross-level coherence, from the object to the meta, and conversely, is a special case of such coherence, and it imports «guidance» of the animal belief by the relevant meta-beliefs (or, in other words, basing of the former on the latter). It bears emphasis that the meta-aptness of a belief, which we have found to be an important factor in its *epistemic* evaluation, requires ascent to a good enough perspective concerning the first level potential attitudes among which the subject must opt (whether he opts with full conscious deliberation or through a less explicit procedure). Coherence among first-level attitudes is not enough. The subject must ascend to a level wherein he assesses relevant risk, whether in full consciousness or less explicitly, and opts on that basis. Included in that analysis is perforce some assessment of one's relevant competence(s) and situation, and this must itself be performed adequately, if it is to yield a fully creditable first-level performance. Its assessment as thus fully creditable is moreover fully epistemic. For it is an assessment based on epistemic standards as to whether belief, rather than suspension of belief, is the proper response to one's situation.

Recibido: 19-Mayo-2015 | Aceptado: 4-Agosto-2015



ERNEST SOSA, es Profesor de Filosofía en la Rutgers University, EUA. Doctor en Filosofía [PhD] por la University of Pittsburgh. Sus principales áreas de interés son la epistemología, la metafísica la filosofía de la mente y la epistemología moral. Entre sus principales publicaciones se cuentan: *A Virtue Epistemology. Apt Belief and Reflective Knowledge*. Vol. I (Oxford: Clarendon Press, 2007), *Reflective Knowledge. Apt Belief and Reflective Knowledge*. Vol. II (Oxford: Clarendon Press, 2009), *Knowing Full Well* (Princeton-Oxford: Princeton University Press, 2011); *Judgment and Agency* (Oxford: Oxford University Press, 2015). Es editor de *Philosophy and Phenomenological Research* y *Noûs*.

DIRECCIÓN POSTAL: Department of Philosophy, Rutgers University, 106 Somerset Street - 525, New Brunswick, NJ 08901, USA. e-mail (✉): ernestsosa@mac.com

COMO CITAR ESTE TRABAJO: SOSA, Ernest. «Knowing full well: the normativity of beliefs as performances». *Disputatio. Philosophical Research Bulletin* 4:5 (2015): pp. 81-94.

© El autor(es) 2015. Este trabajo es un [Artículo. Original], publicado por *Disputatio. Philosophical Research Bulletin* (ISSN: 2254-0601), con permiso del autor y bajo una licencia Creative Commons [BY-NC-ND], por tanto Vd. puede copiar, distribuir y comunicar públicamente este artículo. No obstante, debe tener en cuenta lo prescrito en la *nota de copyright*. Permisos, preguntas, sugerencias y comentarios, dirigirse a este correo electrónico: (✉) boletin@disputatio.eu

Disputatio se distribuye internacionalmente a través del sistema de gestión documental GREDOS de la Universidad de Salamanca. Todos sus documentos están en acceso abierto de manera gratuita. Acepta trabajos en español, inglés y portugués. Salamanca – Madrid. Web site: (✉) www.disputatio.eu