Robust Virtue Epistemology and the Ontology of Complete Competences[1]

Modesto Gómez Alonso[2]

Abstract. In *Judgment and Agency*, Ernest Sosa argues for a triple-S structure of complete competences that includes, besides the innermost seat competence of the agent, her overall intrinsic condition (shape) and the right situational factors for the manifestation of cognitive success to occur. Complete competences are context-sensitive. The question is raised whether epistemic competences are extrinsic or intrinsic dispositional properties, as well as whether knowledge is the manifestation of powers of the actual world or whether it is a matter of what happens in modally close worlds. An ontological background for the context-sensitivity of epistemic competences that is compatible with their intrinsic character is set in the context of a general account of dispositions. An actualist conception of knowledge is also provided. On this view, the modal force of knowledge, far from being captured by non-localized possibilities, is rooted in the directedness of powers. This article, thus, aims at making explicit the ontological underpinnings of knowledge in a way that it is consistent with Sosa's fully dispositional virtue epistemology. To this end, the constitution of complete competences is placed in a proper perspective.
Keywords: Ernest Sosa; anti-luck virtue epistemology; complete competence; epistemic dependence; epistemic dispositions; knowledge; performance normativity; robust virtue epistemology.

[es] La epistemología de virtudes robusta y la ontología de las competencias completas

Resumen. En *Judgment and Agency*, Ernest Sosa argumenta que, para ser completas, las competencias cognitivas han de poseer una estructura tripartita que incluye el asiento o portador de la competencia, la condición general intrínseca del agente, y los factores situacionales que permiten la manifestación de la competencia como éxito epistémico. Dicho análisis plantea varias cuestiones ontológicas, en concreto, acerca de si las competencias epistémicas son (o no) propiedades disposicionales extrínsecas, y en referencia a la explicación de la fuerza modal del conocimiento: bien en términos de potencialidades actuales o de mundos posibles. Los objetivos de este artículo son (i) proporcionar un encuadre ontológico compatible tanto con la sensibilidad contextual de las competencias epistémicas como

con su carácter intrínseco; y (ii) defender una concepción actualista del conocimiento tal que éste se encuentre enraizado en la direccionalidad de las virtudes y no incluya en su análisis posibilidades no-localizadas. Se trata, por tanto, de arrojar luz sobre la constitución de las competencias completas de modo que su análisis sea compatible con la epistemología de virtudes robusta propuesta por Sosa.

Palabras clave Ernest Sosa; epistemología de virtudes anti-suerte; competencia completa; dependencia epistémica; disposiciones epistémicas; conocimiento; normatividad de la actuación; epistemología de virtudes robusta.

Sumario: 1. Introduction; 2. Outsourcing Modality; 3. Costs and benefits; 4 Complete competences; 5. References.

1. Introduction

In *Judgment and Agency*, Sosa (2015) develops his account of knowledge as (broadly) cognitive success due to competence into a rather thorough analysis of how competences, understood as dispositions or abilities to succeed, are constituted. As Sosa sees it[3], cognitive competences have –when complete– a triple-S structure that includes the innermost (seat) competence of the agent, her overall intrinsic state (shape), and the external circumstances for the manifestation of cognitive success to occur (situation). It is also important to notice that a range of *normal* circumstances set a threshold below which S would not count as relevantly competent[4]. Thus, Sosa's analysis is able to capture finer structures of potentiality. Those structures are *context-sensitive*[5]. As thus, they take into account both temporal dispositions that, as inebriation, would mask the agent's more stable innermost ability, and external factors relevant to provide the right opportunity to exercise it.

Context-sensitivity is instrumental in making sense of how the attribution of competence to the same agent in the same situation (one that is unpropitious) may with equal right be held true or false. If we are focusing on the innermost competence (namely, on a basic power to F such as S's basic driving competence), the attribution may be true independently of how things stand in the vicinity of the agent. If, however, we also take into consideration her overall shape (such as her being sleepy and slightly inebriated while driving) and the availability of external conditions suitable for the successful exercise of the competence (if S were driving under hurricane conditions, those situational factors would not be available), the agent might count as having the innermost competence while lacking the relevant, fine-grained, complete competence.[4]

Context-sensitivity, however, needs to be managed with care. Firstly, it may well suggest an ontological distinction between *intrinsic* and *extrinsic* dispositions so

---

[3]   *Cf.* Sosa, E.: *Judgment and Agency*, Oxford, Oxford University Press, 2015, pp. 95-96. [Henceforth J&A].

[4]   The range of *normal* circumstances set a *minimal* threshold for attributing a competence to S. One is not expected to drive in hurricane conditions to count as a good driver, as one is not expected to be infallible (to get knowledge in *all* worlds) to count as epistemically competent.

[5]   Although, as we shall see, this way of describing competences needs of refinement or, better said, it needs to be relativized to the *exercise*, and not to the possession, of the complete competence.

that complete competences would be classified as being part of the latter variety – properties that the subject acquires and loses without undergoing any intrinsic change[6]. This claim would be at odds with one of the paradigm marks of dispositionality: *intrinsicality*. Dispositions being intrinsic properties of their bearers, they are only a matter of how things stand with the cognitive powers of the agent[7].

Second, context-sensitivity could easily be associated with the *success thesis*, namely with the claim that, since S exercises an ability to F only if the action to which the ability is directed is successfully performed, S would lack an ability to F[8] if the situation or shape were not suitable for cognitive success due to the agent's relevant competence[9]. One problem with the success thesis is that it flies in the face of common sense intuitions that seem to conceive abilities (whether general or fine-grained) as *portable properties* of agents and objects. In this sense, it seems, at the very least, strange, to say that a glass wrapped in styrofoam lacks the power *to break when stressed in the right circumstances* (a power that, by all means, would count as a fine-grained disposition). A second source of trouble comes from the single manifestation-type conception of abilities to which this view seems committed. This conception and the 'hub-and-spokes' model that many power theorists[10] propose do not go together well.

Finally, if we shift from competences to successful events such as knowledge, the context-sensitivity of knowledge[11] could be developed into an account that conceives the modality inherent to knowledge as the modality of subjunctive conditionals and easy possibilities, and that, thus sees knowledge as *not being fully anchored* in the cognitive powers of the subject. Supporters of anti-luck virtue epistemology (ALVE,

---

[6]   Jennifer McKitrick has rejected the claim that all dispositions are intrinsic, providing several examples of prospective extrinsic properties: the power of a key to open a *particular* lock, vulnerability, the disposition of water to dissolve the contents of one's pocket when the pocket contains sugar cubes, etc. However, her liberal account of properties, which includes mere Cambridge-properties, makes her proposal hostage to highly questionable commitments. See McKitrick, J.: "A Case for Extrinsic Dispositions", *Australasian Journal of Philosophy*, 81(2), 2003, pp. 155-174.

[7]   For a compelling defence of this view, see Molnar, G.: *Powers: A Study in Metaphysics*, Oxford, Clarendon Press, 2003, pp. 102-110.

[8]   Here 'ability' is shorthand for 'fine-grained ability'. Millar distinguishes between 'ability *simpliciter*' (a general competence that the subject retains even when the situation prevents it from manifesting) and 'fine-grained ability'. It seems as if in his view the success thesis (and, therefore, context-sensitivity) would only apply to the last ability-type. *Cf.* Millar, A.: "Abilities, Competences, and Fallibility", in M.A. Fernández Vargas (ed.), *Performance Epistemology: Foundations and Applications*, Oxford, Oxford University Press, 2016, p. 79. [Henceforth ACP].

[9]   The *success thesis*, which seems to go back at least to Moore, has recently be endorsed by the 'New Dispositionalism' of Vihvelin (2004) and Fara (2008). For a vigorous defence of this claim in the context of epistemology, see ACP. As Millar makes clear, there is a logical connection between the success thesis and the claim that "abilities are abilities with respect to favourable environments." (ACP, p. 64) I will endorse a very different thesis, and that is one in which abilities are pragmatically *selected* with respect to favourable environments. However, pragmatic selection should not be projected to ontology.

[10] For a detailed argument on multi-track dispositions, see Martin, C.B.: *The Mind in Nature*, Oxford, Clarendon Press, 2008, pp. 31-33; and Heil, J.: *The Universe as We Find It*, Oxford, Clarendon Press, 2012, p. 121.

[11] In the present context, the context-sensitivity of knowledge is understood *ontologically* as the claim that knowledge is not built into the *for-ness* of dispositional partners, and that, thus it can only be explained by supplementing the competence condition with a *safety* condition. In this sense, the context-sensitivity of knowledge is equivalent to the notion of *epistemic dependence*. Again, epistemological contextualism is not the issue at hand here. For the concept of epistemic dependence, see Kallestrup, J.; and Pritchard, D.: "Dispositional Robust Virtue Epistemology versus Anti-luck Virtue Epistemology", in M.A. Fernández Vargas (ed.), op. cit. pp. 45-6. [Henceforth RVE/ALVE]

for short) take this route[12]. There are, then, two problems. One is a question about the *nature* of competences (and dispositions in general), which revolves around whether and to what extent they are dependent on external factors. Also, there is the question about the context-sensitivity of *successful events*, which ties in with the issue of whether those events are *manifestations* of powers of the actual world or whether they are, because of their dependence on unactualized possibilities, a matter of what happens in modally close possible worlds. On the latter view, the property of being safe as instantiated by a belief does *not* supervene on cognitive features of the believer, and knowledge is only partially grounded in actuality. Thus, it is not a manifestation of competence. It is important to be clear about two matters. First, at least as Pritchard (2010) formulated ALVE, it is not true that the fact that S safely believes p has to be attributable to an epistemic competence. As Pritchard aptly notes, what is essential for ALVE is that it is a *dual- condition* view, one that incorporates as conditions for knowledge the *ability condition* stressed by robust virtue epistemology (RVE, for short) and the *anti-luck condition* characteristic of safety-based epistemologies, and that, *crucially*, it incorporates them in such a way that "it accords each condition equal weight" as "they are each answering to a fundamental intuition about knowledge"[13]. As we will see below, that Pritchard has not changed his overall position is substantiated by the fact that as recently as 2016 he charges Sosa of implicitly building a safety clause into judgmental competences. The function of such argument is that of showing that even RVE proponents cannot evade the safety condition. But that is quite different from asserting in one's own person what, according to Pritchard, RVE defenders should claim –that safe beliefs supervene on higher-order competences. On the plausible view that ALVE aims at explaining knowledge while avoiding any appeal to the evaluative powers of the epistemic subject, and thus to agency, it would be unreasonable to attribute such account to Pritchard.

More importantly, the two problems above mentioned are related in the following sense: they are rooted in the view that an extrinsic, ontologically independent situation –one where there is *no* predetermined, reciprocal direction of fit among dispositional partners– is crucial for the explanation of successful events such as knowledge, either *directly* as a factor additional to competence, or *through competences* as the causal context in virtue of which competences are acquired and lost. Besides, and as it was previously remarked, the mark of ALVE is a dual-condition view. Let us now take for granted that since ALVE is unable to provide a principled distinction between safe environments and mimics –one that only RVE could provide, it is untenable. Dialectically, it would still be open for its defenders to explain the acquisition of fine-grained cognitive competences in terms of *extrinsic factors*. This strategy would yield a dual-condition view of sorts –one where instead of being an additional factor to competences, the environment is one of the factors that constitute them. On this model, it would be the right kind of situation that makes S to be disposed to get knowledge by using her cognitive ability. The important

---

[12]   For a recent defence of this approach to knowledge, see RVE/ALVE, pp. 31-50. Works developing earlier versions of anti-luck virtue epistemology include Pritchard, D.: *Knowledge*, Basingstoke, Palgrave Macmillan, 2009, pp. 72-85; Haddock, A.; Millar, A.; and Pritchard, D.: *The Nature and Value of Knowledge: Three Investigations*, Oxford, Oxford University Press, 2010, pp. 48-65 [Pritchard's contribution]; Pritchard, D.: "Anti-Luck Virtue Epistemology", *The Journal of Philosophy*, 109(3), 2012, pp. 247-79.

[13]   Haddock, A.; Millar, A.; and Pritchard, D.: op. cit., p. 54.

point is that from a *dialectical* point of view the thesis of the ontological context-sensitivity of dispositions could play the role of a *second line of defence* argument, if not for ALVE itself, at least for a modification of it such that preserves its basic commitment to epistemic dependence. Let us, however, not forget that the issue at hand is that of ontological context-dependence, whether referred to the nature of dispositions or to successful events. Since context-sensitivity is the common ground of ALVE and externalism concerning the nature of abilities, both ways of fleshing it out have to be tackled.

The first objective of this article is thus to meet the challenge of ALVE while sharing with its advocates the view that competences are modally robust across contexts[14].

The second objective of this article is to provide an *ontological background* for the context-sensitivity of epistemic competences such that context-sensitivity would be compatible with the following three theses: (i) cognitive competences are intrinsic properties, (ii) competences (generally) are multi-track dispositions whose powerful nature is not exhausted by a single-track manifestation-type and, thus, are not individuated by a single-type of 'successful' result, and (iii) the modal nature of epistemic virtues is that of potentiality, as it is usually expressed by 'can'.

To this end, it will be instrumental to distinguish the circumstances for the *manifestation* of the disposition to obtain from the *possession* of the disposition, as well as to keep apart the claim that competences are intrinsically directed to highly contextualized manifestations (the *context-directedness* of dispositions, for short) from the claim that competences are ontologically *context-dependent*. The latter distinction corresponds to the difference between the *ontological* and the *semantic* senses according to which competences are context-sensitive. Competences are context-sensitive in an ontological sense as having an intrinsic content and directedness to *reciprocal* manifestation with the right *dispositional partners*, namely with those partners which would fit with the said competence as to result in F (*reciprocity* is thus an intrinsic property of dispositions and powers)[15]. However, let us bear in mind that context-sensitivity is usually a matter of language –and of the pragmatic interests of the community–. The ways of picking out competences via reference to selected manifestations, privileged results, right subjunctive conditionals, and favourable environments should not be projected into an ontological picture of dispositions and abilities.

As I see it, the question is about how to provide a thicker ontological base than meagre potentiality for pragmatic mechanisms to pick out dispositions. It is thus an issue of how to be realistic about the truth-makers of fine-tuned dispositional terms without mapping dispositions according to the context-sensitivity that permeates our

---

[14] As it is expressed by the lemma from Aesop 'Hic Rhodus, hic salta' (Here is Rhodes, jump here!). This motto is commonly used in the literature of dispositions to point to the invariant and portable nature of abilities.

[15] Another way dispositions are context-sensitive in an ontological sense is that they come in *degrees*. This means that they have a proportional nature that, far from being fixed according to how things stand in the particular environment where the object is currently located, is determined by tracking the disposition along a maximal variety of cases and, in Sosa's words, a "sufficient spread of possible shots." (J&A, p. 97) For example, degrees of irascibility are measured along multiple cases. Two subjects can be equally disposed to get angry in a given situation while only one of them would count as irascible (in regards to a sufficient spread of manifestations). For the introduction of *gradability* as an intrinsic feature of potentialities in the literature of dispositions, see Manley, D.; Wasserman, R.: "A Gradable Approach to Dispositions", *The Philosophical Quarterly*, 57(226), 2007, pp. 68-75.

language about powers and abilities. Context-dependence is an artefact of language. There are, however, ontological features of dispositions that make it possible to bridge the gap between bare potentiality and pragmatic (and linguistic) regimentation. The scope of a dispositional term may well be pragmatically selected. However, even so, the *proportional nature* and the *conditions built into the competence for its proper manifestation* are objective features.

To my mind, the ontological model that follows is in keeping with Sosa's "metaphysical explanation"[16] of knowledge, such as it is developed throughout his corpus. If there is a difference, it is in that I am mainly interested here in discouraging a reading of Sosa's recent theory of the constitution of competences such that it would be considered as a substantial departure from his earlier views, and, as thus, his virtue epistemology as drifting into uncongenial accounts. The challenge is to show how context-sensitivity can be introduced while retaining what in my view are Sosa's leading intuitions: (i) that the *modal force* of cognitive competences is entirely grounded in their bases so that external factors have no bearing on the competences the subject possesses and (ii) that the modality of knowledge is anchored in the actuality of epistemic powers capable of producing such manifestation. For this purpose, it is important to bear in mind that when advancing the triple-S structure of complete competence, Sosa focuses on *one aspect of competences*, and thus, that his analysis is *relative to* competences *partially considered* as abilities to succeed.

Here is the plan. In section 2, I will present the epistemic Twin Earth argument against Sosa's robust virtue epistemology which was recently proposed by Jesper Kallestrup and Duncan Pritchard (2016), and I will describe the notion of the epistemic dependence of knowledge that they propose. I will then argue that anti-luck virtue epistemologists are committed to the claims that (i) the modality of knowledge is that of counterfactuals, and (ii) safety clauses explain the modal force of knowledge. In section 3, it is explained how the latter two claims make ALVE unable to avoid the problem of finks, masks and mimics that plagues subjunctive conditionals and provide the appropriate *ordering source* to capture the modal force of knowledge. This failure of anti-luck virtue epistemology to accommodate cognitive competences and thus, to offer a satisfactory account of knowledge, serves as a foil to flesh out my proposal. On this alternative account, *gradability* and *reciprocity*[17] being the intrinsic features that lay bare the right ordering source for abilities and their manifestations, they offer a non-contextual ontological base for context-sensitivity. By way of holding the intrinsic properties of the subject constant, the present proposal removes from masks and mimics their relevance. Finally, in section 4, I will deal with the problem of the individuation of cognitive powers. My aim here is to defend a model according to which multi-track powers are invariant across possible cases and scenarios, as well as to offer a diagnosis of why intuitions on this topic are often fluctuating and conflicting. This article thus aims at making explicit the *ontological underpinnings of knowledge* in a way that it is consistent with Sosa's fully dispositional virtue epistemology. To this end, the constitution of complete competences needs to be placed in proper perspective.

---

[16]   J&A, p. 8.
[17]   Reciprocity is the fit of joint potentialities to mutual manifestation.

2. Outsourcing Modality

In "Dispositional Robust Virtue Epistemology versus Anti-luck Virtue Epistemology" (2016), Kallestrup and Pritchard argue that while Sosa's theory cannot accommodate the epistemic dependence of knowledge, there is an alternative account –anti-luck virtue epistemology– that fares much better in this regard. They put forth a thought experiment that involves duplicates to bring this point home.

Roughly put, the argument runs like this. Let us suppose that S (a chemically ignorant subject) is in full possession of her relevant first-order cognitive competences and that she is located on Earth, where watery stuff[18] is always water (a substance which essentially is H_2O). Let us add that the Earthian agent (S_E) holds a demonstrative belief that that's water on the basis of her reliable perceptual and recognitional competences, and that the belief is true. By way of comparison, let us consider a duplicate of S_E endowed with the same cognitive powers like hers, and who also holds a true demonstrative belief that that's water on the basis of them, but who is situated on Twin Earth, a planet where there are here and there patches of watery stuff that is Twin-Earth water (XXX). Let us add that S_{TE} is situated in a *local area* that happens to be safe and thus, there is no chance for her but to hit upon water in regard to her current position. However, S_{TE} might easily step into a modally close area that contains patches of Twin-Earth water. Provided that the *regional area* is epistemically polluted, the local situation, although narrowly safe, could easily have been unsafe. In view of the foregoing, it seems plausible to claim that while S_T knows that this is water, her counterpart on Twin Earth does not know it. What this comes to is the claim that knowledge is crucially dependent on the propitious nature of the environment in which S is situated, so that there is an unbridgeable gap between cognitive achievement and knowledge[19] –one that is always filled (if at all) by a fortunate environment. As thus, true beliefs due to competence (*apt beliefs*, following Sosa's technical terminology) fall short of knowledge. Competences and their manifestations as apt believing are not *sufficient* for the truth of the corresponding attribution of knowledge[20].

---

[18]   'Watery stuff' could be captured in two-dimensional terms by saying that it is a primary-intension term that picks out whatever property plays the water role in a given world independently of its nature and chemical profile. In this sense, both terrestrial water (H_2O) and Twin-Earth water (XXX) would have a different secondary-intension (nature) while sharing the same primary-intension (role).

[19]   Let us recall that according to the scenario, SE and STE are on a par regarding the reliability of their relevant, first-order competences (they are equally competent). As such, each of their beliefs counts as a cognitive achievement even though only one of them counts as knowledge.

[20]   I am only focusing on the conclusion from the argument for negative epistemic dependence, where the bad situation in which one's performance takes place is the *explanandum*, and one's lack of knowledge is the *explanans*–. Kallestrup and Pritchard's overall conclusion is, however, much stronger (*cf.* RVE/ALVE, pp. 45-6). Cognitive success because of the manifestation of cognitive competence is, for them, neither a sufficient nor a *necessary* condition for the truth of knowledge. Cases of positive epistemic dependence where (allegedly) the explanans is S's possession of knowledge while lacking epistemic skills seem to support the stronger thesis. There is a possible ambiguity here, however. Kallestrup and Pritchard explicitly endorse that in cases such as epistemic Twin Earth competences and their *manifestations* are left intact by unsafe environments (*cf.* RVE/ALVE, p. 39), thus suggesting that, in those cases, success due to competence is at least a necessary condition for knowledge. A challenge for their approach would thus be to provide a *unified account of knowledge* that is able to close the gap between the two varieties of cases at hand. Clearly, the two claims that a safe environment *supplements* aptness for getting knowledge and that it *replaces* aptness to explain knowledge do not go together well.

The first thing to be said of the above argument is that, although it is of the family of fake-barn scenarios, it suggests a more demanding reading of the safety clause than the *local safety* (or unsafety) operative in the standard case of Barney. What the argument suggests is that it is not enough for an environment to be relevantly safe in order to have knowledge that a particular area happens to be safe. It is also required that the agent *could not easily be situated in an unsafe area*, or, in other words, that the situation could not easily be unsafe. Knowledge thus requires a *super-strong safety*, the scope of which goes well beyond the current local area where the agent is placed and thus, makes it false that the agent might easily step into a bad epistemic area. Curiously enough, the epistemic Twin Earth scenario is analogous to one Sosa considers in *Judgment and Agency* where he provides a different explanation of the subject's lack of knowledge[21].

Let me dwell a bit more on the distinction between the *local* and the *regional* areas, and consequently on what a super-strong safety condition –one that the former distinction implicitly suggests–means. It is plausible to see this *expansion* of the situation over which the relevant safety condition quantifies as motivated by cases as the following one–which is a variation of a well-known simile offered by Sextus Empiricus[22]:

> DARK ROOM. Goldie steps into a pitch dark room that, unbeknownst to her, only contains objects made of gold. Such room is merely one among many other similar ones into which she could easily have stepped –rooms that contain no golden object. Due to the contents of the room and to her motor and grasping skills, there is no chance for her but to hit upon gold. Despite all this, her success is in a relevant sense accidental. It was by a stroke of luck that Goldie hit upon gold.

As a way of dealing with this scenario, ALVE defenders would claim that it is not enough for an environment to be safe that a local area happens to be safe. For that, it is also required that the (narrow) situation could not easily be unsafe. In this sense, what would explain Goldie's lack of knowledge is how easily she could have stepped into any of the other rooms. The suggestion is thus that of reading the safety clause as a super-strong safety, one that is met in cases where factors that would make the local area unsafe are remote possibilities not liable to occur. To accommodate apparent counterexamples as DARK ROOM, the focus is shifted from a narrower to a wider area[23].

The second point to note in the epistemic Twin Earth argument is that it hinges on the claim that both S_E and S_{TE} hold beliefs whose accuracy (success) is explained by the exercise of the relevant epistemic competence. It is the fact that S_{TE} forms an

---

[21] Sosa's scenario involves a pilot in training, Simone, who is carried blindfolded to a real cockpit in a real aircraft so that on the target-occasion, she is shooting real targets aloft and having a true belief about that event. However, unbeknownst to her, Simone is sometimes led to a simulation cockpit. For a detailed analysis of this scenario, see J&A, pp. 156-163.

[22] Sextus Empiricus: *Against the Logicians*, Cambridge, Cambridge University Press, 2005, p. 12 [*M* vii 52].

[23] This move is, however, plagued with problems. For one, and lacking a principled way for restricting relevant areas, the wider area to consider might be indefinitely expanded. More importantly, the point of Pyrrhonian cases is not so much to raise the bar on anti-luck conditions as to challenge any variety of epistemic externalism (ALVE included) by pointing that none can properly claim that she has grasped the truth even if her beliefs are relevantly safe. Even within a strongly safe environment, Goldie's belief would be blind. The Pyrrhonian challenge remains uncontested.

apt belief while lacking knowledge that suggests, contra Sosa, that knowledge is not apt believing and thus, that an external anti-luck factor must be included to get the right results.

It is clear, however, how Sosa would respond to this argument. Sosa would offer a diagnosis of his critics' account in terms of the *misidentification* of the competence relevant for *knowledge*. For Sosa, this is a *second-order judgmental competence* of rational agents to evaluate their first-order performances in view of an epistemic perspective, an ability that is manifested as knowledge only if the agent's epistemic performances are guided to aptness by apt judging[24]. The factor that, in Sosa's view, prevents STE from getting knowledge is thus that her belief is deprived of full aptness. Sosa would classify the case neither as an illustration of environmental luck, nor as an example in which the agent is deprived of the relevant ability, but as an instance of *manifestation failure*[25].

Kallestrup and Pritchard think that the above argument flounders. The reasons offered for that conclusion are: (i) that there would be a *first-order epistemic difference* between SE and STE even if both of them were not judging (or were judging badly), a difference that the appeal to the second-order would leave unexplained[26], (ii) that, since we are dealing with perfect duplicates whose judgmental competence would be identical, the thought experiment can be iterated while rendering the same result[27], and (iii) that Sosa must build a strong safety clause into the judgmental competence as a condition for its manifestation as a fully apt belief, thus implicitly yielding the high ground to his opponents[28].

The evaluation of these arguments will take place in section 3. Let us just remark for now that the third consideration seems specially compelling, if only because it might appear that the complex, external conditions for the manifestation of the judgmental competence to occur overlap with the wider situational range over which safety conditions quantify.

What is clear now is that the Twin Earth argument, as it is originally presented, does not necessarily entail ALVE, and thus, it has to be supplemented both with an effective critique of alternative interpretations and with a compelling case for preferring such reading to render the result favored by anti-luck virtue epistemologists. Thus, an ironclad argument is replaced by the assessment of costs and benefits.

Let us ignore for the moment that it is not obvious at all that ALVE picks out the relevant ability for knowledge. Thus, let us concede for the sake of the argument that both SE and STE form apt beliefs, which manifest epistemic competence that would count as knowledge for robust virtue epistemologists. As Kallestrup and Pritchard underline, there is a contrast between the *modal sturdiness* of competences and its lacking in the case of knowledge[29]. However, it seems that such modal sturdiness *must* be transferred from abilities to apt beliefs (their manifestation) under pain of classifying apt beliefs as *epistemically dependent* in the same sense in which

---

[24]   *Cf.* J&A, pp. 150-151.
[25]   Importantly, Sosa would also claim that ALVE does not provide an adequate explanation of why SE knows p. Safety clauses and reliability conditions can be met, while compatible with *guessing*. This is why Sosa endorses a Pyrrhonian bi-level epistemology.
[26]   RVE/ALVE, pp. 40-41.
[27]   *Idem*, p. 41.
[28]   *Idem*, pp. 41-42.
[29]   *Idem*, p. 39.

knowledge is epistemically dependent for anti-luck theorists. However, and since the *manifestation* of the disposition is an external event that depends on how things stand in the vicinity of the bearer of the ability, this proposal cannot work. As a result, the epistemic dependence of interest for ALVE is laid bare as being loaded with a stronger meaning than that which is operative in the conditions for successful manifestation–as depending on a *modal factor* –safety[30]. What is, however, the modality inherent to safety? Painted in broad brushes, it is Lewisian *non-localized possibility* restricted by the usual mechanisms of *centring* and *closeness* to the actual situation under consideration.

The crux of ALVE is that for S to *actually* know p it has to be true *for all the worlds relevantly close to the actual world* that a perfect duplicate of S forms a true belief that p. The same point could be captured by saying that S knows p if and only if, for all the nearby worlds to the target-situation, if S were exposed to the relevant triggers, S would form a true belief. In formulae: K_s *iff* $T_s \,\square\!\!\rightarrow R_s$ (where the modal operator for necessity is the universal quantifier that ranges over the relevant close worlds)[31].

It is apparent from the preceding presentation that Kallestrup and Pritchard consider that the actual world cannot provide for knowledge, and, following closely on the steps of Lewis, that we simply need more worlds added to the actual one to do it. The modality of knowledge is, thus, removed from the competence base, becoming a matter of non-modal facts, that is, of facts that are *superficial* insofar as they are deprived of dispositional *depth*. The truth- makers of statements attributing knowledge to S are events that happen in non-actual worlds so that, at the end of the day, what makes true that S knows p is a bounded collection of counterfactual truths about the subject's (or, better said, about her counterpart's) response to stimuli. Knowledge is abstracted (subtracted) from the dispositional profile of its possessors, as it is made clear by ALVE's claim that knowledge is a target that success due to competence cannot get on its own. The point is that even if ALVE involves a virtue condition which will be met by an agent's belief in virtue of certain facts about her cognitive architecture, such condition is according to this view never sufficient to explain knowledge. Knowledge is again detached from virtue.

Using this model to explain knowledge, one has to focus on some particular *initial conditions* for S and then, while holding S constant, imagine alternative worlds (possibilities) with *similar* initial conditions. Plausibly, those initial conditions would include the situation and belief-forming methods, so that a safety condition quantifies over both relevantly similar situations and relevantly similar belief-forming methods. The similarity of initial conditions thus defines closeness to the

---

[30]   As it will be elucidated below, the modal sturdiness of competences and the modal nature of safety refer to different varieties of modality. As the former picks out a fully *actual* modality rooted in the nature of dispositions, the latter is modal in regard to what happens in close *possible worlds*, that is, in regards to unactualized possibilities whose truth-maker is not a power in the actual world. We could talk of Aristotelian and Lewisian (Humean) possibilities. The logical form of Aristotelian possibilities is 'It is possible for S to F', whereas that of Lewisian modalities is 'It is possible *that*…followed by a sentence'.

[31]   I would have said that according to anti-luck virtue epistemologists, S knows that p if and only if, for all the relevant worlds, if S were exposed to the appropriate triggers, S would hold an *apt belief* (a belief that is true due to her competence), had it not been for the fact that it is not clear whether Kallestrup and Pritchard hold that for some cases, apt belief is necessary for knowledge, or whether they think that for all cases, only the adroitness of the belief is necessary for that (see note 19). To be on the safe side, I have preferred the latter (more liberal) interpretation. My arguments in what follows could be adapted to accommodate both views.

actual world. Let us consider again the case of Sᴛᴇ. If, as in her case, included among the initial conditions is the presence of a patch of twin water to which S has an easy access, then –and even if, as happens, Sᴛᴇ holds a true belief that that's water– it might very easily have occurred that Sᴛᴇ falsely believed that it was water. Thus, the world where Sᴛᴇ holds such false belief is close to the world that is held fixed as actual. Such closeness prevents her from actually having knowledge. It is therefore as if we were focusing on the fact that a glass is standing close to the edge of the counter –so that it could easily be hit, drop and shatter– in order to say something true about what is happening to the glass now.

It is of the greatest significance to bear in mind that the safety clause is nothing else than a way of *outsourcing* the modality of knowledge. For Sosa, however, the actual world provides for knowledge. Through thick and thin, cognitive success is in his view deeply rooted in competences. My question is: How can Sosa be guilty of building a safety clause into the judgmental competence when, as a hardcore actualist, it is what has to actually occur for a reciprocal manifestation of multiple partners to happen, and *not* what might easily happen, that explains knowledge? Manifestations can thicken in regards to their conditions without reaching out to alternative worlds.

3. Costs and benefits

An obvious place to begin the assessment of ALVE and RVE is to recall that ALVE defines knowledge in terms of a biconditional whose right side includes a counterfactual that has to be true for all the worlds where the initial conditions for S are similar. Is this subjunctive conditional able to capture knowledge?

The problem is that, because no restriction of closeness can exclude masks and mimics[32], knowledge is neither sufficient nor necessary for the truth of the corresponding counterfactuals. It is not sufficient because S can know p even when his belief could have easily been false. It is not necessary because the conditional could be true even when S is intuitively deprived of knowledge. The point is that the reductive analysis of knowledge that ALVE proposes fails as it fails the conditional analysis of dispositions –and for the same reasons. There are counterexamples both to the right and to the left side of the biconditional.

---

[32] Those two terms were introduced in the literature about the conditional analysis of dispositions by Johnston (*cf.* Johnston, M.: "How to Speak of Colors", *Philosophical Studies*, 58(3), 1992, pp. 232-33), who also refers to finkish dispositions by the name of *cases of altering* (*Idem*, p. 232). Finkish dispositions are cases where the trigger for the manifestation also causes the dispositional base of the object to alter so that, for instance, a glass that was fragile before t1 does not shatter when stressed because it became hard at t1 through the good offices of a glass-friendly sorcerer. Unlike them, masks are interferers that prevent the manifestation to occur while the object retains its dispositional base. Masks thus undermine the Reformed Conditional Analysis that Lewis proposed to meet the challenge of finks (*cf.* Lewis, D.: "Finkish Dispositions", *The Philosophical Quarterly*, 47(187), 1997, pp. 143-58). As for mimics, these are cases where, despite the absence of the dispositional base, the subjunctive conditional is true. Mimics were introduced in the literature by Smith (*cf.* Smith, A.D.: "Dispositional Properties", *Mind*, 86(343), 1977, pp. 439-45).

It is obvious that, since for ALVE knowledge is not the manifestation of a disposition, the argument from masks and mimics does not operate here as it does in standard discussions about dispositions. However, masks and mimics can be adapted to the present discussion. In this context, they are relevant to show both that knowledge can be possessed in spite of the closeness of interferences and that a modally clear area is not *relevant* for an apt belief to count as knowledge.

Consider, first, how the right side can be false while S has knowledge. Think of any interferer in the wings that could have easily spoilt the subject's epistemic performance. Does such a possibility affect the quality of the performance as it is? Does it affect its actual successful result? Not in the least. If it is true that the belief is apt at all, it is not less true that, in spite of such easy possibility, the agent is creditable for its success.

This much is conceded by ALVE. An obvious rejoinder to the previous consideration is, of course, that unsafe circumstances such as the interferer deprive the agent of knowledge, and thus that the safety clause is a supplement to apt epistemic performance for the purpose of explaining knowledge.

However, is such a rejoinder true? ALVE's response seems intuitive only in so far as one pays exclusive attention to cases such as fake-barn-like scenarios, where it is obvious that the subject lacks knowledge, and *only* to the extent that there is no alternative to ALVE for explaining why it is so. In the view of intuitive alternatives, however, we seem threatened by an impasse. The issue has, thus, to be decided in terms of how intuitive ALVE remains in a variety of cases.

Let us go back to S$_E$ (the good case). As ALVE would have it, S$_E$ gets knowledge because of the absence of twin water. The problem is that there might be many other easy possibilities that could have prevented S$_E$ from getting knowledge –possibilities that should be intuitively excluded to evade the kind of accidentality of interest for ALVE but which cannot be ruled out by isolating a restricted modal area. A demon, for instance, could have easily intervened so as to swap water for twin water, but for a reckless whim[33].

ALVE's proponents would now insist –as they often do for any apparent counterexample– that safety conditions quantify only over relevantly similar initial conditions; the demon-situation is not relevantly similar to the initial situation; and S$_E$ knows p insofar as there is no sufficiently similar world where S$_E$ falsely comes to believe p via a sufficiently similar method.

The first reply to this argument is that the demon is *not included* within the initial situation as it is restrictively described. Which means that it operates (or not) independently of how closely similar are the initial situations under consideration. The situations are abstracted from a broader context, and may thus be identical whatever happens in the wider world. Obviously, proponents of a safety condition would still insist that what would make of the demon an easy possibility –one over which safety clauses quantify– is that the demon- situation is relevantly similar to the initial situation of S$_E$ taken in itself. Since the situations are not similar, the demon does not qualify as an easy possibility. On this view, easy possibilities are relative, and internal to initial situations. The real question is thus whether factors external to the given situation are relevant for safety. And the intuitive answer is *yes*. If the function of safety clauses is that of capturing *objective* easy possibilities, then they

---

[33] Kallestrup and Pritchard seem to rule out far-fetched scenarios such as the demon above mentioned by stipulating that there are no nomological differences between Earth and Twin Earth (RVE/ALVE, p. 38). Two things can be said of this strategy. First, it does not rule out the whole gamut of empirical interferences that would count as being modally close. Secondly, and more importantly, nomological relations –either of the weak variety proposed by Armstrong's 'nomic relations of necessitation' or of the strong dispositional kind– are forbidden by Lewisian modalities. On the latter view, and according to a combinatorial model based on conceivability, there are as many worlds with similar initial conditions in which a demon intervenes after t1 as there are worlds in which no such interference occurs. There is a price to pay for safety.

cannot be arbitrarily limited to quantify over a partial set of initial conditions. The point is that a partial, ad hoc quantification of safety conditions is not able to restrict easy possibilities, and thus, that it fails to capture the intuitive meaning of safety.

Curiously, this is a lesson we learnt from DARK ROOM. The advocates of ALVE were then obliged to expand the initial conditions in order to accommodate cases of the same sort as DARK ROOM, where there is a mismatch between our intuition that S's belief is unsafe in spite of how safe the local area is and the initial conditions that would make of it a safe belief. The trouble comes from the fact that the demon is an exact analogue of DARK ROOM. As such, its threat cannot be met without expanding the initial conditions so that they include a clause ruling the demon out. However, once the initial conditions start to expand so as to exclude factors that would count, if true, as bad situations and easy possibilities, there is no way to non-arbitrarily stop the resulting conjunction to increase. The initial situation would be so wide as to be unmanageable. Which leaves ALVE with two options –either to exclude the safety condition so as to explain knowledge (RVE), or to conclude that knowledge is unfathomable. These are the alternatives: a unified model according to which possible masks and interferers are epistemically irrelevant for explaining either knowledge or the lack of it; or a theory that is unable to restrict closeness relevantly and cannot establish a viable contrast between both cases.

This can be made more apparent by considering how complex the conditions for the manifestation of the judgmental competence are and how many factors, external as well as internal, have to hold together for it to occur. Alternatively, this can be accomplished by bearing in mind that when the agent holds a full, apt belief, those factors could not have existed or have been prevented easily *without the subject's being deprived of knowledge*. The point is that the so-call 'unsafe circumstances' can as well be seen as a collection of fully *actual* facts that explains a manifestation failure, and that 'undermining-knowledge luck' is nothing else than a variety of a *benign luck* that neither deprives nor contributes to the epistemic standing of the agent as it actually is. Alternative worlds are, in short, epistemically inefficacious.

To my mind, the illusion of epistemic counterfactuals stems from the ways in which modal safety and appropriate conditions for the manifestation of knowledge seem to extensionally overlap. There is, however, a great difference between explaining STE epistemic shortcomings by appealing to what is happening in her surroundings (when they are so constituted as to prevent the acquisition of apt belief through the exercise of apt judgment) and explaining them by appealing to what could have happened if the story were different. Let me hasten to add again that, while RVE has the advantage of making possible interferers irrelevant for the possession of knowledge and can thus explain knowledge per se, ALVE, while making them relevant to it, does not have the resources to exclude unwanted worlds from the picture. Thus, ALVE sets standards for knowledge that cannot be plausibly met.

Consider, in second place, how the right side of the biconditional above can be true while S lacks knowledge. For that, let us borrow an scenario from Sosa[34] –one where a veridical mimic is introduced[35]:

---

[34]   J&A, p. 103.
[35]   Mimics come in two varieties: as faking a competence that does not exist and as *veridical mimics* that replace what would count as the manifestation of an extant competence for a success that is not sufficiently attributable to S. In the latter case, the success still counts as accidental.
      Traditional mimics of the first variety are cases in which the disposition ascription is false even if its associated

BENIGN ANGEL. Situation and shape are such that were S to form a perceptual belief p through competence, they would combine together as to render an apt belief. However, a benign angel the aim of which is to *secure* success only if S *acts competently* interferes. Though S comes to form a true belief, its truth is not due to competence.

There are four important things to note about BENIGN ANGEL. For one, mimics provide a causal explanation of the performance's success at the expense of making it *accidental*. Second, cases of mimicking can be *regular* and *systematic*, so that one could form a boundless set of competent, true beliefs compatible with accidentality. Besides, the demon, whose purpose is that of securing success, makes it true that there is no sufficiently similar world where S falsely comes to believe p via a sufficiently similar belief-forming method. Which means that while in the particular case at hand success is *overdetermined*, the demon would also secure success whenever competence would fall short of aptness. Finally, the scenario has been described so as to make salient that competence could still be causally operative for the result to obtain (thus contributing to it, just as for ALVE competences apparently contribute to knowledge while falling short of it) while compatible with mimicking (after all, the demon would not have interfered if S were not relevantly competent).

On first thoughts, one could say that the trouble for ALVE comes from the fact that all of the features above mentioned are attributable to a safe environment, and so, that there is *no difference* at all between a safe environment and a mimic. From this would follow that far from explaining knowledge, safety conditions make luck pervasive. However, proponents of ALVE could object to the apparent similitude by pointing that since BENIGN ANGEL is targeting the first-order ability condition of knowledge –one that far from removing, ALVE supplements–, there is no analogy between standard cases of mimicking and safe environments. They would thus endorse Sosa's diagnostic, and claim that the success of the subject's belief is not sufficiently attributable to the exercise of the subject's epistemic competences, since it depends crucially on the angel.

Granted, there is a disanalogy. But it is enough to shift the epistemic target from apt beliefs to knowledge for seeing its lacking of relevance. Safe environments are mimics insofar as, on ALVE's view, they secure knowledge whenever competence would fall short of getting it –which is *always*. The lesson from BENIGN ANGEL is that success that is not attributable to competence is accidental. This lesson equally applies to apt beliefs as to knowledge. What this comes to is the claim that, for success not to be accidental, it has to be the manifestation of competence. By leaving the gap between apt beliefs and knowledge unfilled, ALVE fails to rule out epistemic accidentality[36].

---

conditional is true. An example might be that of an iron dumbbell that even though it is not intrinsically disposed to shatter into pieces when gently touches a solid surface, it would be immediately zapped by a hovering demon were to touch a surface. By contrast, veridical mimics are cases in which, even if the dispositional ascription happens to be true, this is not the reason that makes the associated conditional true. The scenario of BENIGN ANGEL, which is thoroughly discussed in the main text, serves as a paradigm example of the latter variety of mimics.

[36]   At this point, the proponents of ALVE might opt for denying that there are competences directed to knowledge, and thus, for claiming that there is no conceptual opposition between high-order virtues and mimics. However, this move would still leave BENIGN ANGEL (properly described) as a case of safe environment. And it would make accidentality –which does not disappear because the corresponding competence is uninstantiated– pervasive.

The challenge is now to show whether criticism can be equally levelled against RVE. Recall, in this regard, that Kallestrup and Pritchard find fault in Sosa's view of knowledge as the manifestation of judgmental competence. In their view, Sosa's account, while leaving unexplained the first-order epistemic difference between subjects located in different modal surroundings, fails both to respond to an iterated epistemic Twin Earth argument and to afford a definition of knowledge that does not include safety as a condition for manifestation.

As for the first criticism, the alleged difference is perceived only insofar as one is driven by ALVE. If on the contrary, one takes that the first-order epistemic successes are *manifestations* of first-order dispositions that, as such, and in Sosa's words, are not required to meet any test "of robustness across neighborhood"[37], both Sᴇ and Sᴛᴇ are on a par. There is, obviously, a difference between the epistemic situations of the two agents. But let me ask how they can impact on first-order, narrow abilities which, by stipulation, are part (as well as their successful exercise) of the common profile of Sᴇ and her counterpart. It eludes me how to set up the thought experiment without keeping the first-order epistemic profile constant. Regarding the situational profile, its epistemic difference can as easily be captured in terms of whether the conditions for apt judgment are satisfied as in terms of safety conditions. In any case, this is not an argument contra RVE since it presupposes prior and independent reasons for ALVE.

On the other hand, the second argument stems from a confusion between the possession and the manifestation of competence. Granted that Sᴇ and Sᴛᴇ are exercising the same judgmental ability, there is still a factor that has to be considered for attributing knowledge: whether the conditions for the manifestation of the competence as knowledge are satisfied. Since those conditions include external partners that have to properly fit with the competence for the manifestation to occur, nothing of relevance follows from holding the judgmental competence across scenarios fixed.

The whole question boils down to whether RVE includes a safety clause among the necessary requirements for knowledge. However, it has become clear by now that Sosa's proposal, far from being committed to possible worlds and counterfactuals, is actualist through and through[38]. Knowledge is, for him, *epistemically dependent* on factors external to the agent.

However, those factors are not events that might happen. Epistemic dependence boils down to dependence on the objects that actually are in the agent's surroundings for the right manifestation of her judgmental competence.

To my mind, Sosa drives this point home by distinguishing *proximal* from *distal* dispositions[39]. While proximal dispositions are abilities whose manifestation does not depend on how things stand in the (narrow) neighborhood of the performance, distal dispositions are competences whose target-manifestation is founded on the nature and combination of the right *dispositional partners*. Perceptual abilities,

[37]   Sosa, E.: *Knowing Full Well*, Princeton & Oxford, Princeton University Press, 2011, p. 83.

[38]   *Actualism* refers not only to the physical makeup of objects, but also to their dispositional profile. As a side comment, let us add that the many ontological problems which plague both Pythagoreanism and Nomic views can only be plausibly met by postulating that dispositions are part of the basic furniture of the universe, and endorsing the Limit View, which considers the purely categorical and the purely dispositional as the two sides of the same coin. On the latter view, see Martin, C.B.: *The Mind in Nature*, op. cit., pp. 54-79.

[39]   J&A, p. 101.

which are manifested regardless how bad the (broad) environmental conditions are belong to the first class. Judgmental abilities come, by contrast, with conditions of manifestation wherein the (broad) *situational* factor is all-important. Knowledge being a matter of the mutual manifestation of powers that work in tandem, it is evident why such success cannot be reduced to a first-order success wherein the situational component is (almost) non-existent.

However, it is important to bear in mind that both proximal and distal dispositions are *on the same spectrum* in regard to their (lack of) relation to possible worlds and counterfactuals. They differ in terms of *granularity/reciprocity* from a manifestation-point to a *thickening* of the manifestation conditions such that they include a constellation of reciprocal partners. However, they are equally unrelated to what might have happened. The situational component of distal dispositions is not understood here in terms of the restriction of unwanted worlds. It is, rather, defined in view of the actuality of competences and their manifestations. The ordering source for the modality of intellectual virtues is, thus, grounded in the intrinsic directedness of the different powers towards more or less grained manifestations.

Let me finally add that, though knowledge is the manifestation of co-operative powers, joint potentiality does not entail co-possession. A glass wrapped in styrofoam is less disposed to break than an unprotected glass. However, this lower disposition to break is a property of the glass and not of the styrofoam. Likewise, it is only the rational subject who possesses a disposition to know, and who, when such disposition is actualized, is standing in a relation of knowledge to the target-fact. Knowledge is thus grounded in many powers, while the relevant epistemic virtue is possessed by only one of them.

The emerging picture preserves the context-sensitivity of knowledge while excluding the kind of accidentality relevant in epistemology and avoiding the pitfalls of interferers. It remains to be seen whether as the previous example of the lower disposition of the glass to break when wrapped seems to suggest, the relevant disposition of the agent to know is an *extrinsic disposition* that she gains from her situation, or whether, in keeping with intrinsicality, such disposition is independent of shape and location.

4 Complete competences

Epistemic virtues come in varying degrees. Retentive memory, visual acuity and deductive powers vary from person to person. It is not enough, however, that one would have a minimal potentiality to retain past events in memory for one's memory to count as relevantly reliable. As it is not necessary to have the memory of Funes to possess a good memory. Thresholds with vague limits are settled so as to select virtues whose proportional natures would serve the cognitive interests of an information-sharing species reliably enough in regard to how things normally stand in the actual world. In this sense, virtues (as well as dispositions in general) are relativized to the circumstances and interests of the epistemic community. However, they are relationally specified by reference to an *objective background* of degrees of potentiality. It is a contingent matter that a particular power to F to a certain degree counts as fragility or visual acuity in our context. But the subject's ability to F to the nth degree is a matter of his or her current disposition and not of some situation for

manifestation.

The same goes for reciprocity. We select and individuate dispositions according to the manifestations to which they give rise. We focus on those particular manifestations –and on the situational conditions for them to occur– because they are salient for the regulation of conduct. However, that does not entail that the potentiality for S to F when in the right condition and situation pops in and out of existence depending on whether the conditions of manifestation are satisfied or not. Much less, that the nature of the power could be fully disclosed by our ways to relationally specify it –as operative in circumstances of manifestation that involve such and such other things.

Consider, for instance, an agent whose visual and recognitional powers would count as reliable in normal circumstances. Locate her intrinsic duplicate in a world where on the base of the same powers, she would act very poorly. Let us say that it is a world with such dense atmosphere that she can only identify a medium sized object from ten inches away. It is clear that in such a world, her performances would be extremely unreliable. It is also clear that according to the criteria used in that world to distinguish competences that are relevantly reliable she would plausibly count as sub-competent. However, her perceptual capacities – with the same proportional nature– are retained. They are not manifested as successful performances in such conditions. But even so, they are, in a restricted sense, exercised by the agent. As a matter of fact, those performances would be adroit, since they are grounded in virtues that, in the actual world, would produce reliable results[40].

In my view, it would be wrong to think that the agent, when she moves from the actual to the alien world, acquires a lesser disposition to form true perceptual beliefs. There is no doubt that her ratio of success is much worse than in habitual circumstances. However, it is the degree of the power, which remains constant, that explains the diversity of results. Likewise, the 'lesser disposition to break' of a glass wrapped in styrofoam is not a new dispositional property of the glass. It is the graduated power of the glass that, along a variety of cases, makes it true the high (or low) ratio of relevant manifestations.

Let us consider, by way of contrast, a competent agent from such an alien world that, in the actual world, would also produce a high ratio of successful performances on the basis of her abilities. Far from saying that the agent acquires a new ability because of the new situation, I would say that her intrinsic virtue has a proportional nature whose gradation in the spectrum of the relevant qualitative power is higher than the gradation of the first subject. The more situations in which the power is manifested, the higher its degree is. The important point is that although such differences among both agents are not context-dependent, it captures a way in which capacities, while intrinsic, are more finely-tuned than the so-called 'general abilities'

---

[40]   For Sosa's conception of adroitness as relative to the reliability in the actual world of a virtue that is held constant across different worlds, see BonJour, L.; Sosa, E.: *Epistemic Justification. Internalism vs. Externalism, Foundations vs. Virtues*, Oxford, Blackwell, 2003, pp. 156-58. This claim is instrumental for virtue epistemologists to meet the challenge of the 'new evil demon' proposed by Lehrer & Cohen (see Lehrer, K.; Cohen, S.: "Justification, Truth, and Coherence", *Synthese*, 55(2), 1983, pp. 191-207). Competent agents located in demon scenarios would thus count as 'internally' as well as 'externally' justified when believing falsely that p. External justification comes from the fact that the agent's virtues would be relevantly successful in the actual world (and in favourable circumstances).
It is thus clear how among the epistemological ramifications of the current ontological discussion, justification occupies a privileged position. On Millar's view, one should claim that lacking the relevant fine-tuned virtue, the demon's victim is not externally justified.

are. Similar considerations support the claim that the duplicates of a virtue whose manifestation involves the co-operation of external factors would retain such context-directed potentiality regardless of the satisfaction of the conditions for manifestation. Again, a fine-grained ability could perfectly well be an intrinsic ability, one that is possessed regardless of the state of the world.

On this model, the context-sensitivity of competences is understood in a weak sense either as the virtue's directedness to a manifestation that involves numerous partners or as the explanation for the diversity of manifestations of the same power along cases and circumstances. However, context-sensitivity is usually conceived in a much stronger sense as the context-dependence of abilities and their possession. To my mind, context-dependence is an artifact of language. There is nothing wrong in referring to dispositions via events external to them (manifestations) or via favorable circumstances. The trouble comes from projecting our pragmatic interests to the world. The source of ontological extrinsicality is thus a picture that, deeply ingrained in our language, ceases to be a mode of representation and comes to be perceived as the represented object.

Millar presents the challenge in a succinct and forceful way. There is something strange in saying that S is exercising her ability to ride a bike, and yet she is not riding a bike[41]. He also provides a diagnosis of why –in the case of recognitional abilities– we tend to think that the same ability is exercised both when our recognitional competences are successful and when they are not. There is, in his view, a similarity due to the fact that, if the performance is unsuccessful, there is at least a judgmental act directed at recognition. However, since the relevant ability is lacking, the appropriate first-order recognitional competence is not exercised at all[42].

The problem for the *success thesis* comes from the fact that it must accommodate our common intuition that the agent retains her ability to F even if there is no opportunity to manifest it, an intuition that, in Millar's view, is perfectly right[43]. This is where the distinction between the *general ability* –the ability *simpliciter*, as dubbed by Millar[44]– and more finely grained abilities is invoked. While the former competence is retained by the subject across good and bad scenarios, it is the fine-grained competence that is lacking when shape and/or situation are not adequate for a successful performance[45]. The obvious questions are: What is the nature of the general ability? How are general abilities related to fine-grained abilities?

The first thing to note is that the above theory is developed in three consecutive stages: explanation, reification, and extension of the dispositional base. First, the ability is invoked in order to explain why the result is not a fluke. Then, it is by

---

[41]   ACF, p. 63. By contrast, that the victim of a demon is exercising the same abilities that would yield knowledge in the actual world and yet she does not have knowledge is a claim that Sosa explicitly endorses. See, among others, BonJour, L.; Sosa, E.: op. cit., p. 161. Insofar as justification is truth-conductive, this claim is required to accord to the beliefs of the demon's victim the right kind of justification.
       As for scholastic questions on whether riding a stationary bicycle would not be a counterexample to Millar's remark, note that the set of abilities exercised in riding a bicycle is different from the set required for riding a stationary bicycle. As different as the abilities for archery hunting are different from skills for archery competition.

[42]   ACF, p. 77.

[43]   *Idem*, p. 80.

[44]   *Idem*, p. 79.

[45]   Note that this distinction corresponds to the distinction between innermost competence and complete competence, such as Sosa makes it.

focusing on a single manifestation-type that one gets an ontological conception of the ability as directed to a single type of manifestations. This conception is, however, replaced by a view according to which the general ability (one that is invariant across multiple scenarios) supports the context-dependent fine-grained ability. Thus, it appears as if the general ability were a multi-track disposition, and the fine-grained ability were one of the tracks of the former. For this model to work, it is required that the *content* of the fine-grained ability is not included in the general power. If such content were included in the innermost competence, the relevant ability would not be lacking in those cases where the performance is unsuccessful. The absence of the ability was, however, the only reason to suppose that S was not exercising it in the bad cases.

There are several options here. The first possibility is to conceive of general abilities as bare potentialities. On this view, what the subject retains in the bad cases is either the physical base to acquire the ability when appropriately situated or the iterated capacity to acquire such skill. It is obvious, however, that this response does not work. The intuition that has to be accommodated is that S retains her ability to *F when the conditions for its manifestation are right*, an ability that she does not lose because of the unpropitious circumstances.

There is, however, a second possibility. The general ability is a multi-track ability. Perhaps, this line of reasoning goes, multi-track abilities are, in turn, *determinable properties*. Consider a standard determinable property: color. An object *o* is in its entirety blue to the exclusion of determinates such as red, yellow, green, etc. However, what makes it true that *o* has a color is a disjunction whose disjuncts are that either *o* is blue, red, yellow, or green, etc.[46] Analogously, the general ability could be a disjunctive property. Thus, it would be multiply realizable. On this model, S retains the general disposition even if it is not instantiated as the fine-grained ability to F. This is so because since having a disjunctive property implies that one of the disjuncts is true to the exclusion of the others, the fine-grained ability instantiates to the exclusion of the other determinates.

The problem is that the previous model is not the right way of conceiving multi-track dispositions. It is right to say that one particular manifestation-type of the relevant competence is instantiated to the exclusion of other manifestations. However, having a multi-track competence means having all the single-track competences that its scope intrinsically includes. The victim of a demon retains her epistemic competence to know even though the hostile situation prevents her from manifesting it as apt believing. The relation between multi-track and single-track abilities is not thus a relation between *distinct* abilities so that the former bare power would support the later competences. Much less, it is a relation such that when poorly situated, S would retain the same determinable property that, were S in the right circumstances, would be instantiated as the determinate ability to know. It is the so-called 'determinate ability' that is retained regardless of the situation. As it was suggested above, this ability is directed to a manifestation in which its conditions are *logically* embedded. It is a power to know *only if* shape and situation are right. As thus, it comes as fully specified independently of extrinsic factors[47].

---

[46] Since it would be true that the object has a color independently of which the color is, the truth-maker for its color is not simply its being blue.

[47] In this, I am following the lead of E.J. Lowe (*cf.* Lowe, E.J.: "How *Not* to Think of Powers: A Deconstruction

In my view, the only way to capture multi-track abilities is by appealing to a *multiple aspect theory* of competences. On this account, the manifold aspects of the multi-track power –the whole range of the power– are logically present in each one of its tracks. In this sense, the multi-track competence would neither be the mere mereological (and contingent) sum of the single-track spokes nor a bare and contentless power that helps to individuate the tracks. Both nominalism (as claiming the former) and quidditism (as a position committed to the existence of bare individuals and powers) would thus be avoided.

Following the foregoing, it would be true to say that, in hostile environments, S exercises the same multi-track power that, if S were in the 'right' circumstances, would produce knowledge. However, since each track is directed to combine with different partners for a particular manifestation-type, there is a sense in which S-in-a-demon-world is not *exercising* that power as long as it is *considered as* the power to combine with the right partners for knowledge. S is exercising the same power but in a different way or according to a different aspect of the competence. Also, S is exercising it to the exclusion of other ways of exercising the multi-track power. This claim follows from the fact that there also is a fit between the competence and the corresponding partners in the bad case, and so the ability contributes to the result as directed to such manifestation –and not as directed to a combination that results in knowledge–. The power to know is one of the aspects of the multi-track power, one aspect that S *possesses* even when she is exercising another aspect of the same overall competence.

There is a grain of truth in Millar's account. However, it is important to bear in mind the following: (i) that all the aspects of the multi-track competence are *possessed* by the agent regardless of which one is being *exercised*, (ii) that the multi-track competence does not exist as an underlying substance, and (iii) that single tracks are fully specified independently of the situation. Even if the exercise of the fine-grained competence requires the appropriate partners, so that it is exercised only when it is successful, its possession is independent of exercise and success. The problem for Millar's account comes from a confusion between a fine-grained ability that is *lacking* (false) and a fine-grained ability that, as such, is not being exercised (true).

It is only when, in the abstraction of other single-tracks, we focus on a particular manifestation as the relevant result to which the competence is directed and as the proper 'success' of such ability that the triple-S structure of complete competences emerges. As a contribution to the theory of the role played by epistemic competences in regard to the alethic interests of communities and to the normal circumstances in which their cognitive language-games are rooted, complete competences are invaluable. This should not blind us, however, to the danger of considering a partial perspective as the complete picture of epistemic dispositions and of projecting a mechanism of pragmatic selection into ontology. This is why proper attention should be paid to how Sosa qualifies his analysis as a theory of competences *understood as abilities to succeed*[48]. So understood, Sosa's account is able to disclose the finer

---

of the 'Dispositions and Conditionals' Debate", *The Monist*, 94(1), 2011, pp. 19-33). Lowe's proposal is able to capture the 'internal relation' between the manifestation of a disposition and the conditions for it to occur. To my mind, this model is more congenial with the metaphysical analysis of knowledge proposed by Sosa than the standard analysis in terms of triggers whose relation to manifestation is causal, and so external and contingent.

[48]   J&A, p. 95.

structures of potentiality required for a dispositionalist theory of knowledge to be successful without compromising the metaphysical framework that underpins such a theory.

A virtue epistemology is, thus, compatible with partial considerations that, zooming in on an aspect, make it possible to capture fragments of a pattern that extends through the whole picture.[49]

5. References

BonJour, L.; Sosa, E.: *Epistemic Justification. Internalism vs. Externalism, Foundations vs. Virtues*, Oxford, Blackwell, 2003.

Fara, M.: "Masked Abilities and Compatibilism", *Mind*, 117(468), 2008, pp. 843-865. doi: http://doi.org/10.1093/mind/fzn078

Haddock, A.; Millar, A.; Pritchard, D.: *The Nature and Value of Knowledge: Three Investigations*, Oxford, Oxford University Press, 2010.

Heil, J.: *The Universe as We Find It*, Oxford, Clarendon Press, 2012.

Johnston, M.: "How to Speak of the Colors", *Philosophical Studies*, 68(3), 1992, pp. 221-263. doi: http://doi.org/10.1007/BF00694847

Kallestrup, J.; Pritchard, D.: "Dispositional Robust Virtue Epistemology versus Anti-luck Virtue Epistemology", in M.A. Fernández Vargas (ed.), *Performance Epistemology: Foundations and Applications*, Oxford, Oxford University Press, pp. 31-50.

Lehrer, K.; Cohen, S.: "Justification, Truth, and Coherence", *Synthese*, 55(2), 1983, pp. 191-207. doi: http://doi.org/10.1007/BF00485068

Lewis, D.: "Finkish Dispositions", *The Philosophical Quarterly*, 47(187), 1997, pp. 143-158. doi: http://doi.org/10.1111/1467-9213.00052

Lowe, E.J.: "How *Not* to Think of Powers: A Deconstruction of the 'Dispositions and Conditionals' Debate", *The Monist*, 94(1), 2011, pp. 19-33. doi: http://doi.org/10.5840/monist20119412

Manley, D.; Wasserman, R.: "A Gradable Approach to Dispositions", *The Philosophical Quarterly*, 57(226), 2007, pp. 68-75. doi: http://doi.org/10.1111/j.1467-9213.2007.469.x

Martin, C.B.: *The Mind in Nature*, Oxford, Clarendon Press, 2008.

McKitrick, J.: "A Case for Extrinsic Dispositions", *Australasian Journal of Philosophy*, 81(2), 2003, pp. 155-174. doi: http://doi.org/10.1080/713659629

Millar, A.: "Abilities, Competences, and Fallibility", in M.A. Fernández Vargas (ed.), *Performance Epistemology: Foundations and Applications*, Oxford, Oxford University Press, 2016, pp. 62-82.

Molnar, G.: *Powers: A Study in Metaphysics*, Oxford, Clarendon Press, 2003. Pritchard, D.: *Knowledge*, Basingstoke, Palgrave Macmillan, 2009.

Pritchard, D.: "Anti-luck Virtue Epistemology", *The Journal of Philosophy*, 109(3), 2012, pp.
247-279. doi: http://doi.org/10.5840/jphil201210939

Sextus Empiricus: *Against the Logicians*, Cambridge, Cambridge University Press, 2005. Smith, A.D.: "Dispositional   Properties", *Mind*,   86(343), 1977, pp. 439-445. http://doi.

---

org/10.1093/mind/LXXXVI.343.439

Sosa, E.: *Knowing Full Well*, Princeton & Oxford, Princeton University Press, 2011. Sosa, E.: *Judgment and Agency*, Oxford, Oxford University Press, 2015.

Vihvelin, K.: "Free Will Demystified: A Dispositionalist Account", *Philosophical Topics*, 32 (1/2), 2004, pp. 427-450. doi: http://doi.org/10.5840/philtopics2004321/211