

## Mistakes as revealing and as manifestations of competence

Felipe Morales Carbonell

**Abstract** The final chapter of Elgin's *True Enough* (2017) defends the claim that some mistakes mark significant epistemic achievements. Here, I extend Elgin's analysis of the informativeness of mistakes for epistemic policing. I also examine the type of theory of competence that Elgin's view requires, and suggest some directions in which this can be taken.

**Keywords** understanding, mistakes, informativeness of mistakes, epistemic policies, competence

Knowledge and error flow from the same mental sources, only success can tell the one from the other. A clearly recognized error, by way of corrective, can benefit knowledge just as a positive piece of knowledge can.

ERNST MACH, *Knowledge and Error*

### 1 Introduction

The role of error in the advancement of understanding is one of the transversal themes of Elgin (2017). Her discussion of the topic culminates in chapter 14, on 'Fallibility's Payoff', where the argument is presented in an extended form.<sup>1</sup> There, she tries to make good of the idea that an epistemic agent can recognize that, despite their evidence, they might be mistaken (*fallibilism*). When this thought is applied to the case of knowledge, it leads to an apparent dilemma between skepticism and a disguised

---

Felipe Morales Carbonell  
Center of Logic and Philosophy of Science, Institute of Philosophy, KU Leuven  
Kardinaal Mercierplein 2 – bus 3200  
BE-3000 Leuven / Belgium  
E-mail: felipeandres.moralescarbonell@kuleuven.be

<sup>1</sup> A previous version of the material appears in Elgin (2012).

form of dogmatism; the suggestion is that we should think of fallibilism in terms of understanding. The way she does this is by giving an account of the role of error in the advancement of understanding; fallibilism becomes an active stance of epistemic humility. Along the way, Elgin argues that, on the basis of a form of *holism* about the individuation of beliefs, mistakes are marks of epistemic achievement. But this seems to require a revision of the notion of *competence*, which has been deployed to great effect in recent epistemology (see, for example, Sosa (2015)): if, like Elgin claims, mistakes can be the mark of epistemic achievements, and if those are standing competences, it follows that some manifestations of the relevant competences do not constitute successes. Elgin's view seems to require some weaker link between competences and success than a *naive* approach suggests. How much weaker is an open question.

I proceed as follows: in section 2 I give an overview of Elgin's account of intellectual humility and the importance of error to understanding. My reconstruction introduces a distinction about the types of policy one may take towards error before and after it arises, that I think needs to be added to Elgin's distinction between actual and potential error in order to capture a complete picture of the different ways in which mistakes can be relevant to epistemic policing. In section 3 I examine in more detail a) the way in which the information that can be extracted from mistakes can be used in the application of the types of policies I describe in the previous section, and b) Elgin's argument for the reliability of the informativeness of mistakes. I suggest that it is possible to defend the overall view using slightly different premises than Elgin uses. I also raise the problem of individuating mistakes, and I connect the issue with some recent work in epistemic logic. Finally, in section 4, after I briefly consider the viability of defending the argument for the idea that mistakes signal achievements in independence of Elgin's holistic presuppositions, I consider the line of thought I sketched above about the consequences of having failures as manifestations of competence for the use of the latter notion in epistemology, examine what kind of theory of competence Elgin's view requires, and suggest some directions in which the theory of competence can be taken; in particular, I will suggest that we have to consider as epistemically significant competences that are in some sense neutral to success.

## 2 Fallibilist humble understanding

Fallibilism, in Elgin's proposal, should be characterized in terms of a kind of (complex) stance a subject can take towards a subject matter and the subject's own understanding of that subject matter. This stance, she says, is a form of intellectual humility, for which it should be at least necessary some recognition of epistemic limitations (actual or possible).<sup>2</sup> Elgin's conception of the fallibilist stance, however, enriches this bare conception of intellectual humility by requiring, besides a standing

<sup>2</sup> This rough characterization coincides in the main with Whitcomb et al (2015), who say that intellectual humility is 'having the right stance towards one's intellectual limitations' (p. 516), where 'the right stance' is 'to be appropriately attentive to [those limitations] and to own them'. Church (2016) argues that epistemic humility should also be sensitive to one's strengths, and proposes that intellectual humility is 'the virtue of accurately tracking what one could non-culpably take to be the positive epistemic status of one's own beliefs' (p. 427). Similar points are made by Kallestrup and Pritchard (2016), and by Tanesini

disposition to believe/outright belief that one can be wrong (as an indicative alternative), a set of dispositions or prescriptions to act on that possibility in specific ways, which I will describe in a second.<sup>3</sup> Elgin's epistemically humble fallibilist agents (fallibilist agent's, for short) should not simply accept the possibility that they might be mistaken, but also aim towards epistemic amelioration.

This suggests that fallibilist agents should adopt policies of error management (and make an effort to gain the dispositions to apply them). Elgin distinguishes two types of policies, depending on whether the error(s) in question are *actual* or *possible*. She argues that fallibilism has its real pay-off in consideration of the second case, a point she thinks has been under-appreciated.

To expand on the point, I think it can be helpful to introduce a distinction (which is implicit in Elgin's view, but that she does not exploit directly) between what I will call *ex ante* and *ex post* policies of error management. The former apply before someone has fallen into error (in which case they are *objective ex ante* policies) or has been found to have fallen into error (in which case they are *epistemic ex ante*). For example, one can have an *epistemic ex ante* policy to look for errors in our epistemic commitments. In fact, this is a policy that it seems one *must* take unless one adopted the stance to wait for errors to manifest themselves. A norm that prescribed the uptake of beliefs only in the condition that they are knowledgeable could constitute an *objective ex ante* policy, insofar as part of its rationale could be the minimization of error. While *objective ex ante* policies will be policies of error *avoidance*, *epistemic ex ante* policies will rather oftentimes be policies of error *detection*. On the other hand, *ex post* policies tell us what to do once an error has been found:<sup>4</sup> for example,

---

(2018), who argues that epistemic humility is a complex attitude involving modesty (about one's epistemic successes) and self-acceptance (of one's epistemic failures).

<sup>3</sup> The disjunction prescriptions/dispositions signals two variant reconstructions of the fallibilist stance. In one, fallibilism is characterized by the belief in certain prescriptions on what to do if one falls into error; in this case fallibilism can be identified with a doctrine or higher-order belief about one's understanding. In the other, fallibilism is rather characterized by the possession of some dispositions to act in certain ways relevant to the management of error. Elgin herself says that her argument shows that fallibilism 'is not a higher-order stance toward our understanding of a topic', but that it 'is, or should be, woven into the fabric of understanding' (p. 309), which suggests she prefers the dispositional variant. However, for every dispositional account there will be a doctrine that asserts that understanding should be as if it was determined by the disposition, and *vice versa*. What favours the dispositional view is that it seems to provide a natural solution to the issue of motivating the agent to act as the norms prescribe: the disposition would presumably be expected to intrinsically lead the agent to act in the required way given certain triggers. A purely dispositional account, however, seems problematic: shouldn't one be further disposed to be so disposed (and so on), falling into a regress? How can one even adopt a stance, if so? I would suggest formulating fallibilism in such a way that it entailed the acceptance of the doctrine that one should ameliorate one's understanding's management of error, including the disposition to ameliorate it. Then, we could appeal to more general rational competences for the uptake and application of the relevant norms (this does not necessarily break the potential regress of dispositions, but the fact that the move doesn't require the immediate supposition of *ad hoc* capacities could make it more palatable). It should be noted that the strategy of appealing to more general capacities can be pursued in such a way that no specific dispositions are involved in the description of a fallibilist stance: instead, general dispositions could drive acting according to the fallibilist prescriptions.

<sup>4</sup> In the *ex post* case, the distinction between objective and epistemic policies is difficult to sustain, because while there could be norms that prescribe what to do once one has fallen into error, these cannot be applied unless one also recognizes or comes to believe that one has indeed fallen into error. So, at best, *objective ex post* policies could coincide with *epistemic ex ante* ones.

one could have an *ex post* policy to close any gap one finds in one's understanding. *Ex post* policies will be policies of *amelioration* or *repair*.

To see how the *ex ante/ex post* distinction can help us grasp the character of Elgin's fallibilism, it can be helpful to observe how the distinction mirrors one Rescorla (2009) assumes between norms that apply to acts before their performance (for example, norms that specify how the act should be constituted or performed), and norms which are relevant after it has been performed (which can handle the consequences of the acts, for example). He introduces this point to defend what he calls a non-restrictive dialectical model of assertion, according to which 'constitutive assertoric norms constrain how I must react if someone challenges my assertion, but they do *not* constrain what I should assert in the first place' (p. 100). In our terminology, his view is that there are no (objective) *ex ante* policies for assertion, but only *ex post* dialectical policies of assertion management (assuming that assertions are something like updates to a shared conversational background). How does this apply to Elgin? While she does not outright deny the importance of norms of error avoidance, her fallibilism reduces the stress that might otherwise be put on them. This is important in consideration that her doctrine of error management has as its background her use of the notion of acceptance (cf. Cohen (1992)). Since acceptance is 'always provisional and revocable' (Elgin (2012)), the importance of *ex ante* norms of acceptance is diminished; at the very least, the necessity of having very strict *ex ante* norms is weakened. This, however, does not mean that our attitude towards mistakes should not be forward-looking.

One of the most important insights in Elgin's account is that some errors matter even if they are not (or will not be) realized. Consequently, it is necessary to account for possible or potential error. Certainly, *epistemic ex ante* policies presuppose that there is a salient type of possibility of error. In that case, however, the relevant possibility is merely epistemic: it is the possibility of unknown but *actual* error, and more precisely, *present* error. However, potential error can be salient even if we have no reason to think that we are presently at error. An important example is that we are often concerned with *future* error. We can construe this case as a different example of what *objective ex ante* policies should be concerned about. It should be noted that not all the policies concerning (potential) future error are policies of avoidance. In some cases, for example, it could be unreasonable to think that we can avoid certain future errors. In such cases, our policies should be directed rather to ameliorate our capacity to deal with the fallout of error, that is, they should be directed to the setup of *ex post* policies or to ameliorate our capacity to follow through already established *ex post* policies (these policies of amelioration are themselves *ex ante*). Concern about future error motivates the adoption of higher-order policies. This, in turn, makes the stance of an agent who is concerned with the possibility of error self-improving and self-reinforcing.

The point generalizes. Policies like those I have just described can be driven by concern with potential future error even if there is no guarantee that we will effectively fall into error. Instead of a certainty of error, there might simply be a *risk* of it. Identifying epistemic risk and its grounds is, as it is generally accepted, an important

epistemic project.<sup>5</sup> The risk of error can motivate the adoption of additional *ex ante* policies of amelioration, in this case not of our *ex post* capacity to deal with errors, but of the current state of our epistemic commitments. Consequently, these policies aim to minimize risk.

The type of possibility at play can be even broader. It can be fruitful to examine weaker senses in which errors are possible even if there is no risk of error (for example, if the probability of error is not high). For example, Elgin says we should also look for the weak points ('potential fault lines', p. 299) in our network of commitments. Something can be weak even if it is not at risk. The relevant possibilities in this case are objective, rather than purely epistemic (our original description of *epistemic ex ante* policies should already cover that case). Elgin's notion of vulnerability (p. 307) captures a similarly objective type of possibility.

Are we still, strictly speaking, dealing with mistakes rather than with other diverse forms of epistemic failure? Take the case of vulnerability. In this context it is plausible that it should be conceived primarily as a property of (parts of) networks of commitments. However, these networks are not *only* vulnerable to containing errors: for example, insofar as we rely on our commitments to solve novel problems, they can also be vulnerable to becoming stagnant or infertile. Conceptually, concern with vulnerability does not depend on concern with error. The issue also appears from a different angle. Some policies that Elgin promotes, like double-checking and trying out different perspectives, can be motivated independently by concern with generality or simplicity, which appear to be independent of concern with error.<sup>6</sup> However, this doesn't undermine the relevance of these possibilities for mistakes, however. If we think our epistemic state can be impacted negatively by a certain procedure, it is natural to think that it will be a mistake to follow through with it. The relevant possibilities are features of whatever allows the identification of mistakes; and mistakes, Elgin argues, can only be identified in reference to networks of commitments (I will come back to this form of holism in a second.)

### 3 What mistakes reveal

Concern with error can motivate the adoption of a wide variety of types of epistemic policies. It is important to observe that the content of these policies is sensitive to the properties that errors may have. In Elgin's presentation, the relevance of errors for the application of epistemic policing mainly comes down to the thesis that

---

<sup>5</sup> Pritchard (2017) points out that risk assessment is forwards looking, whereas luck assessment seems to be backwards-looking. Elgin observes that being merely able to assess some unlucky event as such is less good than also being able to identify the factors of risk involved, and is the sign of a worse understanding (p. 298). This emphasizes the importance of adopting certain *ex ante* policies (although in the example Elgin discusses the relevant policies are not concerned with error, but failures of other kinds, like the manifestation of side effects).

<sup>6</sup> In mathematics, it is common to try to prove theorems indifferent ways. Polya (1957) emphasizes this a crucial heuristic for discovery and understanding. In these cases the mathematician is often sure that his initial proof is correct. Nevertheless, they may still be tracking some notion of epistemic failure, as they have come to believe that their commitments are somehow incomplete or otherwise inadequate (in some cases by aesthetic criteria). On this point, see also Elgin (2002).

**(Revealing Revisions)**

Mistakes can be revealing about what commitments should be revised once an error is found

She discusses at length how mistakes can be informative about what beliefs should be dropped besides the ones which have been identified as mistaken and triggered the revision. The content of these policies of revision is doubly dependent on the mistakes. For example, in the minimal case, if our antecedent commitment to  $\phi$  is mistaken, we should (at some point) drop our commitment to  $\phi$ . The prescription depends on the mistake generically (as the type of action it prescribes, in this case to drop the commitment) and more specifically (given the type of action prescribed, it fixes the particulars involved in the prescribed action; in this case, that the commitment to drop is the one that has been found to be mistaken). The policies to adopt can be significantly more complex than this example, as we have seen.

There are more ways in which mistakes can be revealing of the epistemic commitments a subject has and their properties. These would in turn be relevant to further types of policies. More generally, then, we should discuss the conjunction of

**(Revealing)**

Mistakes by some agent  $a$  can be revealing of non-trivial properties of  $a$ 's network of commitments and its contents.

**(Relevance to Policy)**

The properties revealed by mistakes can be non-trivially relevant for the application of epistemic policies (for example, policies of error management).

(Relevance to Policy) plays two roles here. It entails that, if an agent commits to policies which are sensitive to certain information that can be extracted from mistakes, they should also commit to policies that aim to extract that information. This is crucial to the argument that we should pay more attention to the informativeness of mistakes. The thesis also acts as a constraint on what version of (Revealing) should be accepted. Suppose we are antecedently committed to certain policies which are sensitive to some information about our epistemic commitments. Then, in order to make those policies sensitive to the upshot of mistakes, we need to show that mistakes can provide that information. A fully spelled out version of this idea would have to account for the link between the information revealed and the policies:

**(Revealing\*)**

Mistakes can be revealing of a significant portion of the types of properties of networks of commitments and its contents which are relevant to epistemic error management policies.

(The meaning of 'a significant portion' is left vague because we don't have a principled way to fix the class of possible policies we should care about. Regardless, if we replace 'a significant portion' with 'all' the resulting thesis seems too strong.) Note that (Relevance to Policy) is entailed by (Revealing\*); strictly speaking, (Revealing\*) is stronger than the conjunction of (Revealing) and (Relevance to Policy).

As we will see in more detail in a second, Elgin also argues in favour of

## (Revealing Success)

Mistakes can be revealing of epistemic achievement.

For now, it is important to note that the relevant sense of epistemic achievement is formulated in terms of properties of revealed networks of commitments, so besides giving an additional instance of (Revealing), her argument in favour of this thesis is important for the support of (Relevance to Policy) and consequently of (Revealing\*). This mirrors the conceptual point that the assessment of success/achievement (in our case, from the occurrence of mistakes) is derivative from the assessment of more basic properties (in our case, of the commitments revealed by mistakes), that is, that attributions of success require specific grounds. The informativeness of mistakes does not reduce to their revealing of achievement; rather, the latter is an aspect of the former. We can even see how (Revealing Success) could in some cases be an instance of (Revealing\*): perhaps that one has realized an epistemic achievement of a particular kind is informative in ways which are relevant to policies of error management. For example, that one has managed to provide a partial proof of a theorem shows that one's general logico-mathematical commitments are at least sufficiently robust, and suggests, as a constraint on the policies, that those commitments should not be revised wholesale even if the proof as a whole has been found to be mistaken. Of course, one could still ask whether some of those commitments were nonetheless faulty, and if so, which.

One could think that the cases of learning from mistakes that Elgin examines provide support for (Revealing\*), but they don't show by themselves why we should *expect* that thesis to be true. My impression, and I think Elgin's, is that the relevance of the information mistakes provide for epistemic policies is not accidental, even though there are no guarantees that this information will be useful. One way to address this point is to make the support for (Revealing\*) come from inductive generalizations from past failures. In this case, the adoption of policies of information gathering from mistakes is prudential and preemptive: given that in some cases the information provided by mistakes can be revealing of what commitments to revise in non-trivial ways, we should adopt a general policy to extract information from mistakes. In this case, however, it might turn out that the fact mistakes can sometimes be informative in the required sense is a matter of luck. I think that Elgin's own argument in favour of (Revealing Success) shows that the ground for the possibility of informativeness of mistakes can be more reliable than sheer chance.<sup>7</sup>

The argument Elgin presents for (Revealing Success) is deceptively simple. We can reconstruct it as follows:

- (1) It is not possible to make a mistake without already holding beliefs about the topic one makes the mistake about. (Mistakes require beliefs)
- (2) To hold a belief one requires having a network of commitments that satisfy certain conditions *C* of reasonable acceptability and true-enough-ness. (Belief requires Understanding)

<sup>7</sup> Cf. Elgin (2017, p. 306): '... sometimes, we get things wrong in epistemically fruitful ways. Once discovered, such errors provide not just incentives but resources for serious, focused, effective inquiry. By revealing not only that but also where we have got things wrong, they point us in the direction of advancing our understanding. We are lucky we are disposed to make such mistakes.' If I'm right, the appeal to luck here should be treated at least in part as rhetorical.

- (3) To have a network of commitments that satisfy  $C$  constitutes an epistemic achievement.  
(Understanding as Achievement)

It follows that making a mistake entails a form of epistemic achievement, from which (Revealing Success) follows.<sup>8</sup>

The premises are fairly rich in consequences. (1) and (2) together imply that mistakes require the possession of a (relevant) network of beliefs, so that the recognition of a mistake should be informative about the properties of the underlying web of beliefs and other facts about them. Some of these are trivial, such as the fact that there are such networks; indeed, on the supposition that mistakes are themselves beliefs, this follows immediately. However, there are reasons to refuse to make this supposition: for example, it seems plausible that we can be mistaken to take epistemically significant attitudes which are not beliefs, starting from weaker forms of acceptance. If we allow for this possibility, it is no longer clear that every kind of mistake should have the possession of beliefs as its ground. For example, to mistakenly weakly accept something could ‘merely’ require the possession of a network of weak acceptances. In any case, a corresponding revision of (1) shouldn’t affect the overall strategy significantly, since it should be possible to adjust the second (and third) premises accordingly.

As it stands, however, perhaps it would be better to interpret the assumption Elgin makes here as a form of heuristic principle. It seems likely that (regardless of the kind of attitudes that mistakes consist of) the types of question about the supporting attitudes could be transferred to the case of the relevant mistakes. This is obviously true when the type of the supporting and supported attitudes matches. For example, if we can ask what further commitments give us reason to take those attitudes, we should be able to ask the same about the corresponding mistakes. Even if we might not know initially what to ask about the mistakes, we could repurpose admissible questions of the types we know are admissible for the presumed supporting attitudes. Reflection on the answers to the questions generated in this way should help in making explicit the networks of commitments or other attitudes uncovered by the recognition of mistakes (at least to some degree). The information we can gather about the supporting networks of beliefs will depend on our antecedent thoughts on the nature of mistakes, and in turn, the reliability of the informativeness of mistakes will be sensitive to the accuracy of those commitments. On this score, Elgin’s assumptions seem fairly robust. It is important to note that Elgin’s assumption (2) derives from her view on the individuation of beliefs:

(Davidson/Elgin Individuation of Beliefs)

Beliefs are individuated by their position in a rich, textured array of relevant beliefs which are true enough and reasonably accepted.

If we accept this as true, and additionally assume (1), we should expect the transference procedure described above to be reasonably reliable. This result can resist

<sup>8</sup> Write ‘Mst’ for ‘s makes a mistake about t’, ‘Bst’ for ‘s holds beliefs about t’, ‘NsC’ for ‘s has a network of commitments which satisfy condition C’, and ‘As’ for ‘s has an epistemic achievement’. Then, (1) is  $\neg\Diamond(Mst \wedge \neg Bst)$ , (2) is  $\neg\Diamond(Bst \wedge \neg NsC)$ , and (3) is  $NsC \rightarrow As$ . From (1), it follows that  $Mst \rightarrow Bst$ , and from (2) it follows that  $Bst \rightarrow NsC$ . Then, by iterated application of transitivity of  $\rightarrow$ ,  $Mst \rightarrow Asc$ .



some variations on the premise. For example, it could be tempting to state the relevant individuation principle strictly in terms of acceptances, or to say (at least) that the individuation of beliefs depends on their location in networks of acceptances. In either case, it is arguably the *centrality* of the notions of belief and acceptance in the network of epistemically-significant concepts that guarantees the reliability of the procedure: epistemically significant questions will often entail questions about belief and/or acceptance.<sup>9</sup> This suggests that the heuristic can be defended without necessarily appealing to a holistic view of belief or acceptance individuation (for example, by relying instead purely on the conceptual centrality assumption), although in that case I suspect a metaphysical account of mistakes will still be needed.

A reason to insist in this last point is that the dependency relations appealed to in a framework like Elgin's have the ability to at least partially explain the potential relevancy of the informativeness of mistakes for epistemic policing. Consider the case of *ex post* policies. In the simplest case, the heuristic will make admissible questions about the support of mistakes: what assumptions (beliefs/acceptances) did they have their base on, what consequences do they have, what other beliefs do they provide support for, and so on. If the answers to those questions are not immediately available, one could then proceed to try to make them available, or at least find explanations of why those questions cannot be answered from one's current position. This gives a second type of questions whose admissible solutions are parametrized by the information provided or made available for inquiry by mistakes; these questions correspond to a new set of epistemic policies.<sup>10</sup> In this way, we can see that Elgin's assumptions entitle her to hold on to more than (Revealing Revisions): indeed, they seem to provide good reasons to accept at least a limited version of (Revealing\*). At the base of all this lie the supposed dependency relations, which give guarantees about the effectiveness of the heuristic.

Does the support for (1) derive from a similar thought about the individuation of mistakes? If so, something more needs to be said about the conditions in which mistakes are recognized and identified. Elgin briefly touches upon this point when she describes cases where mistakes proved to be informative. The picture she presents in those cases is that mistakes are found when someone finds inconsistencies between, for example, predictions and experimental results or observations. However, the recognition of inconsistencies does not by itself individuate mistakes. In a sense, we should think of the detection of mistakes as progressive: first, we find that *something* is wrong, then we try to figure out *what* it was (or wasn't). To look for mistakes

---

<sup>9</sup> Whether this is acceptable will depend on how we view the structure of epistemic concepts. Other frameworks will have different concepts (like knowledge) as central. However, this does not necessarily undermine the heuristic. If the notions of belief and acceptance exist at all in a framework, the heuristic could still be robust if their place was at least relatively central. A different option would be to distinguish between centrality and 'fundamentality': then, we could claim that, while belief or acceptance are strictly central, other notions are more fundamental. Yet another option would be to limit the scope of the heuristic.

<sup>10</sup> Sosa's (2011) conception of levels of epistemic evaluation could offer a complementary explanation. If mistakes are conceived as a form of performance, the evaluative language associated to performances can be transposed to their case, just like Sosa does in the case of belief. For example, we can consider whether mistaken beliefs or acceptances are inaccurate (as opposed to accurate), incompetent (as opposed to adroit) or inept (as opposed to apt); note that in this case, the severity of the mistakes decreases as we go from inaccuracy to incompetence, and then to ineptitude.

is to look for what *generated* or is *responsible for* the inconsistencies. What underlies (1) is the assumption that in many cases this is precisely the possession of belief-like attitudes. When a prediction is inconsistent with observations we hold on to these and look for the mistake in the grounds for having made the prediction (assuming, of course, that the observation is trustworthy). In other cases it is less clear where we should look for mistakes; in the worst case scenario, only the whole situation where the mistake is recognized can be fully responsible for the inconsistency.<sup>11</sup> Elgin is right in pointing out that an unrestricted holism will make of every case of inconsistency such a *pessimum*; implausibly, such a view cannot distinguish between coexisting mistakes. A more plausible solution to the problem is to say that mistakes are constituted by something like the smallest set of beliefs that yields the inconsistency. In either case we should say that the individuation of mistaken beliefs requires the individuation of further beliefs or acceptances which are at least partially responsible for the inconsistencies that trigger the process of detection of mistakes.

### 3.1 Some more aspects of *ex post* norms

The issue we have just raised can be made more clear if we try to give more precise models of *ex post* policies; I will now try to briefly sketch some important points, that I hope can be illuminating.

Assuming, for simplicity' sake, that our network of commitments ( $\Gamma$ ) can be represented as a set, and using  $U_\Gamma(\alpha, \beta)$  to represent an arbitrary update to our network of commitments with the sets of additions  $\alpha$  and deletions  $\beta$  (such that the updated network  $\Gamma' = \Gamma \cup \alpha \setminus \beta$ ), and  $[_M\phi]\psi$  to formalize that we should do  $\psi$  after  $\phi$  has been detected as a mistake,<sup>12</sup> the policy to immediately drop isolated mistaken commitments can be written schematically as:

(Drop)

$$[_M\phi]U_\Gamma(\emptyset, \{\phi\})$$

A policy  $P_2$  is consistent with another policy  $P_1$  iff following  $P_2$  does not prevent one to succeed in following  $P_1$ . Distinct policies will be more or less conservative depending on the size of the changes they will entail (removing two items from  $\Gamma$  is more conservative than removing two items and adding a dozen; in most cases judgments about the size of the changes will be much more difficult to assess, however). For example, the following policies are consistent with (Drop) but less conservative:

<sup>11</sup> This is closely connected to the Quine-Duhem problem of underdetermination. Unrestricted closure endangers the possibility of determining where errors lie. Cf. Duhem (1954, p. 187): “[...] what the experiment declares stained with error is the whole group of propositions accepted by Newton, and after him by Laplace and Biot, that is, the whole theory from which we deduce the relation between the index of refraction and the velocity of light in various media. But in condemning this system as a whole by declaring it stained with error, the experiment does not tell us where the error lies.”

<sup>12</sup> For our purposes,  $[_M\phi]\psi$  could be defined in terms of the  $[\phi!]\psi$  operator of public announcement logic (PAL), as the special case ‘after a public announcement that  $\phi$  is mistaken, it should be that...’. In a more general setting, we need to distinguish between this and the case where a mistake has been detected but no announcement has been made. An overview of PAL and related formalisms can be found in van Ditmarsch et al (2007).

(Drop with Deletion of Implications)

$$[{}_M\phi]U_\Gamma(\emptyset, \{\psi : \phi \rightarrow \psi\})$$

(Drop with Deductive Closure)

$$[{}_M\phi]U_\Gamma(\emptyset, \{\psi : \psi \rightarrow \phi\})$$

(Drop with Deletion of Implications and Addition of Negative Implications)

$$[{}_M\phi]U_\Gamma(\{\rho : \neg\phi \rightarrow \rho\}, \{\psi : \phi \rightarrow \psi\})$$

Because following (Drop) seems to be the simplest correct policy to follow given a mistake (that is, no correct policy is more conservative than it), I will consider the set of all the policies consistent with (Drop) as the set of all the policies permissible given that mistake. The question is, then, which of these permissible policies should be followed, or equivalently, which closure principle should be adopted for defining the update rule. The selection of policies is a function of the information in the mistake and the network of commitments. It should be noted that the corrections prescribed in these policies are in some sense sufficient to individuate what the mistakes consist of: upon the realization of an inconsistency, the supporting mistakes are precisely what should be corrected (that is, nothing that should be corrected is not a mistake).<sup>13</sup>

Elgin's account suggests that (Drop) is not a universal rule. The less conservative policies also introduce problems. First: as long as the conditional in the closure conditional is classical, updates can be at least partially irrelevant. Suppose we find we are mistaken about the date of a friend's birthday. If we followed (Drop with Deletion of Implications and Addition of Negative Implications), we would be prescribed to add to our network of commitments things like 'the date of our friend's birthday is not that or the moon is made of cheese'. Furthermore, there will be an infinity of such things.<sup>14</sup> Second: there is an implicit assumption that the vulnerability of the network is homogeneous. On the one hand, any point in the network is equally up for grabs, as it were, and further, any mistake will take down as many linked commitments as it logically can. The first issue could be solved by adopting a relevance-friendly conditional. The second seems to require closure principles to be sensitive to the structure of  $\Gamma$ . A promising direction is to make use of a topic-sensitive logic (like in Lewis (1988) and the more recent Yablo (2014)), since it may provide the resources to handle both issues at once (the mereological notions introduced there are structural and seem to capture an important notion of relevance). Berto and Hawke (2018) adopt this type of strategy in their discussion of relative knowability.<sup>15</sup> The connection is not coincidental: part of the problem of the informativeness of mistakes comes from characterizing what is knowable relative to the presence of mistakes.

<sup>13</sup> An alternative way to develop the semantics for  $[{}_M\phi]U_\Gamma(\psi)$  could be to use some form of exact truth-maker semantics, in the way that Fine (2017) does in its treatment of imperatives: the update action  $U_\Gamma(\psi)$  should be in exact compliance with the announcement that ' $\phi$  has been found as a mistake'. Note that the compositional story is more complex than in the imperative case: for example, an update in compliance with  $[{}_M\phi \wedge \psi]$  is not equivalent to an update in compliance with both  $[{}_M\phi]$  and  $[{}_M\psi]$ : that a conjunction is mistaken does not entail that the conjuncts are independently mistaken.

<sup>14</sup> The prescription would be impossible for us to apply fully, so it can be tempting to judge it is poor or deficient in some sense as an action-guiding principle (if it is one at all). Even so, it still seems possible to *try* to follow it. A more difficult situation would arise if even that was prevented; in that case the prescription would not provide any sort of action-guidance.

<sup>15</sup> Cf. also Hawke (2017).

A further issue, raised by Harman (1986), is that belief revision appears to be sensitive to interests and other practical reasons.<sup>16</sup> These can be temporary: Harman gives the example of jurors who, having reached a verdict, are entitled to forget the reasons for it. In the case of mistakes, it seems that there are admissible policies where one extracts information from the mistakes in order to reach a conclusion, but then only retains the conclusion. This suggests that the structure of the policies matters. (Drop) *et al* are assumed to apply immediately. In practice, the policies we may want to follow can take several steps that should be followed in order, and we will drop the mistakes only once we have extracted from them all the information we need (policies consistent with (Drop) can branch out, but they will always converge in the deletion of the mistake). This suggests a further type of restriction on what policies are admissible. The information mistakes provide should be treated as a finite resource, at least in the sense that some of it might become unavailable after taking certain actions; if this is a concern, the policies to follow should drop information as late as possible. Additionally, the fact that information has been extracted or corrected should be treated as additional information, since it will affect the structural integrity of the network of commitments; for example, we might expect commitments which have been revised already to be less vulnerable than an arbitrary non-revised commitment. Our policies of error management affect how steadfast we should be against challenges.

#### 4 Mistakes manifesting competences

As I sketched above, Elgin's argument leads to the idea that mistakes can be marks of achievement. If the networks of beliefs/acceptances underlying meaningful mistakes satisfy the conditions of acceptability that Elgin suggests, they should constitute (in her framework) some form of understanding. It is commonplace to think of understanding as an epistemic state that 1) has distinctive epistemic value, and 2) is the result of epistemic performances. The two ideas combined yield the idea that having understanding is a form of achievement. However, note that in absence of independent reasons for (1), commitment to (2) could 'pump in' the intuition that understanding is valuable. A way this could happen is if we accepted that the constitution of networks of commitments requires epistemic *effort* which deserves to be recognized.<sup>17</sup> This has some problems. It is unclear that the development of networks of commitments can only come to happen through the exercise of effort in any sense that would make us personally responsible for it (sub-personal processes might be sufficient). If it is not required, the evaluation of our epistemic standings as achievements seems to become trivial or at least extremely liberal.<sup>18</sup> In any case, Elgin's argument relies on a different idea: the value of understanding lies in that the acceptability condition

<sup>16</sup> The theme is explored further in Friedman (2017).

<sup>17</sup> On the notion of effort, see Von Kriegstein (2016).

<sup>18</sup> This liberalization of the conditions that could constitute achievement does not prevent having a consistent account of the magnitude of epistemic achievements that fits adequately with the ordinary intuition that achievements are in some sense uncommon: one could say that while getting any achievement is easy, great achievements are hard to come by. A view like this has to be assessed in the balance of conservatism/fruitfulness; in this case, it is unclear that the revisionary view has more expressive power than the ordinary one (although it should not have less expressive power).

can only be satisfied by the performance of responsible epistemic agency (cf. Elgin (2017, ch. 5)), which in turn underlies the possibility of belief.<sup>19</sup> This makes the whole approach considerably more demanding than it may appear at first.

Another case where Elgin seems to be liberal is in her treatment of the capacities that underlie the possibility of mistakes. Her discussion of what we can learn from mistakes suggests that we can treat mistakes as manifestations of capacities which have epistemic significance. We have to be competent enough in order to make meaningful mistakes (p. 301). What are these competences, and how are they epistemically significant? What kind of account of competences does Elgin need?

Some general remarks on the notion of competence are in order. Competences are dispositions. As such, they are often at least partially individuated by what we can call their manifestation conditions.<sup>20</sup> The flammability of a dried leaf is a disposition such that, if the conditions were right (if a lighter was turned on near enough, for example), the leaf would light in fire. In those circumstances, we can say that the fire is the manifestation of the disposition. It is a commonplace that dispositions in general manifest only contingently in the case the trigger conditions are satisfied; they can go unmanifested. If they are manifested, the connection between manifestations and dispositions is modally robust, and also explanatory: *ceteris paribus*, the manifestation would not have easily happened unless the disposition was in place.

Competences, more specifically now, are dispositions to perform. To perform *how*? A natural minimal requirement is that the performances that characterize competences (their *exercises*) should be *effective*: if they were attempted, they would characteristically be carried out to completion. This, however, is probably not sufficient. For competences to be significant, one may think that their manifestations must constitute at least some minimal form of success or signal a positive status. The effective completion of a performance is not inherently a success condition: failures can be complete performances.<sup>21</sup> The required connection to success conditions is characteristic of many competence-based epistemologies. For example, Sosa (2010) endorses the following of a characterization of the notion of competence:

(Well Performing)

Competences are dispositions of an agent to perform well.

As is well known, Sosa has argued that knowledge can be accounted for as belief that is true because of competence, i.e., *apt* belief (cf. Sosa (2007)). More recently (2010; 2011; 2015) he has spelt this idea out saying that knowledge is belief that

<sup>19</sup> Rohrbaugh (2015) develops a notion of ‘inner’ achievement which can be useful here: such achievements apply to ‘inner’ projects like health, or friendship, where ‘our activities do not aim at bringing about some further, as yet unachieved goal, but instead constitute, realize, or prolong the valuable state itself.’ Engaging in such projects is their own achievement. Rohrbaugh suggests that this model of achievement also applies to the epistemic domain, where it pertains to the adequacy of our reasons to believe (p. 1200); instead of ‘believe’, Elgin could say ‘accept’ (see (2017), pp. 97–99).

<sup>20</sup> The traditional view is that dispositions are individuated by a pair of sets of trigger conditions and manifestations. Some, like Vetter (2014), disagree (she argues that the set of manifestations is sufficient). Here, I only need to assume that manifestation conditions are necessary, and remain neutral about sufficiency.

<sup>21</sup> Someone could bite the bullet and accept that the effective completion of a procedure is enough for a minimal sense of success, but I think this trivializes the notion of success.

is true in a way that *manifests* competence. For it to be knowledge, a belief must be true not only *because of* the underlying competence. Rather, the epistemic value of the belief in this case derives from the disposition being realized in a way that characterizes the success condition that is the constitutive manifestation of the relevant competences. Turri (2011) makes a similar suggestion. There, he argues that we should account for Gettier cases as situations in which the relevant competences are not manifested, even though in some sense the true beliefs formed are got *because of* having those competences; this is, in his view, why Gettier cases are not knowledge (*ibid.*, p. 7). The upshot of emphasizing that epistemic competences should be manifested is that some ways in which beliefs can have epistemic value (for example, by being true) are ruled out as properly epistemically significant because they lack the required connection to the supposed underlying competences. The suggestion is that, at least in this case, the competences of epistemic significance should be manifested through the fulfilment of success.

I will try to argue that the notion of competence of epistemological interest extends to cases where having a competence has a significantly more tenuous link to the notion of success; that is, I will try to argue against something like (Well Performing) as a characterization of the competences or dispositions that have epistemic significance.<sup>22</sup> Note, however, that I don't intend this to be a direct criticism of the general approach taken by proponents of competence-based epistemologies; my goal is rather to clarify what notion of competence should be adopted if we accept a picture of the epistemic like Elgin's.<sup>23</sup> Nevertheless, there is a range of positions that would be affected by my argument. I will proceed by considering first the ones that require the strongest link to success, to then move on to views that require a weaker link.

At one extreme, then, we have views that make the satisfaction of success conditions constitutive of exercises of competences. These endorse

(Success Requirement Thesis, SRT)

Epistemically significant competences can only be manifested in success conditions or achievements.

Let's call a disposition that is manifested mostly or characteristically in success conditions *success-oriented* and a disposition that can only be manifested in success conditions *strongly success-oriented*. Then, SRT says that all epistemically significant competences are strongly success-oriented. Although relatively unpopular, something like this has been endorsed by Millar (2008), who claims that perceptual recognitional abilities are success-oriented: according to him, to exercise them at all is to succeed. It is unclear how a view like this can generalize. Indeed, Elgin's discussion of meaningful mistakes seems to provide examples of manifestations of epistemically significant dispositions/competences in failures. These dispositions/competences are not strongly success oriented. It is not obvious how the account can accommodate the intuition that competent failures are possible (Miracchi (2014) raises this point too).

<sup>22</sup> I thank an anonymous reviewer for pressing me on the reach of the argument that follows.

<sup>23</sup> I will remain neutral on whether we should make use of the notion of competence or on whether competence-based accounts are the best explanations for epistemic phenomena. My point is merely that there are constraints on what notion of competence can be used.

One way to do it is to use Turri's distinction between performances that are effective because of standing competences and performances that are effective through the manifestation of competences. The explanation of cases like mistakes could go as follows: mistakes are performances that are effective because of competences, but they are never themselves the manifestation of competences. In this view, competent failures are not manifestations of competence, even though they depend on standing competences. Note that this satisfies Elgin's thought that we need to be competent enough in order to make meaningful mistakes.<sup>24</sup>

However, if we look at some cases of competent failures, it is difficult to deny that the relevant dispositions are being manifested, not merely acting as standing preconditions or constraints. Consider Sosa's example of the capable archer who misses a shot only because of the interference of wind. If the wind had not interfered, the shot would have reached its target because of the archer's performance and their competence. It is implausible to think that the archer's competences extend to the environmental conditions. So it is at least plausible that the shot manifested the archer's competence to shoot at a target.

If we want to allow competent failures to be manifestations of competence, thus, we need to do something different. One way to allow this is to weaken the scope of the requirement for success. Perhaps not all exercises of competences have to be successes, but at least significant portion of them should be:

(Characteristic Success Requirement Thesis, CSRT)

Epistemically significant competences are characteristically manifested in success conditions.

There are at least two ways to pursue this approach, depending on how to understand 'characteristically'.

The first is what I will call generically the *ratio* approach. The idea is that there is some kind of threshold ratio between the number of successful and unsuccessful manifestations: the unsuccessful cannot be so many as to make success uncharacteristic of the competences. The ratio might be spelled out extensionally (of the actual exercises of a competence, most are successes), or modally (with many variations; a simple example is to require that most of the possible exercises are successes). This is the approach taken by Sosa (2015, p. 96–97).

Under this interpretation, it would seem as if dispositions that are rarely successfully exercised should never count as competent. However, this is not adequate for a criterion of epistemic significance: the rarity of success could simply be a measure of how hard the epistemic tasks involved are. A way to recover the intuition that dispositions that can rarely be exercised can be significant is to make the ratio of successes/failures context-sensitive (for example, to the difficulty of the task). Indeed, something like this is necessary to account for the acquisition of competence in a general sense: it is plausible that the acquisition of competence does not guarantee uniform ratios of success in apprentices who are at least competent enough to

<sup>24</sup> A different worry is that Turri's idea was meant to explain successes which nonetheless failed to constitute a higher form of achievement, so it might not be apt for the case of failures.

make meaningful mistakes.<sup>25</sup> This gives an upper bound for the strength of the link between competences and success.

The second approach is what I will call the *disjunctive* approach. In this approach, the exercise of competences is a disjunctive kind: competences can have both successful and failed exercises. Nonetheless, only the successful exercises are characteristic of the competences, because they satisfy all the competence's manifestation conditions. This is the approach endorsed by Miracchi (2014; 2015). She takes unsuccessful cases as degenerate, since they deviate from the ideal manifestation conditions of the competences. In her account, ideal manifestation conditions are factive.

The disjunctive approach does not have any issue handling epistemically significant competences that are rarely successfully manifested: these cases simply witness the ease with which certain competences can deviate from being fully manifested. However, Miracchi's assumption that factive success conditions are characteristic of all epistemically significant competences seems to go against the idea that non-factive conditions can be epistemically significant independently of the significance of factive conditions. From this perspective, competences manifested in unsuccessful conditions should not be explained away in terms of their deviancy from ideal factive conditions, but accounted for in their own terms. That Miracchi's version of the disjunctive picture is incompatible with Elgin's overall picture is not an indictment of disjunctivism, of course (nor against Miracchi's account, if we are not already sold on Elgin's more pluralist presuppositions).

We should be able to go back now in a fuller form to our original question: what sort of theory of competence should we adopt in light of Elgin's discussion of the meaningfulness of mistakes? Some versions of CSRT should be compatible with her account of mistakes (contextual ratio accounts and non-factive disjunctive views are candidates). It shouldn't surprise that Elgin's view retains some link between competences and success conditions: after all, in her account the states underlying mistakes are explicitly achievements.

Nevertheless, I think we can go a bit further in weakening the connection to success.

Both the ratio and disjunctive approaches endorse a basic assumption that

(Possibility of Success Requirement Thesis, PSRT)

It is possible for epistemically significant competences to be manifested in success conditions.

Let's call a disposition that can be manifested in success conditions, *success-apt*. PSRT is, then, the claim that epistemically significant competences are necessarily success-apt. I imagine that the support for this idea comes from two more basic principles: a) that it is impossible for epistemically significant competences to be manifested only in failures, and b) that competences can only be manifested through successes or failures. I have nothing to say against the first principle, but it seems at least within the realm of logical possibility that some competences could be manifested

<sup>25</sup> Cf. Elgin's (2017, pp. 59–61) criticism of Kvanvig's (2003) idea that in some cases (like with incompletely educated children) we attribute understanding only in an honorific sense. If the children's understanding is epistemically significant (*qua* understanding), then their competences should also be epistemically significant.



in conditions that are neither successes nor failures.<sup>26</sup> Such competences would be *success-neutral*. Are such competences possible in a more substantive sense? I will argue that they might.

Suppose we decompose the process by which we arrive at a belief in stages, where the effective completion of a stage is a prerequisite for moving on to the next step. To complete any of these stages one has to be at least competent enough to do precisely that. This gives ‘partial’ completion conditions for each stage: the success condition for the overall completion of the process requires the satisfaction of each of those partial completion conditions. It is natural to think that the competence to complete the overall procedure is constituted by these competences to complete each of the stages, plus the competence to satisfy those in the appropriate order. Now, the completion of a stage in a procedure like this needs not be itself a success-condition: for example, the result of a stage could be wrong without interrupting the overall procedure (for example, the mistaken result could become the ‘input’ of the next stages and propagates).<sup>27</sup> Are these competences epistemically significant? SRT says no; CSRT gives different verdicts. Ratio versions of the CSRT will predict that these competences are epistemically significant if they are characteristically manifested in successes; for example, they would be significant if the fact that the procedure gave a wrong result was merely accidental. Similarly, disjunctive approaches could also predict epistemic significance; for example, in Miracchi’s view this will depend on whether the manifestation of the competence can be understood as a deviation from the best case.

We can picture a simpler account. The overall process would not have reached completion unless the subject would have had the capacity to complete the partial unsuccessful stages. Clearly epistemically significant parts of the procedures depend on them. So if the outcome of the process has epistemic significance, these underlying capacities should be epistemically significant as well. Furthermore, these capacities are *manifested* in the completion of the unsuccessful stages. In this case, completion as a part of an epistemic process *is* enough for epistemic significance. Importantly, it is not necessary that we treat these cases as neither successes nor failures. Transposing to the case of competences, we end up with the sufficiency claim

#### (Completion in an Overall Process, COP)

The completion of parts of epistemically significant processes is manifestation of competences of epistemic significance.

This endorses effectiveness instead of success as a criterion of epistemic significance.<sup>28</sup> Effective manifestations of competences can be success-neutral. So we

<sup>26</sup> Miracchi (2014, pp. 31–33) criticizes an argument by Burge (2011) that every case of success could have been a failure. If we allow the possibility of neutral exercises, a similar argument can be rejected for the claim that for every case of failure, there could have been a success.

<sup>27</sup> It will not necessarily or even inescapably propagate *as a mistake*. This will depend on what kind of cognitive safeguards are in place; we have seen how we can have epistemic policies to minimize risks, and these could minimize the consequences of faulty procedures.

<sup>28</sup> Someone could object that the kind of systematicity requirement that COP endorses (that the exercise of a competence fits in a larger process) should be itself understood as a kind of achievement (for example, in analogy to Elgin’s view on understanding as a form of achievement, which also relies on structural properties of the cognitive states). If we accepted that, COP would collapse into a version of PSRT. However, structure is ubiquitous, so there is no reason to think that it is universally or inescapably valuable as a

should admit at least competences that have neutral manifestations. This is compatible with the rejection of PSRT. I will now offer two cases that potentially offer counterexamples to PSRT; even if they fail to do so, they still should provide something like a lower bound for the link between epistemic significance and success.

In the picture I just sketched, it is not necessary for the overall processes to be successful themselves. Even if belief formation processes end in failure, as in the case of mistakes, it is not impossible that their outcomes could nevertheless serve as the input for further processes (for example, in the application of policies of amelioration or recovery). A way to account for these cases without modifying the account is to describe them as cases where local epistemic processes are embedded in broader ones. This does not mean that the significance of the parts of the local processes should be accounted in terms of the success of the broader processes. In many cases the success of such broad epistemic processes cannot be taken as a *fait accompli*. We have thus to evaluate epistemic significance before we can evaluate success. The significance of these broader processes lies, from the perspective of the underlying processes, merely on the potential of success or in their being oriented/aimed towards success conditions; what witnesses this potential is simply their position in broader, overall processes.

Similarly to how we can recover from mistakes, we can also recover from not being able to complete actions that would constitute part of epistemic processes. Suppose that we made an attempt to make a very involved calculation, but after spending some time trying, we realized that we will not be able to complete it by following the procedure we have been using. We could recover from this situation if we examined the limitations of the unsuccessful procedure and looked for alternative ways to achieve our aims. This may prove fruitful. In cases like this, the relevant competences are simply competences *to try*. This suggests that epistemic significance does not depend on the manifestation of effective competences either (regardless of whether they are success-neutral or not).<sup>29</sup>

## 5 Conclusion

Elgin's account of the epistemic significance of mistakes is very rich and suggestive. In this paper, I have tried to expand and clarify some of the lines of thought that it

---

success. We need stronger reasons to think that in this case the required structure is an achievement. Consider, however, that some intuitively epistemically/cognitively significant processes (significant because we would have to appeal to them to explain our epistemic/cognitive states) could be automatic, and consequently not personally attributable as an achievement. Elgin's argument for the idea that understanding is an achievement seems to require that acceptance is in some sense under personal control (2017, pp. 95–99). But this is not so plausible for the present case. On the other hand, the overall view could be adjusted so that the claim that understanding is an achievement was replaced by the claim that understanding is a manifestation of competence, without the assumption that the relevant manifestations of competence are achievements.

<sup>29</sup> In some cases, like when the requirements for trying are very involved, coming to be able to try should count as an achievement, but in general coming to try is a neutral condition. To count trying as a form of success is to trivialize success (cf. also note 21 above). Offering a defense of this idea is outside of the scope of this paper. If it is not granted, the conclusion we arrive at is simply a statement of the lower bound of success that can be epistemically meaningful, as suggested above. That is success enough.

offers. In particular, I have examined two areas of interest. On one hand, there is the problem of what is the information that we can obtain from mistakes, and how it can be used in our various policies of error management. On the other, I have tried to clarify the way in which the notion of competence could be used in a theory of understanding like Elgin's. While various options are open, I have argued that there are upper and lower bounds. The assumption that epistemic significance comes from success can be weakened in various ways, most importantly by noting that something can be significant without it constituting neither as a success or a failure. The corresponding neutral competences can nonetheless be part of epistemically significant systems. This is important when we consider the issue of what are the minimal conditions for the engagement in epistemic practices. In a world of uncertainties, minimal epistemic significance might be a matter of whether it is admissible to guide our tries by the potential of success.

**Acknowledgements** This research was supported by CONICYT's (Chile) Beca de Doctorado para Estudios en el Extranjero Nr. 72180096. I want to thank my colleagues at KU Leuven, and in particular, I would like to thank Lars Tump and Jan Heylen for their feedback. I would also like to thank two anonymous reviewers for their helpful suggestions.

## References

- Berto F, Hawke P (2018) Knowability relative to information. *Mind* DOI 10.1093/mind/fzy045
- Burge T (2011) Disjunctivism again. *Philosophical Explorations* 14(1):43–80
- Church IM (2016) The doxastic account of intellectual humility. *Logos & Episteme* VII(4):413–433
- Cohen LJ (1992) *An Essay on Belief and Acceptance*. Oxford University Press, Oxford
- van Ditmarsch H, van der Hoek W, Kooi B (2007) *Dynamic Epistemic Logic*. Springer, Berlin
- Duhem P (1954) *The Aim and Structure of Physical Theory*. Princeton University Press, Princeton
- Elgin CZ (2002) Art in the advancement of understanding. *American Philosophical Quarterly* 39:1–12
- Elgin CZ (2012) Ignorancia, error y el avance de la comprensión. In: Campos V, Estany A, Izquierdo M (eds) *Error y Conocimiento: La Gestación de la Ignorancia desde la Didáctica, la Ética y la Filosofía*, Comares, Granada
- Elgin CZ (2017) *True Enough*. The MIT Press
- Fine K (2017) Truthmaker semantics. In: Hale B, Wright C, Miller A (eds) *A Companion to the Philosophy of Language*, Blackwell, pp 556–577
- Friedman J (2017) Junk beliefs and interest driven epistemology. *Philosophy and Phenomenological Research* DOI 10.1111/phpr.12381
- Harman G (1986) *Change in View: Principles of Reasoning*. The MIT Press, Cambridge, Massachusetts
- Hawke P (2017) *The problem of epistemic relevance*. PhD thesis, Department of Philosophy, Stanford University
- Kallestrup J, Pritchard D (2016) From epistemic anti-individualism to intellectual humility. *Res Philosophica* 93(3):533–552
- Kvanvig J (2003) *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press, Cambridge
- Lewis D (1988) Relevant implication. *Theoria* 54:161–174
- Mach E (1976) *Knowledge and Error: Sketches on the Psychology of Enquiry*. D. Reidel Publishing Company, Dordrecht, Holland
- Millar A (2008) Perceptual-recognition abilities and perceptual knowledge. In: Haddock A, Macpherson F (eds) *Disjunctivism: Perception, Action, Knowledge*, Oxford University Press, Oxford
- Miracchi L (2014) *Getting things done*. PhD thesis
- Miracchi L (2015) Competence to know. *Philosophical Studies* 172:29–56
- Polya G (1957) *How to Solve It: A New Aspect of Mathematical Method*, 2nd edn. Princeton University Press

- Pritchard D (2017) Anti-risk epistemology and negative epistemic dependence. *Synthese* DOI 10.1007/s11229-017-1586-6
- Rescorla M (2009) Assertion and its constitutive norms. *Philosophy and Phenomenological Research* 79(1):98–130
- Rohrbaugh G (2015) Inner achievement. *Erkenntnis* 80:1191–1204
- Sosa E (2007) *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. Oxford University Press
- Sosa E (2010) How competence matters in epistemology. *Philosophical Perspectives* 24:465–475
- Sosa E (2011) *Knowing Full Well*. Princeton University Press, Princeton, NJ
- Sosa E (2015) *Judgment and Agency*. Oxford University Press, Oxford
- Tanesini A (2018) Intellectual humility as attitude. *Philosophy and Phenomenological Research* 96(2):399–420
- Turri J (2011) Manifest failure: the Gettier problem solved. *Philosopher's Imprint* 11(8):1–11
- Vetter B (2014) Dispositions without conditionals. *Mind* 123(489):129–156
- Von Kriegstein H (2016) Effort and achievement. *Utilitas* 29(1):27–51
- Whitcomb D, Battaly H, Baehr J, Howard-Snyder D (2015) Intellectual humility: owning our limitations. *Philosophy and Phenomenological Research* XCIV(3):509–539
- Yablo S (2014) *Aboutness*. Princeton University Press, Princeton