

On the Epistemology of the Precautionary Principle: Reply to Steglich-Petersen

J. Adam Carter & Martin Peterson

In recent paper in this journal (2015), we proposed two novel puzzles associated with the precautionary principle. Both are puzzles that materialise, we argue, once we investigate the principle through an *epistemological* lens, and each constitutes a philosophical hurdle for any proponent of a plausible version of the precautionary principle. Asbjørn Steglich-Petersen (2014) claims, also in this journal, that he has resolved our puzzles. In this short note, we explain why we remain skeptical.

INTRODUCTION

The core idea motivating the precautionary principle is that it is sometimes appropriate to take preventative action rather than wait for more information to come in.¹ In a recent paper in this journal (2015), we proposed two novel puzzles associated with the precautionary principle. Both are puzzles that materialise, we argue, once we investigate the principle through an *epistemological* lens, and each constitutes a philosophical hurdle for any proponent of a plausible version of the precautionary principle. Asbjørn Steglich-Petersen (2014) claims, also in this journal, that he has resolved our puzzles. In this note we explain why we remain skeptical.

¹See Aven (2011), Manson (2002), Munthe (2011), Sandin (1999), Sandin (2004) and Peterson (2006) for a sample of some notable interpretations of the precautionary principle.

THE FIRST PUZZLE

The precautionary principle can be understood as composed of three central elements: a damage condition (D), an epistemological condition (E), and a suggested remedy (R). If D and E are satisfied, then R is prescribed, or activated. On this line of thinking, if there is sufficient epistemic confidence that an activity will bring about a damage of a certain degree of severity, then this is sufficient grounds for taking precautionary measures against the activity.² In short, the rule template is: $E \ \& \ D \rightarrow R$.

Any version of the precautionary principle in which E and D are too *easily* satisfied will be implausibly demanding. It would be absurd, for instance, to suppose we must take regulatory action against laws that permit bicycling because there are some (unsubstantiated) claims that bicycling causes some near-negligible bad consequence.³ By contrast, if the conditions for satisfying E and D are too *strict*, the resulting decision rule will be practically inert. The rule will prescribe action against an activity, *a*, *only if* we have something like Cartesian certainty that some *a* causes some maximally destructive damage.

The gist of the first puzzle we posed can be appreciated in the form of a dilemma that arises once we ask: should the epistemic standards that must be met for ‘E’ in the template to be satisfied depend on the severity of the anticipated damage (D)? Here we envisioned two styles of response. Call *invariantism* the view that the question posed be answered negatively, and *contextualism* the view that the question be answered positively. Our ‘first puzzle’ for the proponent of the precautionary principle emerges because *either* strategy seems to lead to a dead end.

Running a *reductio* against invariantism is simple. Suppose there is some body of evidence that building a certain experimental nuclear energy reactor (call this activity ‘*a*’) will melt down in a way that causes (in Case A) ten deaths, or (in Case B) ten billion deaths. How strong must the decision maker’s epistemic position be to rationally mandate precautionary action against *a*? Here the invariantist has a flimsy story: *the same in both cases*. But surely preventative action against *a* is warranted in Case B on epistemic grounds that are not as good as they must be to warrant preventative action against *a* in Case A. After all, Case B threatens to wipe out ten billion lives.

²See Carter & Peterson (2015) for a more detailed presentation.

³Just suppose (for the moment) that quality of one’s epistemic position, as well as the severity of some anticipated damage, can be represented numerically by the real numbers between 0 and 1. (This, note, is in fact an oversimplification of things, though this assumption suffices to make the present point). Now, for *reductio*, imagine a (bizarre) version of the precautionary principle that says: *so long as* (i) one’s epistemic position *e* with respect to the causal connection between some activity, *a*, and an anticipated damage, *d* is such that $e > 0$, and (ii) the ‘badness’ of an anticipated damage (*d*) is such that $d > 0$, then (iii) *R*, where (recall) *R* specifies some regulatory action against *a*. Such a principle, set where *e* and *d* are infinitesimally low, will rationally mandate regulatory action against nearly *all* actions.

The contextualist avoids this problem, but only by incurring another. The contextualist describes the relationship between E and D in terms of negative correlation: the *greater* the severity of the anticipated damage, the *lower* the epistemic standard that needs to be met for the relevant regulatory action to be prescribed. In this respect, the strength of epistemic position that must be met *tracks the practical stakes*.⁴ And this seems to get the right result in the case above.

But the proponent of the precautionary principle who takes the contextualist route (where E is understood as depending on D) walks right in to a trap, or, more weakly, a puzzle. To make this point concrete, consider that on the contextualist route, the rating of the severity of the anticipated damage circumscribes how strong the epistemic position must be for the E condition to be satisfied. Though, as we noted, in situations where an application of the precautionary principle is germane, we have *competing interests* at play, and the extent to which the severity of some anticipated damage is judged is in many cases going to vary dramatically with respect to whose interests are at stake. Therefore, the precautionary principle can be plausibly formulated as a decision rule only if supplemented with some additional *favouring rule*, a rule which adjudicates whose interests are relevant to determining the severity of the damage in question, a severity that for the contextualist is what fixes the relevant epistemic standard that must be satisfied.⁵

At this point the contextualist has no good card to play. In our original paper we argued that the precautionary principle fails to be an appropriately *neutral* decision rule if, built in to what is relevant to assessing the severity of damage (a severity which then, on the contextualist view, determines the strength that must be met for the E condition to be satisfied) is some *particular* interest—say, the interest of developers as opposed to environmentalist groups (groups who might well rate the severity of the anticipated damage very differently). But the contextualist cannot simply retreat to the natural countermove—that of saying that all interests be weighted equally—call this an *egalitarian favouring rule*. This move, we argued, will ultimately incur the same problems associated with invariantism.⁶ We considered, finally, a third option for

⁴The thought that epistemic standards should track practical stakes has obvious commonalities with contextualist approaches in epistemology (e.g. DeRose (1992); DeRose (2009); Cohen (1998)) according to which the epistemic standards that must be satisfied for an individual to count as knowing a proposition are sensitive to the practical stakes, in the context in which the knowledge ascription is made. While we term the approach to the precautionary principle we are envisaging *contextualism*, we should note that contextualism in epistemology is, at root, a *semantic* thesis about knowledge attributions. Our discussion of the precautionary principle is situated within a semantic framework, and the general suggestion that epistemic standards *track* practical stakes is germane to both contextualist and subject-sensitive invariantist semantics for knowledge attributions (e.g. Stanley (2005); Hawthorne (2004)).

⁵Carter & Peterson (2015, 7).

⁶In short, given the *entrenched* political interests of (for example) developers and environmental regulatory bodies, the mean of all relevant standards for damage will typically be insensitive to differ-

the contextualist: to contextualise the favouring rule itself and argued that this merely forestalls, rather than solves, the problem. As no option looks tractable, the proponent of the precautionary principle *qua* decision rule thus has a *bona fide* puzzle on her hands, one that trades importantly on how to think about just how the epistemological condition should be satisfied.

However, according to Steglich-Petersen (2014) the puzzle sketched above should, contrary to what we've suggested, not trouble proponents of the precautionary principle. His rationale takes as a starting point the correct observation that, independently of any epistemic considerations, use of the precautionary principle will involve adjudicating between competing interests involved in the decision scenario (Steglich-Petersen 2014: 3). No arguments here. Steglich-Petersen then observes, also correctly, that *any* decision rule that relies on assessments of costs and values will be subject to some prior adjudication of interests, especially when applied in policy making. And this fact is 'not normally seen as an argument against the usefulness [of the decision rule in question]' nor for adopting an additional favouring rule. As Steglich-Petersen writes:

Such adjudication will typically be done through the normal political processes, sometimes democratic, which are designed to facilitate adjudication of that exact kind ... [But this is not] a problem arising from the contextualist approach to the epistemic standards involved with the precautionary principle, and Carter and Peterson [have not] succeeded in pointing out a new puzzle for the precautionary principle. (Steglich-Petersen 2014: 3.)

In short, Steglich-Petersen has attempted an *undercutting* response to our puzzle. Rather than to show how it can be met, he argues that it doesn't *need* to be met, and he argues for this via the familiar strategy of overgeneralization: Our puzzle should worry a proponent of the precautionary principle only if our puzzle also constitutes a legitimate problem for *any* decision rule that must weight competing interests.

Steglich-Petersen is right that if we had drawn attention to a fact about the precautionary principle which applies *mutatis mutandis* to any decision rule the use of which involves weighing competing interests, then we will not have in doing so highlighted any puzzle that is specific to the precautionary principle. But we reject this characterisation of our argument. In particular, we resist the suggestion that the problem we've posed *vis-a-vis* the precautionary principle is in fact one that overgeneralizes as Steglich-Petersen suggests. We want to be clear that the puzzle we've raised does not

ences in actual damage. But sensitivity to the severity of actual damage is precisely what contextualist approaches can claim, as a key advantage, over invariantist approaches. Retreating to contextualism supplemented with an egalitarian favouring rule surrenders this advantage.

arise because applying the decision rule involves adjudicating *per se* between competing interests involved in the decision scenario. Consider a decision rule that says *for any surplus S, give S to the most needy*. The application of this rule involves adjudicating between competing interests with respect to whom are rightly described as the needy. But this shows nothing perverse about the rule itself, even if there is some difficulty in applying the rule – a difficulty that, as Steglich-Petersen observes, will typically be met democratically.

Our argument, more subtly, attempts to show that there is a special difficulty with respect to how to *formulate* (and not merely *apply*) the precautionary principle, in a way that satisfactorily captures how the epistemological component of the view should be regarded as satisfied. This special difficulty arises in light of the relationship between the ‘D’ and ‘E’ conditions, in the decision rule. As we’ve suggested (both here and in our original paper), we think it is obvious that a formulation of the precautionary principle on which the level of epistemic support that must be attained to satisfy the ‘E’ condition does *not* track practical stakes (e.g. the extent of the anticipated damage) is simply unworkable—as suggested by our *reductio*. However the only natural move left for the purposes of *formulating* the principle in a plausible way is to ‘go contextualist’. Though we systematically showed that each of the available ways of going contextualist looks *prima facie* intractable. And this is problematical for the proponent of the precautionary principle even if we grant with Steglich-Petersen that, in the default case where a decision rule is applied in a way that involves the adjudication of interests, such adjudication can be unproblematically accomplished.⁷

⁷We want to emphasise that simply adverting to the point that (in typical cases) where the application of a decision rule involves weighing interests, such adjudication can be satisfactorily accomplished through democratic political means is not a point that gains any traction for the purposes of spelling out a plausible ‘contextualist’ route to articulating the epistemological condition of the precautionary principle. Such a democratic process will be one that, at least in principle, will aim to weight competing interests equally. However, this approach quickly begins to resemble in the relevant respects what we called the *egalitarian favouring rule*—the favouring rule we envisioned (in our 2014 paper in this journal) on which, in determining whose interests are relevant to assessing the severity of the anticipated damage (a severity that on the contextualist approach to the PP will be what circumscribes the epistemic standards that must be met for the ‘E’ condition to be satisfied), interests are given equal consideration. But the egalitarian favouring rule, we argued, ultimately *relinquished* the advantage the contextualist was originally said to have over the unworkable invariantist approach. The rationale here was that, given the *entrenched* political interests of (for example) developers and environmental regulatory bodies, the mean of the relevant standards for assessing the severity of the anticipated damage will typically remain fixed across cases that vary along damage-relevant dimensions of evaluation such as: how many individuals will be affected, how many carcinogens will be released into the air, etc. Given the contingent fact of entrenched political interests, then, a democratic approach to adjudicating interests vis-a-vis assessing (on a contextualist approach to the precautionary principle) the severity of the anticipated damage will, *mutatis mutandis*, end up in many cases tracing the path of the egalitarian favouring rule. And, again, this path did not end up *retaining* the advantage we claimed the contextualist approach had over

THE SECOND PUZZLE

Everyone agrees that if we were to take precautionary measures against *every* risk that *could* lead to a sufficiently bad outcome, we would have to apply the precautionary principle to far too many (perhaps all) activities. The point of departure for the second puzzle is the plausible idea is that in order to avoid this problem, the precautionary principle has to be formulated in a way that is compatible with the *de minimis* principle, according to which extremely unlikely threats should be ignored, i.e. threats that fall below some appropriate epistemic threshold.

In our original paper, we argued that the best way to express this epistemic condition is to assign some non-zero probability to the hypothesis that the risk is below the *de minimis* threshold.⁸ A decision about how to address a risk should therefore be based on two probability functions: the first describes the *first-order* probability that the risk is *de minimis*, and the other describes the *second-order* probability that the first-order probability is correct.

The second puzzle arises as we try to render this connection between the two probability functions and the decision to treat a risk as *de minimis* more precise. We face a choice between two fundamentally different ways of doing this. The first option is to aggregate the first- and second-order probability functions into a single measure of uncertainty; the second option is to refrain from aggregating the two measures.

Steglich-Petersen agrees with us that the second strategy, to refrain from aggregating the two measures, is implausible for the reason we state in the original paper: if we do so it would become necessary to introduce two or more probabilistic thresholds for when a risk is *de minimis*, but then there will always be cases in which small changes to one of thresholds trigger a “too big” effect on the overall assessment. (See Carter and Peterson 2015: 11-12.)

But what about the first strategy? Why can't we just aggregate the first- and second order probabilities into a combined measure by, for instance, multiplying the two numbers? The problem is that if the first order probability that, say, some substance *x* is carcinogenic is one in a billion, and we are 90% sure that this probability is correct, then the combined probability that the risk is *de minimis* may very well be less than one in a billion. This contradicts the gist of the precautionary principle. Other aggregating principles face similar problems.

Instead of calculating the probability that *x* is carcinogenic, it is possible to avoid the problem by “turning around” the probabilities and calculating the probability that

the invariantist. Thanks to an anonymous referee at *Erkenntnis* for suggestions that have improved the presentation of this discussion.

⁸It should be noted that the second puzzle can also be stated in a non-probabilistic vocabulary. See Carter and Peterson (2015, 10).

x is *safe*. Steglich-Petersen describes this proposal, briefly discussed in our original paper, as follows:

the relevant all-things-considered probability [is] the probability that some target activity is *safe*. This means that the lower the second-order probability, the lower [is] the all-things-considered probability of the activity being safe. This is the right result, since the second-order probability measures the reliability of the first-order assessment, and a lower second-order probability should therefore correspond to a reduced probability of safety. A risk would then be deemed *de minimis* when the probability of the activity being *safe* is sufficiently *high*, or, which is to say the same, when the probability of *harm* is sufficiently *low*. (Steglich-Peterson 2014: 9)

What Steglich-Petersen and we disagree about is whether this response is *ad hoc* or not. We believe it is, but he argues it is not. Although we concede that “turning around” the probabilities would give us the result we want, we think that what is lacking is a *reason* for claiming that this line of reasoning is correct and the alternative one wrong.⁹

Steglich-Petersen believes he has found a reason of the sort we believe to be lacking. He correctly observes that many formulations of the precautionary principle involve some claim about a *reversed burden of proof*, according to which it is not the party who believes that an activity is dangerous who has to prove this, but rather the party who thinks that the activity is safe. A clear example can be found in the Wingspread statement, according to which “the proponent of an activity, rather than

⁹Before proceeding it is worth commenting briefly on Steglich-Petersen’s claim that a low second-order probability does not *always* result in an increased overall risk. He gives the following example:

Suppose that I live in Baltimore and read in the weather forecast that there is a

0.8 probability of thunderstorms, but subsequently come to doubt that I have read the forecast for Baltimore and not Chicago (perhaps I set my phone to sometimes report the one, sometimes the other), and assign it a 0.5 probability that I have read the right forecast. What does this relatively low second-order probability imply about my overall risk of encountering thunderstorms in Baltimore? Nothing, it seems. The second-order probability of 0.5 favors an overall risk higher than the first-order probability of 0.8, as much as it favors an overall risk lower than 0.8. (Steglich-Peterson 2014: 6)

We agree that in this example the low second-order probability of 0.5 does not increase the overall risk. This is because the probability of encountering a thunderstorm in Baltimore and Chicago is (to the best of our knowledge) roughly the same. However, in nearly all cases in which the precautionary principle is discussed, the probability that any two risks are *de minimis* is not the same. The *overwhelming majority* of all risks that reach the public discourse are not *de minimis*. This is, of course, an empirical assumption--but to be fairly uncontroversial.

the public, should bear the burden of proof”.¹⁰ If it is the proponent of an activity who has the burden of proof to show that an activity is safe, then it is—according to Steglich-Petersen—not *ad hoc* to claim that we should calculate the probability that an activity is safe rather than calculating the probability that it is dangerous. We remain skeptical, for two reasons.

Firstly, although we agree that that many formulations of the precautionary principle do actually contain some claim about a reversed burden of proof, it is debatable whether one ought to accept this element of the principle. The mere fact that many defenders of the precautionary principle seem to think that it is the proponent of an activity who has the burden of proof to show that an activity is safe does not entail that they are right, or have some non-arbitrary reason for making this claim. No purely *exegetical* study of what defenders of the precautionary principle say or do not say will help us to solve the second puzzle. What we need is a *reason*.

Secondly, we would like to stress that even if we were to accept Steglich-Petersen’s claim about a reversed burden of proof, nothing follows from this about how we are to aggregate first- and second-order probabilities. The burden-of-proof claim is a point about *who* should provide information about an activity: “the proponent of an activity, rather than the public, should bear the burden of proof”.¹¹ This is not a claim about *how* the relevant information should be aggregated. Therefore, it will not suffice to just point out that some formulations of the precautionary principle state that “the proponent of an activity, rather than the public” has the burden of proof.¹²

To be more precise, even if we think it is the proponent of an activity rather than the public who should bear the burden of proof, it seems clear that the proponent of the activity could fulfill this obligation in at least two different ways. Firstly, the proponent could calculate the probability that the proposed activity is safe. Secondly, he or she could calculate the probability that the activity is dangerous. Under normal circumstances, in which the only relevant probability is the first-order probability, it doesn’t matter which approach is taken. If “safe” is the opposite of “dangerous”, the probability that *x* is safe equals one minus the probability that *x* is dangerous.¹³ However, as we pointed out in our original paper, this type of straightforward reasoning is not applicable when we attempt to aggregate first- and second-order probabilities in to a single measure. Not very much hinges on which aggregation mechanism we choose. The multiplicative rule is simple and fairly attractive, but other aggregation rules will run into similar problems: the influence of the second-order probability of the overall-conclusion will depend of whether we calculate the probability that some-

¹⁰ Wingspread Conference on the Precautionary Principle, January 26, 1998.

¹¹ *Ibid.*

¹² *Ibid.*

¹³ We are aware that this is an oversimplification, but in order to keep the discussion simple we can ignore the possibility of a third option.

thing is risky or safe. What we need is not a technical solution to this problem (that would be easy to construct), but a *reason* for favoring the calculation that gives the intuitively most plausible answer. The upshot of all this is that our two puzzles remain unresolved.

REFERENCES

- Aven, Terje. 2011. "On Different Types of Uncertainties in the Context of the Precautionary Principle." *Risk Analysis* 31 (10): 1515–1525.
- Carter, J. Adam and Martin Peterson. 2015. "On The Epistemology of the Precautionary Principle." *Erkenntnis* 80: 1–13.
- Cohen, Stewart. 1998. "Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery." *Australasian Journal of Philosophy* 76 (2): 289–306.
- DeRose, Keith. 1992. "Contextualism and Knowledge Attributions." *Philosophy and Phenomenological Research*: 913–929.
- . 2009. *The Case for Contextualism: Knowledge, Skepticism, and Context, Vol. 1: Knowledge, Skepticism, and Context*. Vol. 1. Oxford University Press.
- Gardiner, Stephen M. 2006. "A Core Precautionary Principle." *Journal of Political Philosophy* 141: 33–60.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford University Press.
- Manson, Neil. 2002. "Formulating the Precautionary Principle." *Environmental Ethics* 24 (3): 263–274.
- Munthe, Christian. 2011. *The Price of Precaution and the Ethics of Risk*. Springer.
- Peterson, Martin. 2003. "Transformative Decision Rules." *Erkenntnis* 58: 71–85.
- . 2006. "The Precautionary Principle Is Incoherent." *Risk Analysis*, 26: 595–601.
- Sandin, Per. 1999. "Dimensions of the Precautionary Principle." *Human and Ecological Risk Assessment: An International Journal* 5 (5): 889–907.
- . 2004. "The Precautionary Principle and the Concept of Precaution." *Environmental Values* 13 (4): 461–475.
- Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford University Press.
- Steglich-Petersen, Asbjørn. 2014. "The Epistemology of the Precautionary Principle: Two Puzzles Resolved." *Erkenntnis*: 1–9.