

## ORIGINAL ARTICLE

# What's a(t) stake? On stakes, encroachers, knowledge

Peter Baumann 

Department of Philosophy, Swarthmore College, Swarthmore, Pennsylvania, USA

**Correspondence**

Peter Baumann, Department of Philosophy, Swarthmore College, Swarthmore, PA 19081-1306, USA.

Email: [pbauman1@swarthmore.edu](mailto:pbauman1@swarthmore.edu)

**Abstract**

According to subject-sensitive invariantism (SSI), whether S knows that  $p$  depends not only on the subject's epistemic position (the presence of a true belief, sufficient warrant, etc.) but also on non-epistemic factors present in the subject's situation; such factors are seen as "encroaching" on the subject's epistemic standing. Not the only such non-epistemic factor but the most prominent one consists in the subject's practical stakes. Stakes-based SSI holds that two subjects can be in the same epistemic position with respect to some proposition but with different stakes for the two subjects so that one of them might know the proposition while the other might fail to know it. It is remarkable that the notion of stakes has not been discussed much in great detail at all so far. This paper takes a closer look at this notion and proposes a detailed, new analysis. It turns out that there is more than one kind of stakes, namely event-stakes, knowledge-stakes and action-stakes. I discuss several issues that even plausible notions of stakes raise and propose solutions.

**KEYWORDS**

knowledge, practical stakes, pragmatic encroachment, subject-sensitive invariantism

According to subject-sensitive invariantism (SSI), whether S knows that  $p$  depends not only on the subject's epistemic position (the presence of a true belief, sufficient warrant, etc.<sup>1</sup>) but also on non-epistemic factors present in the subject's situation; such factors are seen as "encroaching" on the subject's epistemic standing (see, e.g., Fantl & McGrath, 2009; Hawthorne, 2004; Stanley, 2005; as a precursor see Rudner, 1953, p. 2 as well as some brief remarks in Lewis, 1996, p. 556). Not the only such non-epistemic factor but the most prominent one consists in the subject's practical stakes (see, e.g., Stanley, 2005). We don't need to explain in any detail here what the difference between epistemic and non-epistemic factors is (but see, e.g., Newton, 2020); practical stakes should come out as non-epistemic anyway, no matter what the proposed explanation of the difference is. Stakes-based SSI holds that two subjects can be in the same epistemic position with respect to some proposition but with different stakes for the two subjects so that one of them might know the proposition while the other might fail to know

<sup>1</sup>Counting knowledge itself as part of the epistemic position would be uninformative.

it. The higher or lower the stakes, the more or less demanding the epistemic standards for knowledge; higher or lower stakes make it harder or easier to know. Here is a very well-known pair of cases which can be used to illustrate the importance of stakes for knowledge (DeRose, 1992, p. 913):

*Bank Case A.* My wife and I are driving home on a Friday afternoon. We plan to stop at the bank on the way home to deposit our paychecks. But as we drive past the bank, we notice that the lines inside are very long, as they often are on Friday afternoons. Although we generally like to deposit our paychecks as soon as possible, it is not especially important in this case that they be deposited right away, so I suggest that we drive straight home and deposit our paychecks on Saturday morning. My wife says, “Maybe the bank won’t be open tomorrow. Lots of banks are closed on Saturdays.” I reply, “No, I know it’ll be open. I was just there two weeks ago on Saturday. It’s open until noon.”

*Bank Case B.* My wife and I drive past the bank on a Friday afternoon, as in Case A, and notice the long lines. I again suggest that we deposit our paychecks on Saturday morning, explaining that I was at the bank on Saturday morning only two weeks ago, and discovered that it was open until noon. But in this case, we have just written a very large and very important check. If our paychecks are not deposited into our checking account before Monday morning, the important check we wrote will bounce, leaving us in a *very* bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, “Banks do change their hours. Do you know the bank will be open tomorrow?” Remaining as confident as I was before that the bank will be open then, still, I reply, “Well, no. I’d better go in and make sure.”

Stakes-based SSI can explain the difference in knowledge, given the same epistemic position, as due to a difference in stakes.

We certainly have some intuitive notion of what is (practically) at stake for some subject in some situation (even if a lot of the details might vary from person to person). But still, given the huge amount of sophisticated work on stakes-based SSI, it is remarkable that the notion of stakes has not been discussed much in great detail at all so far (for exceptions, see Anderson & Hawthorne, 2019b, Baumann, 2016a; Fantl & McGrath, 2019; Goldberg, 2019; Russell, 2019; Worsnip, 2015).<sup>2</sup> Since what we can and should say about SSI depends a lot on what the notion of stakes used is, it is important to take a closer look at this notion. I will do so and start with our ordinary notion of stakes and modifying it in a plausible way when necessary. I will be presenting a new proposal and analysis of the notion of stakes. The quantitative question how much is at stake in some situation will play an important role for the clarification of the notion of a stake (however, only up to ordinal rankings, not to anything cardinal). An adequate development of that notion would have to be somewhat close to our ordinary notion but also be able to play a good role for SSI. I will aim at some reflective equilibrium between ordinary usage and the principles of SSI, with more weight on the former.<sup>3</sup> It will turn out that there is more

<sup>2</sup>Gardiner, 2021, pp. 483–484, 488 and 490 touches on the notion of stakes but not in the detail I have in mind here. She also does not claim much of an implication for questions about knowledge in particular (which I am interested in here). Finally, she draws a distinction between the stakes relating to the question whether one’s belief that  $p$  is true or not and the stakes relating to the question whether  $p$  is true or not; I doubt the relevance of such a distinction here, whatever exactly it might amount to.—I will discuss the approaches by Russell and Anderson and Hawthorne below; their views are most relevant to my discussion here. It would, however, go too far to also discuss the other views here, also because they don’t add that much to my discussion.—Thanks to a referee for this.

<sup>3</sup>If some notion of stakes leads to implausible results when applied by SSI, then this might either reflect a weakness of SSI or the inadequacy of the proposed notion of stakes. It will have to be decided case by case which the culprit is.—It seems clear to me that defenders of SSI are invoking a notion of stakes that is sufficiently close to ordinary notions of stakes.

than one kind of stakes. My aim is to clarify and discuss the notion of stakes, not really to discuss SSI more directly; that is a task for another occasion.

It is worth adding here that the notion of stakes does also play important roles for other epistemological theories, like contextualism (see, e.g., DeRose, 1992) or invariantism (see, e.g., Dinges, 2019, which also uses a lot of experimental research). Apart from this, the notion of stakes is also relevant to areas outside of epistemology (in theories of rational decision making, for instance). This should be kept in mind even if the focus here will be on SSI.

It turns out that it is crucial to distinguish between different kinds of stakes because only one type of stakes has relevance for knowledge. I start with a discussion of what I call “event-stakes” (Section 1) which also prepares the presentation of a notion of stakes relevant to knowledge, the knowledge-stakes; this idea of knowledge-stakes also raises some serious problems in need of discussion (Section 2). Section 3 discusses a third type of stakes which I call “action-stakes”; these are often run together with knowledge-stakes but need to be clearly distinguished from them; this distinction also helps to sort out some issues in the recent more detailed discussion of practical stakes.

## 1 | EVENT-STAKES

Consider the following situation. It has been predicted that a tornado will go through the area where I live. It might hit and damage the house where I live or spare it. If the first, then I would be in a very bad situation; not so if the latter. So, a lot is at stake for me in this situation, and there is nothing I or anyone can do about this.<sup>4</sup> For this kind of stakes, the attitudes and actions of the subject don’t matter. We can call them “event-stakes” (similar to Worsnip’s A-stakes; see Worsnip, 2015).<sup>5</sup> The above situation gives us the following outcome matrix (with “C<sub>1</sub>” and “C<sub>2</sub>” for worldly circumstances; the entries indicate utilities):

C <sub>1</sub> (Tornado doesn’t hit)	C <sub>2</sub> (Tornado hits)
fine	very bad

The greater the difference between the outcomes, the more is at stake.<sup>6</sup> This suggests that the following (absolute) value (the modulus) determines the stakes here (with “U” for utility, taking positive values):

$$|U(C_1) - U(C_2)|.$$

We need absolute values here (avoiding negative values for stakes) because, as one should not forget, outcomes can also be good. I might win the lottery (let us assume that would be a good thing):

C <sub>1</sub> (lottery loss)	C <sub>2</sub> (lottery win)
fine	very good

<sup>4</sup>What is at stake is, very plausibly, time-sensitive. With changing weather conditions, for instance, what is at stake at one time might not be at stake at another time. I will take this for granted here and not mention it explicitly in the following. Less trivially, stakes are also issue-relative; what is at stake in some situation is relative to some possibility under consideration. My stakes with respect to an impending tornado are different from my stakes with respect to the possible bouncing of my last cheque.

<sup>5</sup>Worsnip’s A-stakes are close, in different ways, to both my event-stakes and my action-stakes; his W-stakes come close to my knowledge-stakes.

<sup>6</sup>One could argue that the absolute “size” of the utilities also plays a role here and that the relative sizes need to be weighted by absolute sizes. We can leave this complication aside here; it won’t change anything about the basic argument of this paper.

More than just the utility difference seems relevant to the size of the stakes; probabilities also seem to play an important role.<sup>7</sup> The lower the probability that the tornado will hit or that my ticket will win, the lower the stakes. Given that this is plausible, we should weigh (multiply) the utility difference with the probability of the event that would either deteriorate or improve the subject's situation (the default situation). This would be the probability of  $C_2$  here. Hence, we can measure and express event-stakes in the following way:

$$\text{(Event-Stakes)} \quad |U(C_1) - U(C_2)|P(C_2).^8$$

We can generalise to more than two circumstances. Consider the following modification of the above lottery case where the winning possibilities split into two: winning and becoming happy versus winning and ending up miserable (as many real lottery winners do):

$$\begin{array}{ccc} C_1 \text{ (loss)} & C_2 \text{ (happy win)} & C_3 \text{ (miserable win)} \\ \text{fine} & \text{very good} & \text{very bad} \end{array}$$

One plausible way to measure the stakes here would be to add the probability-weighted utility differences for the default circumstance and each alternative circumstance. If  $C_1$  is the default circumstance here, this would give us the following measure for 3-circumstance event-stakes:

$$|U(C_1) - U(C_2)|P(C_2) + |U(C_1) - U(C_3)|P(C_3).$$

Sometimes, there might not be a default situation. In that case, we can choose the situation of the subject before the relevant event happens, the status quo ante, as the comparison point (SQA) for the alternative scenarios ( $C_k$ ). More generally, for  $n$  alternative scenarios we get:

(Generalised Event-Stakes)

$$\sum_{k=1}^n |U(\text{SQA}) - U(C_k)|P(C_k)$$

For simplicity's sake, we may focus on the 2-circumstances case here. This notion of event-stakes captures, as cases like the above illustrate, our ordinary, pre-theoretical notion of stakes quite well (at least in cases where our actions don't play any significant role).

<sup>7</sup>See fn.6.—Pinillos and Simpson (2014, sec. 5.2) report that different probabilities don't influence subjects' judgements about stakes. This is a point where I can see reasons to deviate from peoples' judgements for the sake of a better theory. I cannot go more into methodological issues like this one. See also Yip (2022, p. 1507, fn. 7) who also supports the inclusion of probabilities in our view of stakes.

<sup>8</sup>This implies that a case with an intermediate (not very high, not very low) utility difference but a high (or low) probability for  $C_2$  can present the same degree of stakes as a case with an intermediate probability but a big (or small) utility difference. I think that this is to be expected. Thanks to a referee for bringing up this question.—Perhaps we would also have to take risk aversion and similar phenomena into account; see, for example, Kahneman and Tversky (1979). This would complicate things without affecting the basic points here. Therefore, we can leave this aside.—Should one also include the probability of  $C_1$  here (thanks to a referee for this suggestion)? I don't think so. First, if we assume, as we easily could, that  $C_1$  and  $C_2$  are negations of each other,  $P(C_1)$  is implicitly represented here, even if not in the formula; it equals, of course,  $1 - P(C_2)$ . Second, what would be an alternative principle? Here is a different principle, one about the difference between expected utilities:  $(X) |U(C_1)P(C_1) - U(C_2)P(C_2)|$  (thanks, again, to a referee for this). However,  $(X)$  gives the wrong results in certain cases where my (Event Stakes) gets the right ones. For instance, consider a Tornado case where  $P(C_2)$  (a tornado hitting) = 1 (and thus  $P(C_1) = 0$ ). Then the stakes according to  $(X)$  equal  $U(C_2)$  which is low (given that the utility of  $C_2$  is so low), clearly against what is plausible in this case. On the other hand, if  $P(C_2) = 0$  then the stakes according to  $(X)$  equal  $U(C_1)$  which is high, again clearly against what is plausible in this case. Finally, consider the case where  $U(C_1) = 0.9$ ,  $U(C_2) = 0.1$ ,  $P(C_1) = 0.1$  and  $P(C_2) = 0.9$ . According to  $(X)$  the stakes would equal 0 here which is also very implausible. According to my (Event Stakes), the stakes would be equal to 0.72; this is not only above 0 but also plausibly high. As these examples illustrate, it won't help to add a minus in front of the formula  $(X)$ . The prospects for  $(X)$  and anything like that seem dim; why bet much on it when (Event Stakes) is getting the right results? What if it were the case that  $C_1$  and  $C_2$  are not negations of each other (and the probability of their disjunction  $<1$ )? I don't think this can be the case (given how SSI and the relevant cases are in fact constructed) but even if it were I could easily modify (Event Stakes), for instance by replacing  $P(C_2)$  with  $P(C_2)/[P(C_1) + P(C_2)]$ . We can leave this "side track" aside here. This whole discussion about  $(X)$  also applies, mutatis mutandis, to the following versions of stakes principles.—See also, in a very different context, Colyvan (2008) and Hájek (2006).

Do event-stakes have any relevance for knowledge? Going back to the tornado case above we can imagine the following pair of cases:

*Tornado Case A.* A tornado has been predicted to go through a coastal area. Sue just sold her house on the beach there. She has heard the weather reports and is aware of the small chance that the tornado will hit that house. If that should happen, she wouldn't be in a bad situation at all though she would feel somewhat bad for the new owners (if the tornado weren't to hit, she'd be just fine). Sue is confident that the tornado won't hit the house and she is right.

*Tornado Case B.* A tornado has been predicted to go through a coastal area. Sue just bought a house on the beach there. She has heard the weather reports and is aware of the small chance that the tornado will hit that house. If that should happen, she would be in a very bad situation (if the tornado weren't to hit, she'd be just fine). Sue is confident that the tornado won't hit the house and she is right.

Sue's event-stakes in case B are much higher than in case A (assuming that there is no big difference for her between owning the house or the money from the sale). Can we imagine that in case A Sue does know that the tornado won't hit the house while in case B she doesn't know that it won't hit the house? I find it hard to come up with a verdict here and determine what the SSIist could or should say here. This in itself is interesting; however, we can and should leave this open here (also because event-stakes have no role to play for SSI so far). However, we can use some ideas underlying (Event Stakes) in order to develop a notion of knowledge-stakes.

## 2 | KNOWLEDGE-STAKES

It is uncontroversial that a different kind of stakes (different from event-stakes) has played a much bigger role in the discussion of SSI: stakes where the outcomes and utilities also depend on the subject's actions.<sup>9</sup> The bank cases above fall into this group. We can represent them with the following kind of matrix (with "A" and "B" for feasible acts,<sup>10</sup> and "U1–4" for utilities of outcomes):

<sup>9</sup>There can be cases where the different acts have the same outcomes in each circumstance. And there can be cases where the different acts have different outcomes in each circumstance but with very small differences in utility between the outcomes of the different acts in each circumstance (small differences between the entries in each column). Finally, there can be cases where the latter differences are negligibly small as compared with the differences in the outcomes of the different acts given different circumstances (larger differences between the entries in the rows). For example (thanks to Matthias Brinkmann here), I can check now or later whether I have won the lottery. Whatever I do doesn't make a big difference at all. The big difference is whether I have won or not, no matter what I decide to do. This case falls into both of the last two groups of (non-exclusive) cases mentioned above. All three types of cases mentioned above could reasonably be considered to be cases where the relevant stakes are event-stakes (in a somewhat broader sense of the term). I will leave aside this complicating aspect in the following.

<sup>10</sup>We should not include all possible acts here. If we included, for instance, acts the subject has never considered (or never taken to be live possibilities), then we would get implausible results. Suppose for instance, that in the bank cases there is one possible act which the subject hasn't even thought of—e.g., performing some very complex financial and legal transaction—and which would have no bad outcomes if not  $p$  but a better outcome if  $p$  than the acts the subject is considering. It seems implausible to say that the possibility of such an act lowers the stakes in Bank Case B. So, judgements about stakes only vary between sets of feasible acts. And so will judgements about knowledge. I don't think this further subject-sensitivity creates any major problems. One would, of course, want to say more about what determines the feasibility of acts. But not here. See, however, also Fantl & McGrath (2009, p. 66) and Anderson (2018, pp. 17–20).

	$C_1$	$C_2$
A	$U_1$	$U_2$
B	$U_3$	$U_4$

Here is the matrix for Bank Case A (with “Sat” for “Saturday” and “Fri” for “Friday”):

Bank Case A	Sat open	Sat closed
go Sat	fine	annoying
go Fri	mildly annoying	mildly annoying

And here is the matrix for Bank Case B:

Bank Case B	Sat open	Sat closed
go Sat	fine	disastrous
go Fri	mildly annoying	mildly annoying

In both cases  $U_1 > U_3 = U_4 > U_2$ .<sup>11</sup> The intuitive verdict is that the stakes are much higher in case B than in case A. But what stakes? Well, here it’s the stakes of acting on the proposition (the “*p*-stakes” which I propose to call the “knowledge-stakes”) that the bank will be open on Saturday. The stakes are high in case B because  $U_2$  is much smaller than  $U_1$  and the other utilities; the stakes are low in case A because  $U_2$  is not that much smaller than  $U_1$  and the other utilities. Given that A (going to the bank on Saturday) is the act preferred if *p* (*the bank will be open on Saturday*), it is plausible to say that the stakes for that proposition are determined by the utility difference in the outcomes for A:  $U_1 - U_2$ . Let us go with the absolute value:  $|U_1 - U_2|$ .<sup>12</sup> Again, our intuitive verdicts about stakes also suggest that probabilities play a role here: The stakes in Bank Case B are much higher if there is a good chance that the bank will be closed on Saturday and much lower if there is only a very small chance that the bank will be closed on Saturday. All this suggests (again, leaving aside phenomena like risk aversion) the following measure for knowledge-stakes:

(Knowledge Stakes)  $|U(\text{acting on } p \text{ when } p \text{ is true}) - U(\text{acting on } p \text{ when } p \text{ is false})|P(\text{not } p)$ .<sup>13</sup>

In Bank Case A the knowledge-stakes are:  $|Utility \text{ of being fine} - Utility \text{ of being annoyed}| \times$  the probability that the bank will be closed on Saturday. In Bank Case B, however, the knowledge-stakes are:  $|Utility \text{ of being fine} - Utility \text{ of facing a disaster}| \times$  the probability that the bank will be closed on Saturday. If we assume in the bank cases that the probability of *not p* is not negligible, then we can say that the knowledge-stakes in Bank Case A are quite low while they are quite high in Bank Case B.

Given the strength of the subject’s epistemic position (evidence, etc.) with respect to *p* (that the bank will be open on Saturday) and given the stakes in both bank cases, we can move on to judging the subject’s knowledge of *p* or lack thereof—but only if we have a good idea about what pairings of strength of epistemic position and size of stakes would allow for knowledge and what pairings wouldn’t. Is the subject’s epistemic position strong enough for knowledge that the bank will be open on Saturday in Bank Case A, given the stakes? Are the stakes too high for knowledge (given his epistemic position) in Bank Case B? SSIists have so far not offered any account of the relation between the sufficiency of the strength of epistemic

<sup>11</sup>Usually in the debate about SSI the overall relative standing of  $U_2$  doesn’t receive many comments at all. However, there is no problem assuming it has the lowest value here.

<sup>12</sup>Whether the relevant outcomes lie above or below an indifference point (where the subject is indifferent between having or not having the outcome) does not matter here; only the relevant utility difference does.

<sup>13</sup>See fn. 6 above.

position concerning  $p$ , the size of the knowledge-stakes and knowledge that  $p$ . This is a theoretical desideratum for SSI.<sup>14</sup> However, perhaps we may still go with our intuitive verdict in cases like this one: in Bank Case A the subject’s epistemic position is strong enough for knowledge, given the stakes; in Bank Case B the subject’s epistemic position is not strong enough for knowledge, given the stakes.<sup>15</sup>

One interesting implication of (Knowledge Stakes) is that the stakes for some proposition can be quite high but the stakes for its negation quite low. The stakes for *The bank will be open on Saturday* are high but the stakes for *The bank won’t be open on Saturday* will be low. So, it can be the case that some proposition can be hard to know while its negation can be relatively easy to know. I take it that this is a welcome result.

We thus have an intuitively plausible idea about stakes (knowledge-stakes) that can, perhaps, be applied by SSI to deliver plausible verdicts about knowledge and the lack thereof. But more still needs to be said. First, there could be more than just two circumstances. In analogy to remarks above, the idea of knowledge-stakes can be generalised in the following way (with “Ak” ranging over alternative ways in which  $p$  can be false):

(Generalised Knowledge-Stakes)

$$\sum_{k=1}^n |U(\text{acting on } p, \text{ given } p) - U(\text{acting on } p, \text{ given } Ak) | P(Ak)$$

Generalising from two feasible acts to more than two is a bit more tricky because there might be two or more acts which lead to outcomes of the same utility, given  $p$ .<sup>16</sup> Consider this case. I can think of three ways to spend this Sunday afternoon: have a picnic with friends, take a walk with friends or go to the movies with friends. I have heard the weather reports which say that there is a good, though not very big, chance for rain; I am still confident that it won’t rain and I am right. Here is the matrix (with great > good > fine > somewhat annoying > awful):

Sunday A	No rain	rain
picnic	great	awful
walk	good	somewhat annoying
movies	fine	fine

Since I would go for the picnic option given no rain, the stakes for the proposition that it won’t rain are determined by the utilities and probabilities for the first row of the matrix, the picnic-row. But what if we boost the walk option a bit by adding the playing of some outdoors

<sup>14</sup>There are others: What conception of utilities, practical interests or practical reasons one has (e.g., more or less subjective or objective) makes a difference to one’s verdict about stakes (see Baril, 2019 and also Baumann, 2016b). Similar issues arise with respect to the interpretation of “probability”. We can leave these complications aside here (even though the SSlist will have to make a theoretical choice at some point).—There are also problems if one wants to stick with some principle of epistemic closure. Suppose I’m next to an icy lake, with my arch enemy Arce. My stakes for *The ice is thick enough to hold Arce* are low and I might know that the ice is thick enough to hold Arce. Suppose I do know this and also that we’re the same weight. Hence, I can infer from all this that the ice is thick enough to hold me, too. But my stakes for *The ice is thick enough to hold me* might be too high to allow for knowledge. On this, see also Anderson and Hawthorne (2019a), and a brief remark in Cohen (2019, p. 103) that lends itself to such worries; also see Baumann (2016a, 2016b).

<sup>15</sup>Using intuitive verdicts won’t help in all cases. Consider the proposition  $c$  that significant climate change is already under way. Acting on  $c$  would mean taking several drastic measures which would, however, seem like a big loss if *not-c* were true. How high are the knowledge-stakes then and is our epistemic position strong enough to allow for knowledge that climate change is already under way? While we don’t have clear intuitions about the impact of the stakes on the sufficiency of the epistemic position here, it still seems pretty clear that we know that climate change is under way. It seems that the stakes are simply irrelevant in a case like this. This is something the SSlist should worry about.

<sup>16</sup>This can also happen in 2-act cases but it is more probable in many-act cases.



hopscotch so that it will be as good, absent rain, as the picnic? Only one entry in the matrix would change:

Sunday B	No rain	rain
picnic	great	awful
walk	great	somewhat annoying
movies	fine	fine

What, then, are the stakes for the proposition that there won't be rain? Given what I said above, the stakes would be determined by the utilities and probabilities for the top-ranked act, given no rain. However, there is a tie here: There are two acts with equally good outcomes given no rain but different utilities, given rain. Which act then matters for the determination of the stakes?

I think intuition runs out of steam here. The SSIist should feel free to say one of at least two things. First, one could leave the stakes indeterminate in such a case. But then it might also turn out to be indeterminate whether I know that it won't rain: I might not be in a good enough epistemic position for knowledge given the stakes as determined by the picnic option, whereas I might be in a good enough epistemic position for knowledge given the stakes of the walk option. Leaving it open whether I know that it won't rain constitutes a high price for the SSIist; knowledge just doesn't seem "indeterminate" in such cases. There does not seem to be any motivation independent from adherence to SSI to think that a case like Sunday B makes judgement about knowledge so tricky and difficult.

Probably the SSIist should then go with a second option and determine the stakes in such cases in something like the following way. In Sunday B picnicking has a much worse outcome given *not-p* than walking. Since the two acts' outcomes are equally good given *p*, walking has the higher expected utility (and would thus be the preferable act). We can use expected utility as a tie breaker and let the utilities for walking determine the stakes. So, the stakes for the proposition that it won't rain in Sunday B would be:  $|great - somewhat\ annoying|P(rain)$ . We can use the same procedure if relative gains rather than losses result from choosing the corresponding act under *not-p* circumstances. Finally, there are mixed cases where some combinations of act and circumstance lead to relative losses and others to relative gains. The following matrix for three acts and three circumstances illustrates this kind of case (with higher positive numbers for higher utilities and lower negative numbers for greater disutilities<sup>17</sup>):

	C <sub>1</sub>	C <sub>2</sub>	C <sub>3</sub>
Act 1	4	5	2
Act 2	4	12	-3
Act 3	4	8	0

Again, the best act overall seems to be the one with the highest expected utility (if the probabilities for C<sub>1</sub>, C<sub>2</sub> and C<sub>3</sub> are 0.6, 0.2 and 0.2 respectively, then this would be Act 2). This fits with the idea above that in cases of ties for top act, given *p*, one should use expected utility as a tie breaker and let the act with the highest expected utility determine the stakes.

If it should happen that two acts are tied in both ways (both are top given *p*, and both have equal expected utility), then we're facing a further problem. Consider this 2-act and 3-circumstance scenario (probabilities indicated):

<sup>17</sup>It is only here that I deviate from the rest of the paper and use negative values for disutilities; I do this only for expository purposes.



	$C_1 (p:0.6)$	$C_2 (p:0.2)$	$C_3 (p:0.2)$
Act 1	4	22	-22
Act 2	4	5	-5

The expected utilities are the same for both acts (2.4) but applying the formula for knowledge-stakes above leads to different results: for Act 1 it gives us 44/5; for Act 2 it gives us 10/5. Depending on which act determines the stakes, we get very different results and thus possibly also very different results for our verdict about knowledge or the lack thereof. How serious is this problem?

Perhaps the SSlist can hope that such double-ties are rare and thus don't constitute more than a minor problem for their theory. Or should we, alternatively, assume that this would rather be a problem for the account of stakes proposed here? We shouldn't assume this just because it creates a problem for SSI; we're not presupposing the truth of SSI here. Do we have independent reasons to think that this would be a problem for the account of stakes proposed here? I don't think so, also because we cannot rely—at this level of development of the view—on intuitions any more: They remain silent on such cases.

However, the SSlist can do better than ruling that such cases are simply exceptional. They could have another quite plausible tie-breaker: the degree of risk-averseness of the subject. This would relativise stakes further to the subject but that's right up the general alley of SSI anyway. Here is the idea. If the subject is rather risk-averse, then Act 2 with small variations of utility differences determines the stakes; if the subject is rather risk-friendly, then Act 1 with larger variations of utility differences determines the stakes. We don't have to go into the details here, and we can also admit that there might still be possible cases where a verdict about stakes and knowledge is hard, even with this additional tie-breaker at hand. But at this point we can, I think, be confident that any further resulting indeterminacy wouldn't be worrisome and even to be expected, given the fact that knowledge (and stakes) allow for borderline cases.<sup>18</sup>

### 3 | ACTION-STAKES

It is worth considering two (recent) alternative formulations of the notion of stakes: the one proposed by Jeffrey Sanford Russell and the one proposed by Charity Anderson and John Hawthorne. Both consider not just acting on  $p$ , our target proposition, but also acting on alternative propositions.<sup>19</sup> While (Knowledge Stakes) works even if there is just one feasible act (if that is conceivable or possible), these alternative proposals require at least two alternative feasible acts (could this be an advantage for (Knowledge-Stakes)?). Here is Russell's proposal (see Russell, 2019, pp. 281–284) for the case of two actions (acting on  $p$  and acting on  $not-p$ ) and two circumstances (more complicated cases can be put aside here because this won't affect the basic points to be discussed):

(RUSSELL)

Stakes( $p$ ) =

$U(\text{acting on } not-p \text{ when } not-p \text{ is true}) - U(\text{acting on } p \text{ when } not-p \text{ is true}) /$

$U(\text{acting on } p \text{ when } p \text{ is true}) - U(\text{acting on } not-p \text{ when } p \text{ is true}).$

The stakes of  $p$  thus consist in the ratio of the utility difference for the feasible acts given  $not-p$  and the utility difference for the feasible acts given  $p$ . No indication about the measurement of utilities is given here (whether they are being assigned real values between 0 and 1 or

<sup>18</sup>The remarks in the last few paragraphs can and should also be understood as a proposal for further work for the SSlist. Since this paper aims at a clarification of the notion of stakes, we don't have to do that further work here.

<sup>19</sup>As one would do if one considered opportunity costs or opportunity benefits.

whether disutilities will be assigned negative values). Since it's the utility differences that matter according to (RUSSELL) we could also work with absolute values of the utility differences.

One might notice first that (RUSSELL) doesn't weigh the utilities with probabilities. This seems to me to be less than ideal for the adherent of SSI, for reasons already explained above (see also the brief remark supporting this reservation in Anderson & Hawthorne, 2019b, p. 244).

More important is a second point: (RUSSELL) considers alternative acts. Why do that? This question is also interesting because it will lead us to an important distinction between two different kinds of action-related stakes (in contrast to event-stakes). Consider the following case. I come to a fork in a road. I have to go either left or right (there is no other option). If I go left I'll either be very well off ('heaven') or very badly off ('hell'); the situation is the same for going right—and the probabilities (subjective, epistemic, objective) for LEFT (*heaven is to the left and hell to the right*) and RIGHT (*heaven is to the right, and hell to the left*) are both 0.5.<sup>20</sup> Russell's formula gives us the following verdict about this case:

Stakes(LEFT) =

$$\frac{U(\text{acting on RIGHT when RIGHT is true}) - U(\text{acting on LEFT when RIGHT is true})}{U(\text{acting on LEFT when LEFT is true}) - U(\text{acting on RIGHT when LEFT is true})} = \frac{U(\text{heaven}) - U(\text{hell})}{U(\text{heaven}) - U(\text{hell})} = 1.$$

This is far too low.<sup>21</sup> One can see this if one compares this with (RUSSELL)'s verdict for the proposition that the bank will be open on Saturday (SAT) in the high-stakes version of DeRose's bank case (see DeRose, 1992, p. 913; Russell, 2019, p. 283):

Stakes(SAT) =

$$\frac{U(\text{acting on not - SAT when not - SAT is true}) - U(\text{acting on SAT when not - SAT is true})}{U(\text{acting on SAT when SAT is true}) - U(\text{acting on not - SAT when SAT is true})}.$$

For Bank Case B the numerator is quite big while the denominator is comparatively small. The stakes are thus high and much bigger than 1 according to (RUSSELL) (for a similar example, see Russell, 2019, pp. 284–285). So, the stakes for LEFT (*heaven is to the left*) are much smaller than the stakes for SAT (*the bank will be open on Saturday*). However, this is very implausible and not the result any SSIist would want to accept. Using (Knowledge Stakes) above, however, together with plausible assumptions about heaven, hell and bouncing checks, gives us the verdict that the stakes for LEFT are much higher than for SAT. And this seems right.

So, one might be tempted to say that (RUSSELL) doesn't get the stakes right for certain kinds of propositions like LEFT. However, only to say this would be misleading—and not only because no account can be expected to cover all cases. Rather, this would be misleading because there is something (RUSSELL) is getting right even if it is missing something else. There is an interesting difference between a case like the bank case above and a case like the heaven-and-hell case. While the expected utilities (and the utility differences across circumstances) for the feasible acts differ drastically (given certain assumptions) in the bank case, they don't differ at all in the heaven-and-hell case. As far as the agent's choice is concerned, nothing from their perspective distinguishes between going left from going right in the heaven-and-hell case whereas going to the bank on Saturday is quite different in this respect from going on Friday. One can put this difference in terms of "action-stakes", that is, the stakes that are relevant and matter to, or are at least to some degree sensitive to, principles of making the best decision. More precisely, the action-stakes are higher (lower) the more (less) of a significant gap there is between the different feasible acts in their decision-theoretic ranking (by expected utility or utility differences if probabilities can be assigned, by some other principle if not). In the

<sup>20</sup>If the chances are 50–50, then, plausibly, the subject cannot have knowledge one way or the other. We can put this aside here because this case is only used to illustrate differences in verdicts about stakes.

<sup>21</sup>Remember that since utilities are not normalised in the case of (RUSSELL), stakes aren't either.

heaven-and-hell case the action-stakes are very low while they are very high in the bank case. (RUSSELL) gets this right (not going into the details here). However, since we're interested in SSI here and the stakes relevant to knowledge (or the lack thereof) we should have a different notion of action-related stakes, namely one that is captured by our (Knowledge Stakes) which seems to offer the better account of knowledge-stakes.<sup>22</sup>

In other words, if one is interested in knowledge-stakes—as adherents of SSI are primarily—then one should go for something like (Knowledge Stakes), that is, a formula that only deals with the target proposition. If, however, one is interested in action-stakes, then one should go for something like (RUSSELL), that is, a formula that compares in some way the utilities for the target proposition and for some alternative act(s). By the way, knowledge- and action-stakes can diverge, even by a lot, but often they don't and rather converge.

Anderson and Hawthorne propose a notion of stakes very similar to Russell's (but see also the brief explanation in Anderson & Hawthorne, 2019a, p. 112, which seems closer to (Knowledge Stakes)): The stakes of  $p$  are determined by the difference between the "maximum regret" for acting on  $p$  (the maximum loss under any circumstance for acting on  $p$  vs. not acting on  $p$ ) and the maximum regret for acting on *not-p*. This proposal amounts to something very much like (RUSSELL) (see Anderson & Hawthorne, 2019b, pp. 245–246; for some criticism see Goldberg, 2019 and the reply to Anderson and Hawthorne by Fantl & McGrath, 2019).<sup>23</sup> It gives verdicts on the bank cases and on the heaven-and-hell case which are very similar to those by (RUSSELL). So, similar points will apply here as for (RUSSELL).

Anderson and Hawthorne make an interesting objection to the likes of (Knowledge Stakes). Suppose (I am modifying their example a bit) I have to choose between two envelopes A and B. I know that they both contain the same amount of money and that they could contain either 1 or 1000 dollars. I am confident, though not completely sure, that it's 1000 dollars; I also have some evidence for this and I am right about this. What then are the (knowledge-) stakes for the proposition that they contain 1000 dollars (1 K)?

I know that which envelope I choose doesn't affect the outcome and also doesn't make a difference to the utility of the outcome. There is no decision-theoretic advantage to either of the two acts. In that sense one should expect the stakes (action-stakes) to be low. And indeed, they equal 0, according to both of Anderson and Hawthorne's proposals (see Anderson & Hawthorne, 2019b, pp. 245–246); (RUSSELL) doesn't even deliver a verdict because "0/0" is undefined.

What does (Knowledge-Stakes) tell us here? We need to distinguish two cases here. First, if the envelopes contain different amounts of money (against the assumption in Anderson and Hawthorne's example), then the stakes for 1 K are quite high (if we can identify a unique top-ranked act; if not, there can be an issue of indeterminacy to be solved; see above). According to SSI (if it incorporates (Knowledge Stakes)), it would then make sense to say that in such a case it's hard to be in a position to know, whereas it is much easier in a situation where the monetary differences are small. The divergence of verdicts by the alternative notions of stakes discussed above and by (Knowledge Stakes) fits nicely with the distinction between action-stakes and knowledge-stakes: The former are very low (minimal) while the latter are high. If what matters for knowledge, according to the adherent of SSI, are the knowledge-stakes, then (Knowledge Stakes) is giving the right verdicts again. However, if, second, the amounts of money in the two envelopes is the same, as in Anderson and Hawthorne's example, then (Knowledge-Stakes)

<sup>22</sup>So, I argue that even though Russell doesn't get the knowledge-stakes right, he is getting something else right: the action-stakes. The same holds for Anderson and Hawthorne (see below).

<sup>23</sup>Anderson and Hawthorne (2019b, pp. 245–246) also mention an alternative: The stakes are determined by the differential between the maximal gap between the utilities of outcomes for different acts, given  $p$ , on the one side, and the maximal gap between the utilities of outcomes for different acts, given *not p*, on the other side. According to this, the stakes are high in Bank Case B and low in the heaven-and-hell case. In general, the latter proposal is quite close to Anderson and Hawthorne's first proposal or to Russell's, too. Anderson and Hawthorne do not seem fully committed to any of these proposals.

agrees with the alternative accounts offered by Anderson and Hawthorne, and Russell. However, this is not surprising. A defender of (Knowledge Stakes) shouldn't go this far and rather insist that the stakes under discussion here are event-stakes: that they concern a case where taking any of the feasible acts has no significant effect on the outcome (nor its utility). And in that case, it would be very doubtful, to say the least, whether the stakes have any connection with knowledge (see above).

## 4 | CONCLUSION

If one wants to defend SSI, then one has to say more about the nature and function of the non-epistemic subject-sensitive factors. The notion of stakes can use more clarification. The above contains a new proposal of thinking about stakes which is, I think, quite different from what has been proposed so far. As I hope has become clear in this paper, at no point does the account proposed here differ in any (significant) way from ordinary usage or a defensible reflective equilibrium between ordinary usage and epistemological theory. It also turns out that the stakes-based way of explaining subject-sensitivity has issues which need to be addressed (see above). How much this helps SSI is a further question; SSI faces other challenges, too. Apart from that, there is the possibility of defending SSI without relying (so much) on stakes.<sup>24,25</sup>

## ORCID

Peter Baumann  <https://orcid.org/0000-0003-0991-6288>

## REFERENCES

- Anderson, C. (2018) Putting fallibilism to work. In: McHugh, C., Way, J. & Whiting, D. (Eds.) *Normativity: epistemic and practical*. Oxford: Oxford University Press, pp. 12–25.
- Anderson, C. & Hawthorne, J. (2019a) Pragmatic encroachment and closure. In: Kim, B. & McGrath, M. (Eds.) *Pragmatic encroachment in epistemology*. London etc: Routledge, pp. 107–115.
- Anderson, C. & Hawthorne, J. (2019b) Knowledge, practical adequacy, and stakes. In: Gendler, T.S. & Hawthorne, J. (Eds.) *Oxford studies in epistemology*, Vol. 6. Oxford: Oxford University Press, pp. 234–257.
- Baril, A. (2019) Pragmatic encroachment and practical reasons. In: Kim, B. & McGrath, M. (Eds.) *Pragmatic encroachment in epistemology*. London: Routledge, pp. 56–68.
- Baumann, P. (2016a) Knowledge across contexts. A problem for subject-sensitive invariantism. *Philosophical Review*, 55, 363–380.
- Baumann, P. (2016b) *Epistemic contextualism: a defense*. Oxford: Oxford University Press.
- Cohen, S. (2019) Pragmatic encroachment and having reasons. In: Kim, B. & McGrath, M. (Eds.) *Pragmatic encroachment in epistemology*. London: Routledge, pp. 101–106.
- Colyvan, M. (2008) Relative expectation theory. *The Journal of Philosophy*, 105, 37–44.
- DeRose, K. (1992) Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52, 913–929.
- Dinges, A. (2019) *Knowledge, stakes and error: a psychological account*. Frankfurt: M. Klostermann.
- Fantl, J. & McGrath, M. (2009) *Knowledge in an uncertain world*. Oxford: Oxford University Press.
- Fantl, J. & McGrath, M. (2019) Clarifying pragmatic encroachment: a reply to Charity Anderson and John Hawthorne on knowledge, practical adequacy, and stakes. In: Gendler, T.S. & Hawthorne, J. (Eds.) *Oxford studies in epistemology*, Vol. 6. Oxford: Oxford University Press, pp. 258–266.
- Gardiner, G. (2021) Relevance and risk: how the relevant alternatives framework models the epistemology of risk. *Synthese*, 199, 481–511.
- Goldberg, S.C. (2019) Stakes, practical adequacy, and the epistemic significance of double-checking. In: Gendler, T.S. & Hawthorne, J. (Eds.) *Oxford studies in epistemology*, Vol. 6. Oxford: Oxford University Press, pp. 267–278.

<sup>24</sup>As they say: this would be a topic for a different paper.—It would also go beyond the limits of this paper to explore the connections between the proposal made here and different theories of rational choice (thanks to a referee for this).

<sup>25</sup>I am grateful for discussion and comments by Matthias Brinkmann and Alexander Dinges. Juan S. Piñeros Glasscock and Baron Reed commented on an earlier version of this paper at the Central Division Meeting of the *American Philosophical Association* in 2020—many thanks to them. I am also grateful to referees and to audiences at the University of Köln and the University of Osnabrück.

- Hájek, A. (2006) In memory of Richard Jeffrey: some reminiscences and some reflections on *the logic of decision*. *Philosophy of Science*, 73, 947–958.
- Hawthorne, J. (2004) *Knowledge and lotteries*. Oxford: Clarendon.
- Kahneman, D. & Tversky, A. (1979) Prospect theory: an analysis of decision under risk. *Econometrica*, 47, 263–291.
- Lewis, D. (1996) Elusive knowledge. *Australasian Journal of Philosophy*, 74, 549–567.
- Newton, L. (2020) The epistemic significance of modal factors. *Synthese*, 199, 227–248.
- Pinillos, N.Á. & Simpson, S. (2014) Experimental evidence supporting anti-intellectualism about knowledge. In: Beebe, J.R. (Ed.) *Advances in experimental epistemology*. London: Bloomsbury, pp. 9–43.
- Rudner, R. (1953) The scientist *qua* scientist makes value judgments. *Philosophy of Science*, 20(1), 1–6.
- Russell, J.S. (2019) How much is at stake for the pragmatic encroacher. In: Gendler, T.S. & Hawthorne, J. (Eds.) *Oxford studies in epistemology*, Vol. 6. Oxford: Oxford University Press, pp. 279–286.
- Stanley, J. (2005) *Knowledge and practical interests*. Oxford: Clarendon.
- Worsnip, A. (2015) Two kinds of stakes. *Pacific Philosophical Quarterly*, 96, 307–324.
- Yip, B. (2022) Assertion, stakes and expected blameworthiness: an insensitive invariantist solution to the bank cases. *Erkenntnis*, 87, 1501–1519.

**How to cite this article:** Baumann, P. (2024) What's a(t) stake? On stakes, encroachers, knowledge. *Theoria*, 90(1), 109–121. Available from: <https://doi.org/10.1111/theo.12512>