

‘Knowledge and Objective Chance’¹

John Hawthorne & Maria Lasonen-Aarnio

For *Williamson on Knowledge*

1.

We think we have lots of substantial knowledge about the future. But contemporary wisdom has it that indeterminism prevails in such a way that just about any proposition about the future has a non-zero objective chance of being false.^{2, 3} What should one do about this? One, pessimistic, reaction is scepticism about knowledge of the future. We think this should be something of a last resort, especially since this scepticism is likely to infect alleged knowledge of the present and past. One anti-sceptical strategy is to pin our hopes on determinism, conceding that knowledge of the future is unavailable in an indeterministic world. This is not satisfying either: we would rather not be hostage to empirical fortune in the way that this strategy recommends. A final strategy, one that we shall explore in this paper, is one of reconciliation: knowledge of a proposition is compatible with a subject’s belief having a non-zero objective chance of error.⁴ Following Williamson, we are interested in tying knowledge to the presence or absence of error in close cases, and so we shall explore the connections between knowledge and objective chance within such a framework.

We don’t want to get tangled up here in complications involved in attempting to formulate a necessary and sufficient condition for knowledge in terms of safety. Instead, we will assume the following rough and ready necessary condition: a subject knows *P* only if she could not easily have falsely believed *P*.⁵ Assuming that easiness is to be spelt

¹ This joint project was prompted by the discovery that we had independently come up with the ‘High Chance – Close Possibility Argument’ as well as the same taxonomy of possible solutions. The choice of formulation stays close to the version in Lasonen-Aarnio’s DPhil thesis. We are most grateful to Timothy Williamson for extended discussion, which helped sharpen the central arguments and draw our attention to the broader significance of the contrast between global risk and local excellence.

² We are not concerned in this paper to explore the connection between knowledge and epistemic notions of chance, of the sort encoded by epistemic uses of modals. Note that in seeking reconciliation between knowledge and objective chance of error, we are not claiming that ordinary utterances of the form ‘I know *P* and it might be that not-*P*’ would be true, since it is epistemic modals that figure in such claims. Even if knowledge that *P* is compatible with an objective chance of error, that does not mean it is compatible with an epistemic chance of error.

³ Not all notions of objective chance require indeterminism. For example, the conception of chance that one finds in statistical mechanics makes no such requirement. We are interested in how such notions relate to knowledge. But in this paper we pursue the more limited goal of exploring how to salvage knowledge of the future in the face of indeterminism. (For more on chance in statistical mechanics, see Albert 2000).

⁴ For reasons that will become clear below, a non-zero objective chance of error is not quite the same as a non-zero objective chance that the proposition believed is false.

⁵ Safety theorists are aware that it is slightly more accurate to mention methods in formulating a safety requirement: One knows *P* only if one could not easily have believed *P* falsely by a relevantly similar method. Nothing we say turns on this, so for the sake of exposition, we set aside this refinement. Further, safety principles gain more power if we extend the principle as follows: One knows *P* only if one could not easily have had a false belief using a relevantly similar method (where the false belief may be in a different

out in terms of close possible worlds, a subject knows P only if there is no close possible world in which she falsely believes P .⁶ (We shall call the set of close possible worlds the ‘safety set of worlds’.)

2.

If the safety theorist wants to avoid widespread scepticism about knowledge of the future, he must be careful to disambiguate modal locutions such as ‘could easily have been the case that’ and ‘there is a close possible world in which it is the case that’. In particular, not every world in which the same or very similar conditions obtain at a relevant time can be close. The problem is most acute in indeterministic worlds. But first let us introduce some terminology, in particular a distinction between *branching* and *close* possible worlds.

Branching possibilities are possible worlds sharing their histories with the actual world up to and including some time t . More precisely, let $H_{w,t}$ be a proposition stating the entire history of a world w up to and including a time t , and entailing the laws of nature at w . In other words, $H_{w,t}$ is necessarily equivalent to the conjunction of propositions about w up to and including t together with the laws of nature.⁷ A branching possibility at a time t and world w is any possible world in which $H_{w,t}$ is true.⁸ Indeterminism can be stated in terms of the notion of a branching possibility: a world is indeterministic if and only if it has branching counterfactual possibilities. Obviously, a counterfactual possibility can be branching at a time t but no longer be branching at a later time t' . For instance, at a time t before a coin was tossed twice, there were branching possibilities in which the coin landed heads on both of these two tosses, but once it has landed tails on the first toss, possibilities in which it lands heads on both tosses are no longer branching.

Close possibilities are just close, in whatever sense of close is relevant for knowledge. We will allow the closeness of a world to vary from one time to another, and speak of close possibilities at a world w and time t . Relativising closeness to times allows for danger and safety to be time-relative.⁹

So far nothing has been said about whether being a branching possibility at a time t entails being a close possibility at t . One might certainly be tempted to accept the

proposition). We are inclined to endorse some such principle, but have no need for the extra explanatory power here.

⁶ Williamson (2000: 126-127) uses a formulation along these lines, though he speaks of cases rather than worlds.

⁷ We will assume that at any time, the chance of the laws of nature changing is 0.

⁸ If there were only finitely many branching possibilities at a time t , then we could say that branching possibilities at t are those possibilities that have a non-zero chance of being the case, and that a proposition is true in some branching possibility at t if and only if it has a non-zero chance of being true at t . However, if there are infinitely many equally probable branching possibilities, then each must either have an infinitesimally small chance, or a chance of 0.

⁹ See Williamson (2000: 124) on the idea that safety is time-relative. Following Williamson’s own emphasis on cases, it may ultimately be best to articulate ideas about safety in terms of centered worlds. We shall return to this theme later.

entailment. Consider extremely unlikely and bizarre ‘quantum’ events such as the event that a marble I drop tunnels through the whole house and lands on the ground underneath, leaving the matter it penetrates intact. The situation can be re-described as a cosmic lottery with immensely many tickets. In this lottery, holding a winning ticket means having one’s marble tunnel through the house. Re-describing the situation as a kind of lottery invites thinking of the actual world as being surrounded by a sphere of equally close worlds, among them worlds in which the marble does tunnel. But for any event e with a non-zero chance, it is possible to partition the space of possibilities into equally likely subcases, some being ones in which e occurs.¹⁰ One might thus be led to the following principle, where ‘ $\text{Ch}_{w,t}(P)$ ’ stands for ‘the chance of P in w at t ’:

Chance – Close World Principle

For all worlds w , times t , and propositions P , if $\text{Ch}_{w,t}(P) > 0$, then at t in w there is a close possibility in which P .

This principle is disastrous for knowledge of a chancy future. For assume that at t a subject believes P , the proposition that her marble won’t tunnel through the house, but that $\sim P$ has a non-zero chance of being true. Then, the conjunction $(\sim P \ \& \ H_{@,t})$ has a non-zero chance of being true at t , since the chance of $H_{@,t}$ at t is 1. Because all worlds in which $H_{@,t}$ is true are branching possibilities at t , there is a non-zero chance of being in a branching possibility in which $\sim P$. By the above principle, some such branching possibility is close. But the subject believes P in all branching possibilities at t . And so there will be a close possible world in which the subject falsely believes P .

If the safety theorist wants to avoid scepticism about knowledge of the future, he is pressured to deny the *Chance – Close World Principle*.¹¹ Denying this principle amounts to denying that all branching possibilities at a time t are close at t . For assume that there is a world w , time t , and proposition P such that $\text{Ch}_{w,t}(P) > 0$, but P is not true in any close world at t . This entails that the conjunction $(H_{w,t} \ \& \ P)$ isn’t true in any close world at t . But the conjunction $(H_{w,t} \ \& \ P)$ is true in some branching worlds at t . So there are branching worlds that are not close and generally, being a branching possibility at a time t doesn’t entail being close at t (or, indeed, at any other time).

What about the other direction of the entailment, from being a close possibility at a time t to being a branching possibility at t ? The safety theorist also has ample reason to deny this entailment. Assume that a subject believes P at t , and P concerns events at t or at earlier times (i.e. the present or the past). If the safety set of worlds for the subject’s belief in P only consists of branching close possibilities at t , then the truth of P would seem to guarantee that her belief is safe – for normally, holding everything fixed up to and including t will both hold fixed everything a subject believes at t , and the truth of all propositions concerning t and earlier times.¹² To make meeting the safety requirement non-trivial, the safety set of worlds for the subject’s belief must contain worlds that are not close and branching at t . One might indeed want to allow for close worlds that aren’t branching at *any* time. Nevertheless, *at least* close branching worlds at t belong to the

¹⁰ Cf. Vogel (1990: 21–22); Hawthorne (2004: 14–16).

¹¹ Similar issues arise if one wishes to save the truth of ordinary counterfactuals in a chancy world. For discussion, see Hawthorne (2005b).

¹² But see below for knowledge of contingent a priori truths.

safety set of worlds for a subject's belief in a proposition P at t , whatever else might be part of that set.

In this section we have considered a distinction the safety-theorist needs to draw in order to accommodate knowledge of the future in indeterministic worlds. In particular, the safety theorist must allow for branching possible worlds that are not close. Our main argument will be to show that given a plausible connection between chance and modal closeness, this move isn't sufficient to avoid scepticism about knowledge of the future. However, before presenting our argument, we briefly discuss an alternative argument relying on a closure principle for knowledge.

3.

The argument from Closure Under Conjunction makes use of the following closure principle:

Closure Under Conjunction

For any subject s , and any propositions P, Q , if s knows P and s knows Q and s comes to believe $(P \& Q)$ based on competent deduction from P, Q , while retaining knowledge of both P and Q throughout, then s knows $(P \& Q)$.

Assume that a subject Suzy knows each proposition in a set $S_p = \{P_1, P_2, \dots, P_n\}$ of high-chance propositions. Assume, for instance, that S_p consists of all the propositions that Suzy knows about the future. Suzy retains knowledge of each proposition, and competently deduces conjunctions of these, finally arriving at the conjunction $(P_1 \& P_2 \& \dots \& P_n)$ at a time t . In so far as she satisfies the antecedent of Closure Under Conjunction at each step, Suzy now knows $(P_1 \& P_2 \& \dots \& P_n)$. But assume also that the chance of $(P_1 \& P_2 \& \dots \& P_n)$ at t is very low. There seems to be something seriously wrong with allowing subjects to know such low-chance propositions.¹³

Care must be taken in attempting to capture the intuitive epistemological principle that is violated by allowing subjects to know low-chance propositions such as $(P_1 \& P_2 \& \dots \& P_n)$. On a somewhat tempting diagnosis, the problem is just that they are highly unlikely. And so:

For all worlds w , times t , subjects s , and propositions P , if $\text{Ch}_{w,t}(P)$ is low, then s does not know P at t in w .

But this principle has counterexamples created by knowledge of contingent a priori truths.

For instance, assume that a lottery draw is about to take place, and Suzy fixes the reference of 'Lucky' as 'the winner of the lottery'. Suzy knows that the lottery has 1000000 tickets, each owned by a different person, that it is fair, and that exactly one person will win. And assume that the winner is in fact John. Despite the fact that its

¹³ Hawthorne (2005a) gives this argument.

reference was fixed by a description, ‘Lucky’ is a singular term, and the proposition expressed by ‘Lucky will win the lottery’ is not the necessarily true proposition that whoever will win the lottery will win the lottery but rather, a contingent, singular proposition about John. In so far as Suzy can have this singular proposition as the content of her propositional attitudes in the situation described, and ‘the winner of the lottery’ picks out different people in different branching worlds at t , Suzy believes different singular propositions in different branching worlds. Moreover, at a time t prior to the draw she will be able to have a priori knowledge of the singular proposition expressed by ‘Lucky will win the lottery’, despite the fact that at t the chance that Lucky, i.e. John, will win the lottery is very low.¹⁴

To avoid ruling out knowledge of low-chance contingent a priori truths, the epistemological principle that is violated by allowing subjects to know overwhelmingly unlikely conjunctions must be revised. This can be done by formulating a principle in terms of the chance of a belief-episode of a subject expressing a true proposition, rather than the chance of the proposition in fact expressed by the belief-episode:

Low Chance

For all worlds w , times t , subjects s , belief-episodes B , and propositions P , if at t s 's belief-episode B expresses proposition P , at t the chance that B expresses a true proposition is low, and at t s is not inadmissibly connected to the future, then s does not know P at t in w .¹⁵

The clause about inadmissible connections restricts the principle to cases in which there are no time-travellers from the future, clairvoyance by backwards causation etc. *Low Chance* is compatible with knowledge of low-chance, contingent a priori truths. For though at t the chance that Lucky (i.e. John) will win the lottery is low, Suzy's belief-episode does not have a low chance of expressing a truth. In different branching possibilities at t ‘Lucky’ will refer to different persons, in each world to whoever wins the race in *that* world. And in none of these branching worlds does Suzy's belief-episode express a false proposition.

It is also worth noting in passing that low-chance, contingent a priori truths create trouble for the so-called Principal Principle.¹⁶ This principle entails that if at a time t a subject is certain that the chance of a certain outcome is x , and is not inadmissibly connected to the future, then the credence she assigns to the outcome at t should likewise be x . Assume that in the case described Suzy is certain, and knows, that for any one ticket in the lottery, the objective chance of that ticket winning is 0.000001. She is also certain that whoever ‘Lucky’ refers to, the objective chance of that person winning at t , prior to the draw, is 0.000001. By the Principal Principle, the credence Suzy should assign to the proposition that Lucky will win the lottery is 0.000001, since she is certain that that is the objective chance that Lucky will win. But this isn't right, for given that she is certain that the lottery has exactly one winner, and that ‘Lucky’ refers to this winner, the credence

¹⁴ Consider similarly known propositions of the form [P iff actually P].

¹⁵ We speak of ‘inadmissibly connected’ instead of the more usual ‘inadmissible evidence’ because we do not wish either to deny knowledge of the future or that knowledge is evidence.

¹⁶ See Lewis (1986).

Suzy ought to assign to the proposition that Lucky will win is 1. To say the least, the Principal Principle needs to be revised in light of contingent a priori knowledge.¹⁷

Setting such issues aside, there are various ways of responding to the argument from Closure Under Conjunction. One is to fault the argument's reliance on closure.¹⁸ Another is to impose a constraint on what a subject can know at any one time that prevents problem-generating cases from arising. In particular, if $\{B_1, B_2, \dots, B_n\}$ is the set of all belief-episodes giving rise to beliefs in propositions that a subject knows at a time t , then the chance at t that at least one of these belief-episodes expresses a false proposition cannot be high. This allows for knowledge of the future, but restricts how much a subject can know at any one time. It also prevents subjects from pooling together their knowledge, since the conjunction of everything various subjects know might have a low chance.

Neither of these responses will help with the argument we present below. Our argument doesn't rely in any way on closure. It is also compatible with the suggested restriction on what can be known by a subject at any one time, since it only assumes that one subject can know one proposition about the future. The argument does rely on the assumption that there should be no restriction on how many subjects can know things about the future at any time t . If there is knowledge of the future in the first place, then it should be possible for each subject in a very large set of subjects to know some high-chance proposition about the future, even if the conjunction of all the known propositions has a very low chance of being true. It also relies on a principle stating that high-chance propositions are true in some close worlds.

4.

What we call the *High Chance – Close Possibility* argument rests on the following principle connecting high chance and modal closeness:

High Chance-Close Possibility Principle (HCCP principle)

For any world w , time t , and proposition P , if $\text{Ch}_{w,t}(P)$ is high, then there is a close branching possibility at t in w in which P .

This principle states that high-chance propositions are true in some close worlds. It entails that at any time t , the chance of being in a world that is close at t is high. For let $\text{CloseBranching}_{w,t}$ be the proposition that is true in all worlds that are close and branching at t in w . The High Chance – Close Possibility Principle is equivalent to a principle

¹⁷ There is similarly no Principal Principle style connection between objective chance and evidential probability.

¹⁸ Assume a threshold model of outright belief on which assigning a sufficiently high credence to a proposition entails believing it. Someone might try to deny that the antecedent of Closure Under Conjunction can ever be satisfied when a conjunction is low-chance, simply because rational subjects never believe such conjunctions. The problem with this response is that though credences filter down across multi-premise entailments, they might not do this quickly enough, since the credence a subject assigns to a proposition need not equal its chance.

stating that for any world w and time t , $\text{CloseBranching}_{w,t}$ has a high chance.¹⁹ Hence, the chance of being in a world that is both close and branching at t is high, which entails that the chance of being in a world that is close at t is high.

Prima facie at least, it looks as though the safety-theorist ought to subscribe to the HCCP principle. It is overwhelmingly natural to think that high-chance events could easily have occurred. If there is a high objective chance of a washing machine breaking down, it is hard to deny that it could easily break. If there is a high chance that I will die climbing a certain route, then it is not safe for me to embark on the climb. Similarly, mutatis mutandis, for belief and error. Assuming the modal analysis – which is part of the framework of safety theory – events that could easily occur, or events that are not safe from not occurring, will occur in some close worlds. This validates the HCCP principle.

The problem created by the HCCP principle is the following. Suppose that I have just dropped a marble. Assume, contra the sceptic, that I can know, while it is in midair, that it will land on the floor. Moreover, this piece of knowledge shouldn't, it seems, depend on what else is going on in the world outside what we can assume to be a closed system consisting of me and my marble – in particular, it shouldn't depend on how many other subjects hold beliefs about other falling marbles. The existence of such subjects should have no effect on my epistemic position. Then, it looks to be nomologically possible for there to be a very large set of propositions $S_p = \{P_1, P_2, \dots, P_n\}$, a world w , a time t , and a set of subjects $s = \{s_1, s_2, \dots, s_n\}$ such that

- i) at t in w , each proposition $P_i \in S_p$ is known by exactly one subject $s_i \in s$,
- ii) each proposition $P_i \in S_p$ is about a time t' after t ,
- iii) at t the chance of the conjunction of all propositions in S_p , $(P_1 \& P_2 \& \dots \& P_n)$, is low.

(For example, consider a world where many subjects have just dropped a marble, each believing simultaneously that their recently dropped marble is floor-bound. Suppose that each marble has a non-zero chance of tunnelling through the floor and that there is probabilistic independence.²⁰ If the number of subject-marble pairs is large enough, the conjunction of all the propositions believed by the subjects about their marbles will be highly unlikely.)

¹⁹ Proof: To see that the HCCP principle entails that $\text{Ch}_{w,t}(\text{CloseBranching}_{w,t})$ is high, assume that $\text{Ch}_{w,t}(\text{CloseBranching}_{w,t})$ is low. Then, $\text{Ch}_{w,t}(\sim\text{CloseBranching}_{w,t})$ is high. By the HCCP principle, $\text{CloseBranching}_{w,t}$ is false in some world that is close and branching at t . But $(\text{CloseBranching}_{w,t})$ is the proposition true in all close branching worlds at t . Hence, the assumption must be false and $\text{Ch}_{w,t}(\text{CloseBranching}_{w,t})$ must be high. To see that the proposition that $\text{Ch}_{w,t}(\text{CloseBranching}_{w,t})$ is high entails the HCCP principle, assume that $\text{Ch}_{w,t}(\text{CloseBranching}_{w,t})$ is high and that the HCCP principle is false: $\text{Ch}_{w,t}(P)$ is high, but P is not true in any close branching world at t . Then, P is only true in worlds in which $(\sim\text{CloseBranching}_{w,t})$ is true, and P will entail $\sim\text{CloseBranching}_{w,t}$. Then, $\text{Ch}_{w,t}(\sim\text{CloseBranching}_{w,t}) \geq \text{Ch}_{w,t}(P)$, and $\text{Ch}_{w,t}(\sim\text{CloseBranching}_{w,t})$ is high. But this contradicts our assumption that $\text{Ch}_{w,t}(\text{CloseBranching}_{w,t})$ is high. Hence, the HCCP principle is equivalent to the claim that $\text{Ch}_{w,t}(\text{CloseBranching}_{w,t})$ is high.

²⁰ The simplifying assumption of independence is not essential for our argument. Suppose for example that instead of the tunnelling possibilities being probabilistically independent, there was 'entanglement' that disrupted independence. So long as the chance of any given marble tunnelling is low but the chance of at least one marble tunnelling is high, the argument in the text applies.

By safety, if each of our subjects knows the relevant proposition in S_p , then none will falsely believe that proposition in any world that is close and branching at t . In the present case, this requires that for any subject s_i , proposition P_i will be true in all close branching worlds at t . Here is why. For reasons given above and having to do with contingent a priori truths, it cannot be assumed that a subject holds exactly the same belief at t in all worlds that are branching possibilities at t : externalism about the contents of thoughts extends outside a subject's skin not only to the environment, but also to the future. But not all thoughts supervene on future facts. If I now believe that this marble is floor-bound, there is no reason to think that just which content I entertain depends on the future. If each subject in the case described above is to satisfy safety, then the relevant marble-belief of each subject must be true in all branching worlds at t . Hence, if i) above is true, and each proposition in S_p is known by some subject at t , the conjunction ($P_1 \& P_2 \& \dots \& P_n$) must be true in all close branching worlds at t .

Because the conjunction ($P_1 \& P_2 \& \dots \& P_n$) has a low chance of being true at t , then by the HCCP principle its negation, ($\sim P_1$ or $\sim P_2$ or \dots or $\sim P_n$), will be true in some close branching world at t : at least one of the marbles tunnels in a close branching world at t . Then, by safety, not every subject can know of his or her marble that it will come to rest on the floor, and our assumption is false.

It is essential to the problem-generating case that each subject is in the same epistemic position. This is hard to challenge. It would be ad hoc to select a privileged set of knowing subjects. But denying all of the subjects knowledge would be equally implausible, for it would make knowledge depend on facts that it seems it ought not to depend on. In particular, how could facts such as how many other subjects hold beliefs about the future have any impact on whether one subject knows that her marble won't tunnel? Odd counterfactuals would come out as true: 'Had those other subjects held beliefs about their marbles, I would not have known that my marble would come to rest on the floor'. (There is, of course, the sceptical alternative of denying that subjects can ever know propositions with non-zero chances of being false. But we are currently in the anti-sceptical business.)

Despite this, someone might bite the bullet and allow that what a subject knows depends on how many other subjects there are in her world holding beliefs about the future, placing the following restriction on the set of propositions that can be known by subjects at a time: the conjunction of all the propositions in the set must have a high chance at that time.²¹ This would make knowledge much more sparse than it appears to be, vindicating at least partial scepticism. And the problem of choosing just which beliefs constitute knowledge is by no means minor. Do I know that my marble is floor-bound or that my pen is? That I am meeting John on Tuesday or that I am attending a seminar on Wednesday?

Similar problems arise for knowledge of the present and past, though in those case trouble cannot be created *just* by appeal to the HCCP principle, for if a proposition P is true at a time t , then its chance at t is 1, and there is no branching world at t in which it is false. Rather, what would be needed is some principle connecting high chance at a time t with close possibilities at a slightly later time t' . Suffice it to say that in many cases such connections seem very plausible. For instance, if a lottery draw has recently taken place

²¹ For the ingenious model in Williamson (2005: 485–487) to work as a way of reconciling widespread knowledge with widespread risk, it would have to impose some such requirement.

and I won, but the result has not been announced, and I believe all along that I am winner, there is still now a close world in which I falsely believe I won, given that I had a high chance of losing prior to the draw. In light of this, the case discussed above could be altered in the following way: it is 12.01, and for each subject, a different marble was dropped at noon somewhere where it could not be observed by anyone. At 12.01 each subject believes of his or her marble that it is on the floor.

5.

We have been discussing cases in which each belief in a large set of beliefs is held simultaneously. But perhaps the lesson can be extended to diachronic cases. In particular, there are chains of knowledge acquisition, transmission and preservation stretching across time that are globally risky despite consisting of locally excellent steps. Because it is highly likely that something goes wrong at some step in the chain, one might worry that some principle very much like the High Chance – Close Possibility Principle forces the safety theorist to fault some step in the chain. But, the worry is, this would mean giving up some plausible principle concerning knowledge, such as single premise closure. Devotees of closure are pushed to transmit knowledge across chains that are, overall, highly risky.

Here are a few examples of the sorts of chains we have in mind. First, the *deduction chain*. Assume that at t_0 a subject knows a true proposition P_0 , and deduces from it a proposition P_1 . She is an excellent deducer, but prior to the deduction, there is a 0.000001 chance that her deductive capacities will lead her astray. For whenever she is about to perform a deduction, a random lottery occurs in her brain, and in the unlucky case she infers a falsehood without realising her predicament. Now she goes on to perform 99999 more successful deductions, reaching a true proposition P_n . At t_0 there is a high chance of going wrong at some step of a 1000000-step deduction. Further, let us assume that at t_0 there is a high chance of forming a false belief as the result of a 1000000-step deduction. If close branching worlds at t_0 are a subset of the safety set of worlds for the subject's belief in P_n , then the safety theorist will have to say that the subject does not know proposition P_n . Consequently, single premise closure will at least be less than perfectly true: there will be a step in the deduction at which the subject knows a proposition P_i , competently deduces from it a proposition P_{i+1} , but fails to know proposition P_{i+1} .

Similar cases are easy to construct for testimonial transmission of knowledge from one subject to another, and preservative memory. In the latter case, a subject knows a proposition P at a time t_0 , and retains a belief in P over a long period of time, up to and including t_n . During each i -length interval of time beginning from t_0 there is a non-zero chance that the subject's preservative memory malfunctions in a way which leads her to form a false belief. Assume that at t_0 there is a high chance that some malfunction of memory will occur during the period $t_0 \dots t_n$ and further, that at t_0 there is a high chance that by t_n the subject's preservative memory will have produced in her a false belief. Call this the *memory chain*. Again, if branching worlds at t_0 are a subset of the safety-class of worlds for the subject's belief in P at t_n , then she will fail to know P at t_n . If the safety theorist denies that the subject knows P at t_n , one might worry that some plausible

principle along the following lines will be violated: if a subject knows a proposition P at t_i , and competently preserves P in memory between t_i and t_{i+1} , then at t_{i+1} the subject knows P .

However, in such diachronic cases it is more difficult to find a plausible principle connecting chance and modal closeness as a means to showing that some step in the chain fails to transmit or produce knowledge. Just like chance, closeness and safety are time-relative. For instance, assume that yesterday there was a high chance that at noon today I am drugged and undergo a hallucination of a tree. Despite this, at noon today worlds in which I am drugged need not be part of the safety set of worlds. If the person who would have drugged me happened to die and at noon I veridically perceive a tree, then there will be no close world with relevantly similar initial conditions in which I form a false belief. Hence, it cannot be simply assumed that in the chain cases described close branching worlds at a time t_0 are a subset of the safety set of worlds for a belief a subject holds at a time t_n .

Whether or not knowledge is allowed to transmit across the sorts of chains described, there are puzzles to solve. In each chain-case described, some knowledge transmission principle forces attributing knowledge across the chain. These principles state that knowledge is extended over competent deduction, competent preservation of belief in memory, or competent testimonial transmission. If knowledge doesn't transmit across the chains described, some seemingly very plausible principle about knowledge such as single premise closure has to give. If it does, various oddities will have to be dealt with. Among these are peculiar asymmetries of the following sort, created by the seemingly magical effect had by being appropriately linked to a knowledge-transmitting or knowledge-producing chain.

Take the deduction chain described above. Let us assume that Suzy undergoes such a chain of deductions, and that she also knows the relevant facts about the chain. At t_0 she knows that there is a high chance that if she attempts to carry out 1000000 deductions, she will end up believing a falsehood. For simplicity, assume also that it is certain that she will enter a chain involving 1000000 inferences. By single premise closure, as long as Suzy in fact deduces competently throughout, it looks as though in the end of the chain, at t_n , she knows proposition P_n .²² Now consider an onlooker, John. At t_0 it is certain on John's evidence that Suzy will enter a 1000000-step inferential chain. John knows the relevant facts about Suzy's deductive abilities, and at t_0 it is highly likely on John's evidence that the proposition in the end of Suzy's chain is false. Moreover, John has no intuitions about the truth of the propositions deduced by Suzy, is incapable of deducing, and could not spot possible mistakes made by Suzy. At t_n , John learns that the chain has led Suzy to believe proposition P_n . It looks as though at t_n the probability on John's evidence that the proposition in the end of Suzy's chain is true is exactly the same as the probability on his evidence at t_0 that the proposition in the end of Suzy's chain is true, and equals the probability at t_n on John's evidence that P_n is true. John, unlike Suzy, is not in a position to know P_n .

Suppose that John is disposed to believe whatever Suzy comes up with at the end of long deductions by eavesdropping. Presumably this does not generate knowledge any

²² Those that are inclined to deploy the rhetoric of defeaters as an all purpose fix to epistemological problems will surely be inclined to trot out that rhetoric here. We believe that the concept of a defeater is both woefully underdeveloped and woefully overdeployed, but here is not the place to pursue the matter.

more than John could get knowledge from a highly unreliable barometer (albeit one that sometimes delivers correct information thanks to myriad locally excellent steps). Safety does nothing to explain the asymmetry between John and Suzy. Given that the truth of both Suzy's and John's beliefs depend on Suzy's deductions, the truth of the relevant beliefs march in step across close possible worlds. If Suzy's beliefs are safe, so are John's. So why does John not know P_n ? The peculiarity of the situation is only sharpened by assuming a transmission principle for testimony stating that if a subject s knows a proposition P , and another subject s^* comes to believe P solely based on s 's competent testimonial transmission of P , then s^* knows P . For then John could come to know P_n based on Suzy's testimony, even if he could not know by eavesdropping! Further, prior to Suzy's testimony, John would not be in a position to know that if Suzy testifies that P_n , then P_n . But once the testimonial transmission has taken place, John is in a position to know P_n . Chain related issues are indeed bewildering, but we shall not pursue them further here.

6.

The High Chance – Close Possibility Principle combines with the following assumption to create sceptical trouble: knowledge of P entails that there is no close possible world in which the subject holds a relevant false belief.

There are three main non-sceptical options for the safety theorist. The first is to liberalise safety by conceding that knowledge is tolerant to holding a false belief at a small proportion of close worlds: it is enough that a subject avoids false belief in *most* close worlds. The second is to revise the modal account of safety by framing safety not in terms of possible worlds, but in terms of subject-centered worlds. The third is to deny the HCCP principle, thereby forsaking any systematic connection between objective chance and modal closeness. In the remainder of this paper, we briefly outline these options. (We leave it as an exercise to the reader to figure out why adding a dose of contextualism to safety doesn't help.)

(i)

Let 'most worlds safety' refer to the modified safety requirement on which knowledge is compatible with holding false beliefs in some close possible worlds. The most worlds safety theorist is much better equipped to accommodate knowledge of the future in indeterministic worlds. Within such a framework, there is no need to give up the *Chance – Close World Principle* in the first place.

Most worlds safety at least encourages one to think that one can know that a ticket in a lottery will lose (so long as it does lose). But perhaps this result is to be welcomed. After all, it is desperately difficult to explain why such propositions cannot be known even though much of our alleged knowledge is subject to similar risks.

A potentially more worrying entailment is that multi premise closure fails. For even if P is true in most close worlds and Q are true in most close worlds, the conjunction ($P \& Q$) might fail to be true in most close worlds. Though we cannot argue for the claim here, the most worlds safety theorist even has trouble with single premise

closure, for the simple reason that even competent deduction from a single premise can bring with it an element of risk or danger.²³

If multi premise closure is false, then any reasoning from more than one premise becomes problematic, no matter how infallible the deductive capacities of a subject are. And even the simplest practical reasoning typically proceeds from two premises:

1. P .
2. If P , I should φ .

Therefore,

3. I should φ .

Most worlds safety allows situations to arise in which a subject knows P ; knows that if P , she should φ ; competently deduces and comes to believe that she should φ from these premises, but does not know that she should φ .

And there are further problems. A safety requirement for knowledge, we take it, is supposed to eliminate the problematic sort of epistemic luck involved in Gettier-cases. But the revised safety requirement seems unable to rule out standard Gettier-cases, for the reason that any belief in a true, high-chance proposition looks to be safe.²⁴ But beliefs in high-chance propositions can be Gettiered. Take the following example. A pyromaniac is about to strike a match. At a time t prior to striking the match she infers, and thereby comes to believe, that it will light when struck from her knowledge that it is a dry match of a brand that has always lit for her when dry and struck. There is a small chance that the particular match she holds won't light by friction when struck. And in fact, the match doesn't light by friction. But it lights nevertheless, because of a burst of rare Q -radiation.²⁵ This looks very similar to Russell's stopped clock case in which a subject fails to have knowledge of the time based on having looked at a normally reliable clock that just happened to stop 12 hours before. The pyromaniac has a justified belief that is due to a bout of luck, but seems to lack knowledge. Nevertheless, her belief is actually true, and true in most close worlds, since in most close worlds the match lights in the normal way by friction.

(ii)

Upon closer inspection, the idea that knowledge requires avoiding false belief in close worlds seems implausible to begin with. When evaluating whether a subject could easily have held a false belief, we are ordinarily only interested in whether things could easily have been a certain way as regards *that subject*. The worry is that even if closeness is relativised to times, global similarities elsewhere can outweigh radical, local dissimilarities regarding the situation of one subject. For instance, allowing one subject to be envatted, but keeping everything else fixed as far as possible, might make for a possible world that is, overall, close to the actual world.

²³ This is argued in work in progress by Lasonen-Aarnio ('Single Premise Deduction and Risk').

²⁴ Lasonen-Aarnio (*forthcoming*).

²⁵ The case is modified from Skyrms (1967: 383).

In effect, Williamson gives modal truth-conditions for safe belief in terms of close cases rather than worlds. He states that a case is ‘like a possible world, but with a distinguished subject and time’, or what Lewis calls a ‘centered world’.²⁶ Centering on times was already implicit in our discussion above, and now the idea is to center further on subjects.

We will follow Lewis in thinking of centered worlds as pairs of worlds and space-time points in those worlds.²⁷ For instance, $\langle (l, t), w \rangle$ is a centering of world w under the space-time point (l, t) . Some space-time points are occupied by subjects. Let such occupied centered worlds be *subject-centered worlds*. Surrounding each subject-centered world is a possibility-space consisting of subject-centered worlds with relevantly similar centerings, occupied by the same subject. The new safety requirement is stated in terms of avoiding false belief in close subject-centered worlds. To get the right anti-sceptical results, the closeness-relation for subject-centered worlds must assign special weight to propositions believed by the subject occupying the centre of a world. So, for instance, if Suzy is to know that her marble is floor-bound, there can be no close subject-centered world in which her marble tunnels, though there might well be close worlds in which other subjects’ marbles tunnel. Indeed, assuming some analogue of the High Chance – Close Possibility Principle, some tunnelling will have to take place in close subject-centered worlds.²⁸

One problem with the resulting position is that as long as it is assumed that high-chance propositions are true in some close subject-centered worlds, there is still a limit to how much any one subject can know at any one time. In particular, the chance of the conjunction of everything a subject knows at any one time cannot be low, whether or not the subject believes the conjunction. (As a result, different subjects cannot pool their knowledge together.) Allowing the epistemic status of a subject’s belief to depend in this way on how many other beliefs the subject holds in propositions with non-zero chances of falsity isn’t as absurd as allowing the epistemic status of a subject’s belief to depend on facts about how many *other* subjects hold beliefs in propositions about the future. But the sceptical worry remains that knowledge would become sparse. And is there any non *ad hoc* way of deciding just which propositions a subject knows at any one time?

Assume, for instance, that Suzy holds beliefs about very many marbles at one time t . At t each marble has exactly the same low chance of tunnelling, but the chance that at least one will tunnel is high. Moreover, each of Suzy’s beliefs is based on the same sort of evidence. It would seem *ad hoc* to allow for some, but not all, of her beliefs to count as knowledge. But it would also be odd not to allow for her to know *any* of the propositions she believes about where marbles are headed. For then the following odd circumstance would obtain: had she only formed beliefs about n marbles, and not $n + m$ marbles, she would have known of each of n marbles that it would come to rest on the floor.

A natural development of the subject-centered approach gives special weight to propositions that have special bearing on the life of the subject in the centre of the world. If I believe of many people that they will live through the night, and have the same kind

²⁶ Williamson (2000: 52).

²⁷ Lewis (1983: 149).

²⁸ Here is a general form of the principle applied to subject-centred worlds: For any time t , location l , world w and proposition P , if $\text{Ch}_{w,t}(P)$ is high, then there is a centered world close to $\langle (l, t), w \rangle$ in which P .

of evidence for each, then it is the ones nearest and dearest that are known by me to be survivors. This protects much of the knowledge we care about, though at the cost of delivering very odd sounding truths: ‘I know Bill, my friend, will live through the night, but not Fred, since I don’t care about him.’ ‘I would have know that Fred would live through the night had I cared more about him.’ Nor will such a move make the restriction problem go away entirely. After all, my friendships may be too plentiful.

The problems we have considering might be escaped by some more centering. If the propositions believed by a subject must be assigned special weight in any case, why not also center around belief-episodes? Let *super-centered* worlds be worlds that are centered not only around a subject and time, but also a belief episode of the subject. A subject’s belief in a proposition P is safe if and only if there is no close super-centered world in which the relevant belief-episode gives rise to a false belief (or a mere illusion of content).²⁹

One likely casualty of super-centering is multi premise closure. Even if a subject knows each of P_1, \dots, P_n , there is no guarantee that she is in a position to know their conjunction. For when evaluating whether she knows the conjunction ($P_1 \& \dots \& P_n$), we center around a belief-episode that is distinct from any of the belief-episodes giving rise to the subject’s beliefs in P_1, \dots, P_n . Moreover, if one adopts a suitable analogue of High Chance – Close Possibility while using super-centering to safeguard widespread knowledge, failures of multi premise close are bound to occur. This is far from devastating. Given the alternatives that precede and follow, an anti-sceptical solution grounded in super-centering is a serious option.

(iii)

A final option is simply to reject the High Chance – Close Possibility Principle. This would make safety powerless to prohibit knowledge that contravenes Low Chance. Allowing such knowledge risks disrupting intuitive connections between knowledge, ease of mistake and danger. After all, as noted above, it is hard to deny that if there is a high chance of an event of a certain type occurring, then an event of that type could easily have occurred. And it is similarly hard to deny that high chance events are in danger of occurring.³⁰ Rejecting the HCCP principle threatens to sever the very connections between knowledge and objective danger that were seemingly integral to motivating a safety requirement on knowledge in the first place. (Note that such motivations were front and center in Williamson’s own discussions of safety.³¹) Of course, it is conceivable that one could reject the HCCP principle but still prohibit low chance knowledge on grounds other than safety. Unless combined with a much more widespread scepticism, this would make multi premise closure untenable. Moreover, it would considerably dilute the explanatory work that safety was fit to perform in epistemological theorizing.

²⁹ Interestingly, once one has super-centering, it is rather less clear whether one needs the distinction between close worlds in which relevantly similar methods are deployed and those in which they are not. Why not just make methods a criterion of closeness of super-centered worlds?

³⁰ It is at least less clear that in the ordinary sense of hope and danger, any non-zero chance event has a hope of/is in danger of occurring. (After all, I may say that there is no hope of you finding a ring that you dropped on a vast beach, and that there is no danger of dying from a simple medical procedure, owing to the fact that the relevant chances are so small they can be discounted.)

³¹ See, for example, Williamson (2000: 123-124).

Conclusion

The sceptical pressures posed by chance are not to be underestimated. It is very tempting to surrender to scepticism about knowledge of the future when faced with chance-based considerations, thereby initiating descent into a more widespread sceptical abyss. The epistemically brave of heart will have to super-center, embrace knowledge in the face of danger or take refuge in something like most worlds safety. Such choices cannot be made by casual head counting on judgments about cases. After all, it is clear that resisting scepticism will require giving up a range of highly intuitive judgments about knowledge, which in turn casts doubt on a simple case-driven methodology. The best we can do is to reflect on what is structurally important about knowledge to our cognitive lives.

Bibliography

Albert, David Z.

2000 *Time and Chance*, Harvard: Harvard University Press.

Hawthorne, John

2004 *Knowledge and Lotteries*, Knowledge and Lotteries, Oxford: Oxford University Press.

2005a 'Knowledge and Evidence', *Philosophy and Phenomenological Research*, Vol. LXX, No. 2, pp. 452-458.

2005b 'Chance and Counterfactuals', *Philosophy and Phenomenological Research*, Vol. LXX, No. 2, pp. 396-405.

Lasonen-Aarnio, Maria

2007 Review of *Epistemic Luck*, *European Journal of Analytic Philosophy*, Vol. 2.

Lewis, David

1983 'Attitudes *Dedicto* and *De Se*' (and postscript), in *Philosophical Papers*, Volume I, Oxford: Oxford University Press, pp. 133 – 156.

1986 'A Subjectivist's Guide to Objective Chance', *Philosophical Papers*, Volume II, Oxford: Oxford University Press, pp. 83 – 113.

Pritchard, Duncan

2005 *Epistemic Luck*, Oxford University Press, Oxford.

Sainsbury, R. M.

1997 'Easy Possibilities', *Philosophy and Phenomenological Research* Vol. 57, No. 4, pp. 907 – 919.

Skyrms, F. B.

1967 'The Explication of "X Knows that p"', *Journal of Philosophy* 64: 373 – 389.

Sosa, Ernest

1999a 'How to Defeat Opposition to Moore', in Tomberlin (ed.) *Philosophical Perspectives* 13, pp. 141–153.

1999b 'How Must Knowledge Be Modally Related to What Is Known?', *Philosophical Topics* 26: 373–384.

Vogel, Jonathan

1990 'Are There Counterexamples to the Closure Principle?', in M. Roth and G. Ross (eds.), *Doubling: Contemporary Perspectives on Skepticism*, Dordrecht: Kluwer Academic Publishers.

Williamson, Timothy

1994 *Vagueness*, London & New York: Routledge.

2000 *Knowledge and its Limits*, Oxford: Oxford University Press.

2005 'Reply to Hawthorne', *Philosophy and Phenomenological Research*, Vol. LXX, No. 2, pp. 476-487.