A Chancy Grim Reaper Scenario

Ben Bavar • 13 Apr 2025

The Classic Grim Reaper Paradox

GRIM REAPER: Fred is alive at 10:00am. For every positive integer n and number $\tau(n)$ of minutes that have elapsed since 10:00am such that $\tau(n) = 60/n$, there is a distinct Grim Reaper. This Reaper swings its scythe at Fred $\tau(n)$ minutes past 10:00am if and only if no Reaper does so earlier. (Henceforth I shall use "swing(s)" to mean "swing(s) its scythe," except when I feel like writing the longer version. In the scenario at hand, the scythe may be assumed to be swung at Fred.) For any time t between 10:00am (exclusive) and 11:00am (inclusive), Fred is alive at t if and only if no Reaper swings earlier than t.

Suppose we count the Reapers from 1 onward, beginning with the Reaper poised to swing exactly one hour past 10:00am. The n^{th} Reaper is the one poised to swing $\tau(n)$ minutes past 10:00am. There are infinitely many Reapers poised to swing before the n^{th} Reaper does.

There are only two possible cases, which I've named SUPPOSITION I and SUPPOSITION II. Either supposition leads to contradiction, as shown below.

SUPPOSITION I: there exists a positive integer n such that some Reaper swings $\tau(n)$ minutes past 10:00am.

Let i be one such positive integer n.

1. The i^{th} Reaper swings.
2. If the i^{th} Reaper swings, then for no positive integer j does the $(i+j)^{\text{th}}$ Reaper swing.
3. If for no positive integer j does the $(i+j)^{\text{th}}$ Reaper swing, then the $(i+1)^{\text{th}}$ Reaper does and does not swing.
4. Therefore, the $(i+1)^{\text{th}}$ Reaper does and does not swing.

That the i^{th} Reaper swings follows from SUPPOSITION I and the definitions of i and "the i^{th} Reaper." According to SUPPOSITION I, there exists a positive integer n such that some Reaper swings $\tau(n)$ minutes past 10:00am. But since i

is defined as one such positive integer n (i.e., one such that some Reaper swings $\tau(n)$ minutes past 10:00am), there must be a Reaper that swings $\tau(i)$ minutes past 10:00am. This Reaper can only be the one we call "the i^{th} Reaper," as this is by definition the only one poised to swing $\tau(i)$ minutes past 10:00am. Thus, the i^{th} Reaper swings.

Premise 2 is a consequence of the stipulations of GRIM REAPER. Any given Reaper swings at Fred if and only if no Reaper swings at Fred earlier. Thus, the i^{th} Reaper swings $\tau(i)$ minutes past 10:00am only if no Reaper swings less than $\tau(i)$ minutes past 10:00am. But for any positive integer j, the $(i+j)^{\text{th}}$ Reaper is poised to swing less than $\tau(i)$ minutes past 10:00am. To see this, note that $i + j > i$ for any positive j. So for positive j, the duration $\tau(i+j) = 60/(i+j)$ minutes is less than $60/i = \tau(i)$ minutes. Therefore, the i^{th} Reaper swings only if the $(i+j)^{\text{th}}$ Reaper does not, for any positive integer j.

I'll give an example to illustrate the justification for premise 3. Let's say $i = 1$. Then $i + 1 = 2$, which makes Reaper 2 the $(i+1)^{\text{th}}$ Reaper. So to establish premise 3, we need to argue that, without loss of generality:

3*. If for no positive integer j does the $(1+j)^{\text{th}}$ Reaper swing, then Reaper 2 does and does not swing.

Now, from the supposition that the $(1+j)^{\text{th}}$ Reaper does not swing for any positive integer j, it follows that neither Reaper 2 nor Reaper 3 nor any Reaper assigned an integer greater than 1 swings. Thus we see already that the second conjunct ("Reaper 2 does not swing") of the consequent of 3* follows from the antecedent of 3*. What about the first conjunct, i.e., "Reaper 2 does swing"? As observed a moment ago, no Reaper assigned an integer greater than 1 swings, assuming the antecedent of 3* is true. Thus, neither does any Reaper assigned a number greater than 2 (e.g., Reaper 3, 4, or 5). But if no Reaper assigned a number greater than 2 swings, Reaper 2 does, for Reaper 2 swings if no Reaper swings before it does. So we get the contradictory result that Reaper 2 does, yet does not, swing.

The conclusion of the argument above follows from the premises by modus ponens/conditional proof. So we've shown that SUPPOSITION I entails a contradiction. Next we derive a contradiction from SUPPOSITION II.

SUPPOSITION II: there exists no positive integer n such that any Reaper swings $\tau(n)$ minutes past 10:00am.

1. None of the Reapers swings.
2. If none swings, one does.

3. So, one of the Reapers swings.
4. Therefore, none of the Reapers swings, but one of the Reapers does.

SUPPOSITION II, conjoined with GRIM REAPER, straightforwardly implies premise 1. By stipulation, for each Reaper there is some positive integer n such that the Reaper swings $\tau(n)$ minutes past 10:00am, if the Reaper swings at all. If there exists no such n for any Reaper, none of the Reapers swings.

But if none swings, then some Reaper r (in fact, every Reaper) is such that none of the ones poised to swing before r swings does so. So r swings, for a Reaper swings as long as no Reaper has done so yet. This establishes premise 2.

We may apply modus ponens to infer 3 from 1 and 2. Then, to reach the desired contradictory conclusion, we simply conjoin 1 and 3.

The Unsatisfiable Pair Diagnosis

Causal finitists such as Rob Koons claim that the GRIM REAPER scenario is impossible because no infinite causal regress is possible. But Joe Schmid and Alex Malpass beg to differ. According to Schmid and Malpass's *unsatisfiable pair diagnosis* (UPD) of infinitary paradoxes like GRIM REAPER, what makes these scenarios impossible is rather that each scenario requires two conditions to be met that cannot be met at once. Here's one way to loosely characterize those conditions in the case of GRIM REAPER: i) there is an infinite causal regress, and ii) each cause in the regress is the (non-)occurrence of a distinct Grim Reaper's swing, with the swing occurring if and only if (and because) none of the infinitely many earlier Reapers' swings occurs. Given (i) and (ii) we run into the contradictions discussed above. But (i) entails no apparent contradiction on its own, so it's plausibly possible despite the impossibility of certain infinite regresses.

In reply to the UPD (which Alex Pruss calls "the Inconsistent Pair response" to the Grim Reaper paradox), Alex Pruss develops a new version of the Grim Reaper scenario, to which he believes no satisfactory Inconsistent Pair response can be given. It goes something like this:

UNRELIABLE REAPER: Fred is alive at 10:00am. Let $\tau(m) = 60/m$ for every integer m. There exists some nonnegative integer i such that

a) for every integer $n > i$, a distinct Grim Reaper swings its scythe at Fred $\tau(n)$ minutes past 10:00am if and only if Fred is alive then, and

b) for any time t between $10{:}00\mathrm{am}$ (exclusive) and $\tau(i+1)$ minutes past $10{:}00\mathrm{am}$ (inclusive), Fred is alive at t if and only if no Reaper's swing before t kills him.

The probability that the n^{th} Reaper's swing would kill Fred is $1/n^p$, where p is the same positive real number for every Reaper.

It is unclear whether Pruss intends for there to be a distinct Reaper for every positive integer, or merely one for every integer greater than some nonnegative integer i. But in either case, we can describe the situation by saying, as I have, that there's a distinct Reaper for every integer greater than some nonnegative integer i. If there is a distinct Reaper for every positive integer, then $i = 0$. Otherwise $i > 0$.

I added stipulation (b) to avoid going down a rabbit trail. Without this stipulation, Fred may be alive at t even though some Reaper kills him before t. Fred might have died earlier than t and been resurrected by t, such that he could be killed again. But that's irrelevant in this dialectical context. We are exclusively interested in the paradox that might arise when we suppose a Reaper attempts to do something if and only if no Reaper has already succeeded in doing that thing. So, really, whether Fred is alive is of no significance to us. What matters is whether a Reaper previously caused him to die. For this reason, rather than stipulating (b) separately from (a), we could just replace "Fred is alive then" with "no Reaper's swing has already killed Fred" in (a).

If we modify the scenario, we can even ignore altogether whether any Reaper's swing kills Fred. Let's say there's no one around for the Reapers to kill, but they're all eager to swing their scythes anyway. Now suppose that, the greater the n, the weaker the n^{th} Reaper. More precisely, the probability is $1/n^p$ that the n^{th} Reaper would succeed in swinging its scythe (i.e., manage to lift it and make a swinging motion with it). Suppose in addition that the n^{th} Reaper attempts to swing if and only if no Reaper successfully swings earlier. Then there may be a chance that roughly the same paradox we teased out of Grim Reaper will arise.

But how? So far we haven't even explicated how my first formulation of Unreliable Reaper is paradoxical. So let's do that. For the rest of this post, I'll stick with that initial formulation.

We shall use the locution "Reaper n non-fatally swings" to say that Reaper n swings its scythe at Fred but fails to kill Fred thereby. To say that Reaper n does not non-fatally swing is thus not to affirm that Reaper n's swing kills Fred, but to deny that Reaper n swings and yet Fred survives. This denial is consistent with Reaper n's refraining from swinging altogether. Let j be some nonnegative integer i such that (a) & (b). UNRELIABLE REAPER leads to contradiction on the following supposition.

SUPPOSITION: there exists no integer $n > j$ such that, for every integer $k \geq n$, the k^{th} Reaper non-fatally swings.

Let $x > j$ be an integer.

1. For some integer $k \geq x$, the k^{th} Reaper does not non-fatally swing.
2. If, for some integer $k \geq x$, the k^{th} Reaper does not non-fatally swing, then there exists an integer $n > j$ such that, for every integer $k \geq n$, the k^{th} Reaper non-fatally swings.
3. So, there exists an integer $n > j$ such that, for every integer $k \geq n$, the k^{th} Reaper non-fatally swings.
4. Therefore, there does and does not exist an integer $n > j$ such that, for every integer $k \geq n$, the k^{th} Reaper non-fatally swings.

SUPPOSITION entails that, for every integer $n > j$, there exists some integer $k \geq n$ such that the k^{th} Reaper does not non-fatally swing. Therefore, since by stipulation x is an integer greater than j, there is an integer $k \geq x$ such that Reaper k does not non-fatally swing. Let y be one such integer k, so that Reaper y does not non-fatally swing. Observe that, by definition, any Reaper that does not non-fatally swing either does not swing or kills Fred upon swinging. This means that the y^{th} Reaper either does not swing or kills Fred upon swinging.

Suppose the y^{th} Reaper does not swing. Then Fred must not be alive when the y^{th} Reaper would have swung; otherwise, it would have swung! But if Fred is not alive when the y^{th} Reaper would have swung, there exists an integer $n > y$ such that the n^{th} Reaper fatally swings. Let z be one such integer n. For every integer $n \geq z + 1$, the n^{th} Reaper non-fatally swings. So, there exists an integer $n > j$—namely $z + 1$—such that, for every integer $k \geq n$, the k^{th} Reaper non-fatally swings.

Suppose instead that the y^{th} Reaper kills Fred upon swinging at time t. In that case, Fred must be alive at t. But if he's still alive then, every Reaper that swings earlier than t non-fatally swings. In other words, for every integer $k \geq y$, Reaper k non-fatally swings. Recall that $y \geq x > j$.

So, no matter what, there exists an integer $n > j$ such that, for every integer $k \geq n$, the k^{th} Reaper non-fatally swings. This directly contradicts SUPPOSITION, according to which no such n exists. So, just as planned, we've deduced a contradiction.

Pruss's Objection to the UPD

Notice that we only get that last contradiction by assuming SUPPOSITION is true. This is the crucial difference between UNRELIABLE REAPER and GRIM REAPER; the latter scenario gives rise to a contradiction no matter what. In UNRELIABLE REAPER it's logically possible to avoid paradox, since a Reaper is not guaranteed to kill Fred just as long as it swings its scythe at him. So in UNRELIABLE REAPER, some ways of filling in the details give us a possible situation; others don't—the ones that generate contradictions. Taking into consideration this feature of the thought experiment, Pruss tries to back the proponent of the UPD into a corner by showing they're committed to an inexplicable modal discontinuity. In particular, he suggests, though not as explicitly as he might've, that the only available Inconsistent Pair response to UNRELIABLE REAPER entails UNRELIABLE REAPER is impossible if $p \leq 1$ but possible if $p > 1$. Let's see if he's right. (You may rightly be curious what he finds counterintuitive about this entailment and whether that intuition of his should be trusted, but that's a discussion for another time.)

His central argument is this. The paradox, i.e., contradiction, in UNRELIABLE REAPER can only be avoided if SUPPOSITION is false. Consequently, ¬SUPPOSITION is necessarily true, which makes it 100% likely. But, at the same time, if $p \leq 1$, the probability of ¬SUPPOSITION is zero. Thus, owing to the inconsistency and concomitant impossibility of a necessary event's being both 0% likely and 100% likely, UNRELIABLE REAPER is impossible if $p \leq 1$. But since the probability of ¬SUPPOSITION needn't be (or is not) zero if $p > 1$, the UPD proponent should hold that UNRELIABLE REAPER remains possible if $p > 1$, due to its consistency in that case.

The thought is that ¬SUPPOSITION is necessary because SUPPOSITION cannot be true without contradiction, and contradictions cannot be true; this is the sense in which "we have to avoid paradox," to borrow Pruss's phrasing. Naturally we might think that whatever is necessary, i.e., whatever cannot fail to obtain, is 100% likely. For if it were less than 100% likely, there would be some chance, and thus some possibility, of its failing to obtain. But nothing necessary possibly fails to obtain.

That much is straightforward. But why does Pruss think ¬SUPPOSITION has a probability of zero in the event that $p \leq 1$? Here's why.

Recall that SUPPOSITION, put simply, is the proposition that there is no Reaper, r, such that all Reapers poised to attack Fred before r try and fail to kill him. We'll say those Reapers are "earlier than" r. So ¬SUPPOSITION, the negation of SUPPOSITION, is the proposition that all (the infinitely many) Reapers earlier than some Reaper non-fatally swing. To determine the probability of ¬SUPPOSITION, we need to consider the probability of the failure of each member of an infinite subseries of the series of Reapers.

While presenting his Unreliable Reaper scenario, Pruss tells us that "[t]he probability that the n^{th} Reaper's swing would succeed in killing Fred is $1/n^p$." It may appear that only one event is involved, but in fact two are. There's the event T_n of the n^{th} Reaper's *trying* to kill Fred by swinging its scythe at him, as well as the event S_n of the n^{th} Reaper's swing's killing Fred, i.e., *succeeding*. Thus, "[t]he probability that the n^{th} Reaper's swing would succeed in killing Fred" is ambiguous, considered in isolation. It may mean the unconditional probability of the conjunction $T_n \wedge S_n$, or it may mean the probability of S_n conditional on T_n. Thus, $1/n^p$ is either a) the chance that the n^{th} Reaper swings its scythe and thereby kills Fred or b) the chance, given that the n^{th} Reaper swings at Fred, that this attempt to kill Fred succeeds.

The latter is perhaps the most charitable reading of Pruss. Had he meant the former, we would run into contradictions of a kind Pruss altogether neglects to mention. If you're curious what contradictions those are and how they would undermine Pruss's critique of the UPD, see Appendix A.

There are countably infinitely many Reapers you might single out for consideration. No matter which Reaper you choose, infinitely many Reapers are poised to attack Fred before that one does. And each of those earlier Reapers has a $1/n^p$ chance of killing Fred, given that it tries to kill him. For this reason, no matter which Reaper, r, you choose, the probability is zero that all Reapers earlier than r try and fail to kill Fred, supposing all Reapers (even those later than r) try. Moreoever, as Pruss points out in a parenthetical remark, the union of *countably* many—even countably *infinitely* many—zero-probability events itself has probability zero. We have such a union here. For each of the countably infinitely many Reapers, there's a distinct hypothetical event of the failure of all Reapers earlier than that one, representable as the intersection, or conjunction, of all the individual failures. We can take the union of all those zero-probability events, which represents their disjunction. That union, having countably many events in it, must itself have probability zero, conditional on all the Reapers'

swinging. That is, there's zero chance that no Reaper kills Fred, assuming absolutely every Reaper in the series swings. Fred's being killed by no Reaper has a *conditional* probability of zero—the condition being that every single Reaper swings. Whether it has an unconditional probability of zero is another matter, to be addressed shortly. At any rate, if you're unconvinced of the probabilistic claims made thus far (and you probably shouldn't be convinced), see Appendix B.

There needn't be any tension between an event's being necessary and its having a *conditional* probability of zero. For even if every necessary event must have an unconditional probability of one, a necessary event may have a conditional probability of zero, relative to certain conditions. Any event E, be it necessary or contingent, has a probability of zero conditional on E's nonoccurrence, denoted by $\neg E$, or conditional on a conjunction (intersection) including $\neg E$ as a conjunct. Moreover, the event of, say, Pruss's posting the Unreliable Reaper thought experiment on his blog in April of 2025 (though contingent) has in some sense an unconditional probability of one now that it has occurred, but it has a probability of zero on the supposition that Pruss has yet to post Unreliable Reaper on his blog. The supposition is obviously false, but that doesn't mean we can't assign a probability based on that assumption. To do so is just to ask how likely it would be, in some hypothetical universe U (where September comes after April as it does in our universe), for Pruss to post the scenario in April of 2025 if he had yet to post it by September of the same year in U.

The general point of the previous paragraph seems to bear scrutiny, but it's doubtful that it applies to the case at hand. The event whose probability is in question is that, for some $j > i$, all Reapers assigned numbers greater than or equal to j (henceforth "Reapers j onward") swing non-fatally. Arguably, if this particular event has a probability of zero conditional on Reapers j onward's swinging for some $j > i$, it must also have an unconditional probability of zero. For this event partly consists in the condition that, for some $j > i$, Reapers j onward swing; that is to say, the event includes the condition relative to which we're assigning the event a probability. Thus, whatever it is about the condition that makes the event 0% likely, that is an intrinsic feature of this latter event, which makes it 0% likely independently of conditionalization.

There's another worry for Pruss's argument though. It's not actually necessary that there be some $j > i$ such that Reapers j onward swing non-fatally. Thus, there's no need to explain how there being such a j is necessary despite being 0% likely. Here's a way to see that this zero-probability event needn't be necessary. Note two distinctions: a) that between avoiding paradox *in a Grim*

Reaper scenario and avoiding paradox, and b) that between an event's being necessary and some condition's necessitating it. What is commonly believed is that, necessarily, no contradiction is true. That is the sense in which we ostensibly must avoid paradox; we don't need to *narrowly* avoid paradox, as we would were a situation like UNRELIABLE REAPER to arise without leading to some contradiction that could easily result. It's possible for no such situation to arise in the first place. If none did, we could avoid paradox despite the nonoccurrence of the zero-probability event that Reapers j onward swing non-fatally. We also know that, necessarily, *if* the UNRELIABLE REAPER story is true (and paradox avoidance is necessary), *then* there is some $j > i$ such that Reapers j onward fail. In other words, it's impossible for these two things to be true at once: 1) there's no such j, and 2) the UNRELIABLE REAPER scenario actually happens. For short we say (2) "necessitates" the negation of (1) (as does (1) the negation of (2)). It does *not* follow from this that it's *necessary simpliciter* that there be some $j > i$ such that Reapers j onward fail.

Why do I say there being such a j is *not* necessary, rather than suspending judgment on whether it is? I say so because there's a plethora of possible situations where no such j exists. Granted, if any of these possibilities is actualized, the UNRELIABLE REAPER story is false. But regardless of whether that story is true, these alternate scenarios are possible, and thus they may have nonzero probability.

One alternate possibility is that there are no Reapers. That's not only possible; it's reality! This is an *alternate* possibility because, if there are no Reapers, there are no Reapers j onward that fail for any $j > i$. Another possibility is that none of the Reapers swings. But if none swings, none swings non-fatally. So none fails, i.e., makes an unsuccessful attempt to kill Fred. And if none fails, there's no $j > i$ such that Reapers j onward fail. It's even conceivable that all the Reapers succeed, insofar as it's conceivable that Fred is resurrected each time a Reaper kills him.

Therefore, even if $p \leq 1$, no zero-probability event need be necessary. And if none is necessary, we may infer from any event's having probability zero that it does not have probability one (as it plausibly would were it necessary). So the UNRELIABLE REAPER story is not rendered impossible when we include in that story the supposition that $p \leq 1$. As long as there's no other inconsistent pair in UNRELIABLE REAPER, and I see no reason to think there is any other, those who accept the unsatisfiable pair diagnosis should be open to the idea that UNRELIABLE REAPER is possible even if $p \leq 1$.

Appendix A

On the modest assumption that $n = 1$ for some Reaper (that is, that there's a Reaper 1), it follows that, for that Reaper, $1/n^p = 1/1^p = 1$ for any p. But then there's a 100% chance that Reaper 1 kills Fred. Observe as well that there's a $1/2^p$ chance of Reaper 2's killing Fred, and that $1/2^p > 0$ for any positive real p. But as long as the probability of Reaper 1's killing Fred and that of Reaper 2's killing Fred are both nonzero and the two events are mutually exclusive, the probability that one or the other kills Fred is greater than the probability that Reaper 1 does, and greater than the probability that Reaper 2 does. We've seen that these two probabilities are indeed nonzero. The two events are also mutually exclusive, as only one Reaper can kill Fred; Fred can only die once. So the probability that Reaper 1 or Reaper 2 kills Fred is greater than the probability that Reaper 1 does. Ergo, the probability that Reaper 1 or Reaper 2 kills Fred is greater than 100%. But nothing can be over 100% likely. Therefore, that Reaper 1 or Reaper 2 kills Fred is, yet isn't, over 100% likely. There's our contradiction.

In fact, we don't even need to assume there is a Reaper 1, though that simplifies the math. Suppose $p = 1$. Then as long as there's a Reaper 2, we have that it's over 100% likely that Reaper 2, 3, or 4 kills Fred; to be exact, the probability is $1/2 + 1/3 + 1/4 = 6/12 + 4/12 + 3/12 = 13/12$. And if there's a Reaper 3, it's over 100% likely that Reaper 3, 4, 5, 6, or 7 kills Fred. And so on. Even changing p does not necessarily rescue us from absurdity.

It's problematic for Pruss if his scenario leads to contradiction in all these cases. His thesis is that the only available Inconsistent Pair response to Unreliable Reaper entails, counterintuitively, that Unreliable Reaper is impossible if $p \leq 1$ but possible if $p > 1$. But if my previous argument goes through, there's another response available to the UPD proponent: that Unreliable Reaper is *possible* if and only if

i. $p > 1$, and
ii. the Reapers are numbered $x_1, x_2, x_3, \ldots$, where $1/x_1^p + 1/x_2^p + 1/x_3^p + \cdots \leq 1$.

I doubt if this proposal would strike Pruss as counterintuitive in the same way. It doesn't have the prima facie odd consequence that the scenario is impossible whenever p is 1 but possible whenever p is 1.000001—or, more broadly, that a small numerical difference makes a big modal difference.

So on this reading of "[t]he probability that the n^{th} Reaper's swing would succeed in killing Fred is $1/n^p$," Unreliable Reaper fails to threaten the UPD. But I don't think Pruss would have made such a glaring oversight in

devising his scenario. In all likelihood, what Pruss means by "[t]he probability that the n^{th} Reaper's swing would succeed in killing Fred" is not the unconditional probability $P(T_n \wedge S_n)$ but rather the conditional probability $P(S_n|T_n)$. That is, he means the probability, conditional on the n^{th} Reaper's swinging at Fred, of the swing's killing Fred. And if that equals $1/n^p$, then the probability $P(\neg S_n|T_n)$ of the swing's *not* killing Fred, given that the swing occurs, is $1 - 1/n^p$.

Appendix B

If we've interpreted Pruss correctly, then, as noted at the end of the previous appendix, we have $P(S_n|T_n) = 1/n^p$ and, consequently, $P(\neg S_n|T_n) = 1 - 1/n^p$. Pruss also tells us what the probability is, given an integer $i \geq 0$, that every Reaper beginning with Reaper $i + 1$ fails to kill Fred, on the supposition that each Reaper in that subseries tries. He says it's zero. Let's see why that is.

We denote the event that none succeeds by $\bigcap_{n=i+1}^{\infty} \neg S_n$, which is a more compact version of the infinite sequence of intersections $\neg S_{i+1} \cap \neg S_{i+2} \cap \neg S_{i+3} \cap \ldots$ Likewise, we write $\bigcap_{n=i+1}^{\infty} T_n$ to denote the event that all Reapers beginning with Reaper $i + 1$ swing, i.e., try to kill Fred. In probability/measure theory any event representable as a conjunction p is the intersection of the events whose occurrence p's conjuncts assert. With the information provided in UNRELIABLE REAPER, we may deduce the conditional probability $P\left(\bigcap_{n=i+1}^{\infty} \neg S_n \middle| \bigcap_{n=i+1}^{\infty} T_n\right)$ that Fred is not killed by the time Reaper $i + 1$ would've swung—that all Reapers earlier than and including Reaper $i + 1$ fail (that is, Reaper $i + 1$ fails, Reaper $i + 2$ fails, Reaper $i + 3$ fails, and so on ad infinitum).

The probability

$$P\left(\bigcap_{n=i+1}^{\infty} \neg S_n \middle| \bigcap_{n=i+1}^{\infty} T_n\right)$$

equals

$$P\left(\neg S_{i+1} \middle| \bigcap_{n=i+1}^{\infty} T_n\right) P\left(\neg S_{i+2} \middle| \bigcap_{n=i+1}^{\infty} T_n\right) P\left(\neg S_{i+3} \middle| \bigcap_{n=i+1}^{\infty} T_n\right) \ldots$$

on the assumption that, for any positive integer m,

$$P\left(\neg S_{i+m} \middle| \bigcap_{n=i+1}^{\infty} T_n\right) = P\left(\neg S_{i+m} \middle| \bigcap_{n=i+1}^{\infty} T_n \cap \bigcap_{j<i+m} \neg S_j\right).$$

(I'm indebted to ChatGPT for the insight that this assumption is necessary.) This last equation states that the probability that Reaper $i + m$ does not succeed in killing Fred, given that all Reapers earlier than and including Reaper $i + 1$ try, equals the probability of the same given the same supposition *plus* the supposition that all Reapers later than Reaper $i + m$ do not succeed. Does this equation hold? Honestly, I doubt it. Conditionalizing on the additional event that no Reaper succeeds in killing Fred after Reaper $i + m$'s attempt arguably makes a difference to the probability, in that it gives us reason to suspect no Reaper even tries to kill Fred after Reaper $i + m$ swings. For a Reaper likelier doesn't succeed if it doesn't try at all than if it tries. And if no later Reaper tried, that would entail that Reaper $i + m$ succeeds (on our assumption that Reaper $i + m$ tries, which means all earlier Reapers try and fail); otherwise, some later Reaper would try to make up for Reaper $i + m$'s failure. By this reasoning, it would appear the additional condition gives us some evidence against $\neg S_{i+m}$, lowering its probability.

The elephant in the room is that, by specifying that we have in mind conditional probabilities rather than unconditional ones, we have made all the more plausible the suggestion that these probabilities are subjective, or epistemic. Indeed, I'm not sure there is any precedent in the philosophical or mathematical literature for defining or interpreting conditional probabilities as objective, like the chance of a coin's coming up heads supposedly is (as it's not merely a matter of how confident one actually is or should be that the coin will come up heads, but of something independent of one's mind and minds in general, like how frequently the coin would come up heads were it flipped infinitely many times). So, until someone goes to the trouble of elucidating what it might mean for these conditional probabilities to be objective, we had better stick with the conventional epistemic reading.

But suppose the foregoing equation is true, despite the worries I've raised. Then, to show $P\left(\bigcap_{n=i+1}^{\infty} \neg S_n \middle| \bigcap_{n=i+1}^{\infty} T_n\right) = 0$, we need the further assumption that

$$P\left(\neg S_{i+1} \middle| \bigcap_{n=i+1}^{\infty} T_n\right) P\left(\neg S_{i+2} \middle| \bigcap_{n=i+1}^{\infty} T_n\right) P\left(\neg S_{i+3} \middle| \bigcap_{n=i+1}^{\infty} T_n\right) \ldots$$

equals

$$P(\neg S_{i+1}|T_{i+1})P(\neg S_{i+2}|T_{i+2})P(\neg S_{i+3}|T_{i+3}) \ldots$$

or, in more compact notation, $\prod_{n=i+1}^{\infty} P\left(\neg S_n|T_n\right)$. Essentially, we need it to be the case that, for any positive integer m, once the occurrence of the $(i + m)^{\text{th}}$ Reaper's attempt/swing is fixed, the occurrence of any other Reaper's attempt

makes no difference to the probability of the $(i+m)^{\text{th}}$ Reaper's non-success. In other words, there must be no difference between $P(\neg S_{i+m}|T_{i+m})$ and $P\left(\neg S_{i+m}\middle|\bigcap_{n=i+1}^{\infty} T_n\right)$. But worries similar to the previous ones crop up here. Of course, something slightly different's going on; we're now comparing the probability that Reaper $i+m$ doesn't succeed assuming it tries to the probability of the same assuming all Reapers earlier than and including Reaper $i+1$ (including Reaper $i+m$) try. The extra assumptions that appear in the latter probability's condition have to do with attempts, whereas the extra assumptions we conditionalized on previously had to do with successes. Nonetheless, adding these assumptions seems to make a difference, supposing as we have been that these conditional probabilities are epistemic. In particular, we can be absolutely sure that Reaper $i+m$ fails to kill Fred if some later Reaper tries to kill him (and some later Reaper does given $\bigcap_{n=i+1}^{\infty} T_n$, provided that $m > 1$). But assuming merely that Reaper $i+m$ tries to kill him, we can only have a credence of $1 - 1/n^p$ that Reaper $i+m$ fails. Admittedly, this doesn't establish that the equation is false in cases where $m = 1$. But as long as it's false in every other case, we have a problem.

Anyway, bracket all those concerns. Just take the equations for granted. We continue as follows:

$$\begin{aligned}\prod_{n=i+1}^{\infty} P\left(\neg S_n|T_n\right) &= \prod_{n=i+1}^{\infty} \left(1 - P(S_n|T_n)\right) \\ &= \prod_{n=i+1}^{\infty} \left(1 - \frac{1}{n^p}\right)\end{aligned}$$

For those who don't know, I should mention that $\prod$ is like sigma, i.e., the summation notation $\sum$, only it represents a product, not a sum. To calculate $\prod_{n=i+1}^{\infty} (1 - 1/n^p)$, we multiply all the values of $(1 - 1/n^p)$ resulting from substituting for n all integers, one by one, in the interval $[i+1, \infty)$.

What is the value of this infinite product—this multiplication of infinitely many numbers? If we assume $i = 0 \wedge p = 1$, the answer is easy:

$$\begin{aligned}\prod_{n=i+1}^{\infty} \left(1 - \frac{1}{n^p}\right) &= \prod_{n=1}^{\infty} \left(1 - \frac{1}{n}\right) \\ &= \left(1 - \frac{1}{1}\right)\left(1 - \frac{1}{2}\right)\left(1 - \frac{1}{3}\right)\cdots \\ &= (0)\left(\frac{1}{2}\right)\left(\frac{2}{3}\right)\cdots = 0\end{aligned}$$

In fact it's trivial to calculate the product for any positive (or negative, for that matter) real value of p, so long as $i = 0$. In that case, for $n = 1$, we get $1/n^p = 1/1 = 1$, since one to any power is one. Thus:

$$\begin{aligned}\prod_{n=i+1}^{\infty}\left(1-\frac{1}{n^p}\right) &= \prod_{n=1}^{\infty}\left(1-\frac{1}{n^p}\right)\\ &= \left(1-\frac{1}{1^p}\right)\left(1-\frac{1}{2^p}\right)\left(1-\frac{1}{3^p}\right)\dots\\ &= (1-1)\left(1-\frac{1}{2^p}\right)\left(1-\frac{1}{3^p}\right)\dots\\ &= (0)\left(1-\frac{1}{2^p}\right)\left(1-\frac{1}{3^p}\right)\cdots = 0\end{aligned}$$

We end up with zero once again. But what happens if our nonnegative integer i is greater than zero? Then the smallest value of n, in any multiplier of the form $(1 - 1/n^p)$, is two or greater. Bearing this in mind, we can write a proof by cases that the infinite product is still zero provided that $p \leq 1$.

Suppose $p = 1$. In that case the result of $1 - 1/n^p$ is a positive fraction whose magnitude increases indefinitely as n approaches infinity. But even though it increases indefinitely, it does not become arbitrarily large. It is always less than one. Now, if you take any positive number, and multiply it by a positive fraction less than one, you get a smaller number than you had initially. So with each new multiplier m in the infinite product $\prod_{n=i+1}^{\infty}(1 - 1/n^p)$ (starting from the multiplier $1 - 1/(i + 1)^p$, and proceeding to $1 - 1/(i + 2)^p$ and so on), the product of all multipliers up to and including m shrinks. Call any such product of all multipliers up to and including a given multiplier a "partial product," as it is merely part of the infinite multiplication expression.

Not only does the partial product shrink indefinitely as multipliers are introduced, but its magnitude becomes arbitrarily small. (We'll consider an argument for this momentarily, albeit a half-assed one.) That is, it gets arbitrarily close to zero. Thus, the infinite product *is* zero, as it is by definition equal to the limit of the partial product as the number of multipliers goes to infinity.

How do we know the partial product gets arbitrarily close to zero as the number of multipliers goes to infinity? This follows from the fact that, for any $n \geq 2$, you can get arbitrarily close to zero by multiplying $(n - 1)/n$ by itself arbitrarily many times (or, more precisely, raising it to an arbitrarily large power). For multiplying $i/(i + 1)$ by itself arbitrarily many times is effectively what happens in the infinite product $\prod_{n=i+1}^{\infty}(1 - 1/n)$ (and we've stipulated

that $i \geq 1$, which makes
$i + 1 \geq 2$). We multiply $1 - 1/(i + 1)$ by $1 - 1/(i + 2)$, and then multiply the result by $1 - 1/(i + 3)$, and so on. Put another way, we multiply $i/(i + 1)$ by

$$\left(\frac{i+1}{i+2}\right)\left(\frac{i+2}{i+3}\right)\left(\frac{i+3}{i+4}\right)\ldots$$

Though $(i + 1)/(i + 2) > i/(i + 1)$, there can be no doubt that, if we multiply $(i + 1)/(i + 2)$ by $(i + 2)/(i + 3)$ (which is less than one), multiply their product by $(i + 3)/(i + 4)$ (again, less than one), and so on, then sooner or later we'll get a number no greater than $i/(i + 1)$. And in multiplying $i/(i + 1)$ by a bunch of numbers whose product is $i/(i + 1)$ or less, we in effect multiply $i/(i + 1)$ by itself or something smaller. Furthermore, given that this sequence of multipliers is infinite, we can be confident $i/(i + 1)$ is effectively squared infinitely many times. It just takes more multipliers every time, since the farther along you get in the sequence, the closer the multipliers get to one, and the less the partial product shrinks with the introduction of each multiplier.

I'm aware this argument doesn't really work, but it's the best I feel I can do at the moment. (Of course, if there's no legitimate proof to be given, that works in my favor, not Pruss's.)

Previously we assumed for the sake of argument that $p = 1$. We now imagine ourselves to have shown informally that, in that case, the infinite product is zero. Next, to complete our proof by cases, let's assume $p < 1$. (That p is less than one is the only possibility that remains, given our presupposition that $p \leq 1$.) Recall that Pruss stipulated $p > 0$. So the exponent p must be a positive proper fraction x/y. (In order for this fraction to be positive, it must be that x and y are both positive or both negative. But that's not pertinent.) I'm no mathematician, so this is all I managed to come up with on my own. At least it's semi-useful to the end of proving the desired conclusion. I had ChatGPT take care of the rest:

Prove that

$$\prod_{n=i+1}^{\infty}\left(1-\frac{1}{n^p}\right)=0$$

for $i \geq 1$ and $p<1$.

Thought for 4 seconds

I'll give a short, rigorous proof under the natural assumption $0<p<1$ (if $p \leq 0$ the factors $1-n^{-p}$ are not in $(0,1)$ for large n, so the statement needs reinterpretation).

Claim. For any integer $i \geq 1$ and any real p with $0<p<1$,

$$\prod_{n=i+1}^{\infty}\left(1-\frac{1}{n^p}\right)=0.$$

Proof.

1. **All factors are in $(0,1)$.**
   For $n \geq 2$ and $0<p<1$ we have $n^p>1$, so $0<1-n^{-p}<1$. (Since $i \geq 1$, every factor in the product has $n \geq i+1 \geq 2$.)
2. **The corresponding p-series diverges.**
   Consider the series $\sum_{n=i+1}^{\infty} \frac{1}{n^p}$. Because $0<p<1$, the integral

$$\int_{i+1}^{\infty} x^{-p}\,dx=\lim_{R\to\infty}\frac{x^{1-p}}{1-p}\Big|_{x=i+1}^{x=R}$$

$$\int_{i+1}^{\infty} x^{-p}\,dx=\lim_{R\to\infty}\frac{x^{1-p}}{1-p}\Big|_{x=i+1}^{x=R}$$

   diverges to $+\infty$ (since $1-p>0$ and $R^{1-p}\to\infty$). By the integral test (or a standard comparison), $\sum_{n=i+1}^{\infty} n^{-p}=\infty$.
3. **Take logarithms and compare.**
   The logarithm turns the product into a sum:

$$\ln\left(\prod_{n=i+1}^{\infty}\left(1-\frac{1}{n^p}\right)\right)=\sum_{n=i+1}^{\infty}\ln\left(1-\frac{1}{n^p}\right).$$

   For $0<x<1$ the inequality $\ln(1-x) \leq -x$ holds. Applying this with $x=n^{-p}$ gives

$$\ln\left(1-\frac{1}{n^p}\right) \leq -\frac{1}{n^p}.$$

   Summing over $n \geq i+1$,

$$\sum_{n=i+1}^{\infty}\ln\left(1-\frac{1}{n^p}\right) \leq -\sum_{n=i+1}^{\infty}\frac{1}{n^p}=-\infty.$$

4. **Conclude the product is zero.**
   The sum of the logarithms is $-\infty$, so the logarithm of the product is $-\infty$. Exponentiating yields

$$\prod_{n=i+1}^{\infty}\left(1-\frac{1}{n^p}\right)=\exp(-\infty)=0.$$

This completes the proof for $0<p<1$. (If $p=1$ the p-series diverges too and $\prod(1-1/n)=0$ by the same log-inequality argument; if $p \leq 0$ the factors are not of the form $1-x$ with $0<x<1$ for large n, so the product behaves differently and the statement must be adjusted.) ∎

What's neat is that, despite my forgetting to state my underlying assumption that $p > 0$, GPT picked up on the fact that such an assumption is appropriate in this context. This is perhaps a bit of evidence in favor of GPT's possessing true intelligence and my lacking it.

If the above proof works, we have, for all $i \geq 0$,

$$P\left(\bigcap_{n=i+1}^{\infty} \neg S_n \middle| \bigcap_{n=i+1}^{\infty} T_n\right) = 0,$$

but what we want is

$$P\left(\bigcup_{i=0}^{\infty} \bigcap_{n=i+1}^{\infty} \neg S_n \middle| \bigcap_{n=1}^{\infty} T_n\right) = 0.$$

The way to reach this conclusion, presumably, is to establish that, for all $i \geq 0$,

$$P\left(\bigcap_{n=i+1}^{\infty} \neg S_n \middle| \bigcap_{n=1}^{\infty} T_n\right) = 0,$$

and then appeal to the principle that a union of denumerably many events, each of whose probabilities conditional on C is zero, itself has probability zero given C (assuming, as ChatGPT said we must, that no individual event E in the union's probability changes when we add to C the events in the union which have smaller indices than E's). I leave that proof as an exercise for the reader.