Auszüge aus: **Martin Rees: Our Final Hour**, New York 2003
(Unterstreichungen: Marcel Hänggi. Halbfett: Kapiteltitel. In Klammern: Seitenzahlen)

Martin Rees ist Professor für Kosmologie und Astrophysik in Cambridge (Nachfolger auf Newtons Lehrstuhl), Astronomer Royal, war 2005 bis 2010 Präsident der Royal Society und ist seit 2012 Ehrendoktor der ETH Zürich.

(119)

Our "Final" Experiment?

Physicists aim to understand the particles that the world is made of, and the forces that govern those particles. They are eager to probe the most extreme energies, pressures, and temperatures; for this purpose they build huge and elaborate machines: particle accelerators. The optimum way to produce an intense concentration of energy is to accelerate atoms to enormous speeds, close to that of light, and crash them together. It is best of all to use very heavy atoms. A gold atom, for instance, has nearly two hundred times the mass of a hydrogen atom. Its nucleus contains 79 protons and 118 neutrons. A lead nucleus is heavier still, containing 82 protons and 125 neutrons. When two such atoms are crashed together, their constituent protons and neutrons implode to a density and pressure far higher than what they were when they were packed into a normal gold or lead nucleus. They may then break up into still smaller particles. According to theory, each proton and neutron consists of three quarks, so the resultant "splat" releases over a thousand quarks. The conditions of a particle accelerator replicate, in
(120)
microcosm, those that prevailed in the first microsecond after the "big bang," when all the matter in the universe was squeezed into a so-called quark-gluon plasma.

Some physicists raise the possibility that these experiments might do something far worse than smashing a few atoms, like destroying our Earth, or even our entire universe. Such an event is the theme of Greg Benford's novel COSM, where an experiment at Brookhaven laboratory devastates the accelerator and creates a new "microuniverse" (which remains, comfortingly, encased within a sphere small enough to be carried around by its graduate-student creator).

An experiment that generates an unprecedented concentration of energy could—conceivably, but highly implausibly— trigger three quite different disaster scenarios.

Perhaps a black hole could form, and then suck in everything around it. According to Einstein's theory of relativity, the energy needed to make even the smallest black hole would far exceed what these collisions could generate. Some new theories, however, invoke extra spatial dimensions beyond our usual three; a consequence would be to strengthen gravity's grip, rendering it less difficult than we previously thought for a small object to implode into a black hole. But the same theories suggest that these holes would still be innocuous, because they would erode away almost instantly, rather than tugging in more stuff from their surroundings.

The second frightening possibility is that the quarks might reassemble themselves into a very compressed object called a strangelet. That in itself would be harmless: the strangelet would still be far smaller than a single atom. However, the danger is that a strangelet could, by contagion, convert anything else it encountered into a strange new form of matter. In Kurt Vonnegut's novel Cafs Cradle a Pentagon scientist produces a new form of ice, "ice nine," that is solid at ordinary tempera-

(121)

tures; when it escapes from the laboratory it "infects" natural water, and even the oceans solidify. Likewise, a hypothetical strangelet disaster could transform the entire planet Earth into an inert hyperdense sphere about one hundred metres across.

The third risk from these collision experiments is still more exotic, and potentially the most disastrous of all: a catastrophe that engulfs space itself. Empty space—what physicists call "the vacuum"—is more than just nothingness. It is the arena for everything that happens: it has, latent in it, all the forces and particles that govern our physical world. Some physicists suspect that space can exist in different "phases," rather as water can exist in three forms: ice, liquid, and steam. Moreover, the present vacuum could be fragile and unstable. The analogy here is with water that is "supercooled." Water can cool below its normal freezing point if it is very pure and still; however, it takes only a small localised disturbance—for instance, a speck of dust falling into it—to trigger supercooled waters conversion into ice. Likewise, some have speculated that the concentrated energy created when particles crash together could trigger a "phase transition" that would rip the fabric of space itself. The boundary of the new-style vacuum would spread like an expanding bubble. In that bubble atoms could not exist: it would be "curtains" for us, for Earth, and indeed for the wider cosmos; eventually, the entire galaxy, and beyond, would be engulfed. And we would never see this disaster coming. The "bubble" of new vacuum advances as fast as light, and so no signal could forewarn us of our fate. This would be a cosmic calamity, not just a terrestrial one.

These scenarios may seem bizarre, but physicists discuss them with a straight face. The most favoured theories are reassuring: they imply that the risk is zero. But we cannot be one hundred percent sure what might actually happen. Physicists can dream up alternative theories (and even write down equa-

(122)

tions) that are consistent with everything we know, and therefore cannot be absolutely ruled out, but that would allow one or other of these catastrophes to happen. These alternative theories may not be frontrunners, but are they all so incredible that we needn't worry?

Back in 1983, physicists were already becoming interested in high-energy experiments of this kind. While visiting the Institute for Advanced Study in

Princeton, I discussed these issues with a Dutch colleague, Piet Hut, who was also visiting Princeton and subsequently became a professor there. (The academic style of this institute, where Freeman Dyson has long been a professor, encourages "out of the box" thinking and speculations.) Hut and I realised that one way of checking whether an experiment is safe would be to see whether nature has already done it for us. It turned out that collisions similar to those being planned by the 1983 experimenters were a common occurrence in the universe. The entire cosmos is pervaded by particles known as cosmic rays that hurtle through space at almost the speed of light; these particles routinely crash into other atomic nuclei in space, with even greater violence than could be achieved in any currently feasible experiment. Hut and I concluded that empty space cannot be so fragile that it can be ripped apart by anything that physicists could do in their accelerator experiments. If it were, then the universe would not have lasted long enough for us to be here at all. However, if these accelerators became a hundred times more powerful— something that financial constraints still preclude, but which may be affordable if clever new designs are developed—then these concerns would revive, unless in the meantime our understanding has advanced enough to allow us to make firmer and more reassuring predictions from theory alone.
The old fears resurfaced more recently when plans were announced, both at the Brookhaven National Laboratory in the US and at the CERN laboratory in Geneva, to crash atoms to-
(123)
gether even more forcefully than had been done before. The director of the Brookhaven Laboratory at the time, John Marburger (now President Bush's scientific advisor), asked a group of experts to look into the issue. They did a calculation along the lines of the one that Hut and I had given, and offered reassurance that there was no threat of a cosmic Doomsday triggered by tearing the fabric of space.
But these physicists could not be quite so reassuring about the risk from strangelets. Collisions with the same energy certainly occur in the cosmos, but under conditions that differ in relevant respects from those of the planned terrestrial experiments; these differences could alter the likelihood of a runaway process.
Most of the "natural" cosmic collisions happen in interstellar space, in an environment so rarefied that even if they produced a strangelet, it would be unlikely to encounter a third nucleus, so there would be no chance of a runaway process. Collisions with Earth also differ in an essential way from those in accelerators, because incoming nuclei are stopped in the atmosphere, which does not contain heavy atoms like lead and gold.
Some fast-moving nuclei, however, impact directly on the Moon's solid surface, which does contain such atoms. Such impacts have occurred over its entire history. The Moon is nonetheless still there, and the authors of the Brookhaven report proffered this indisputable fact as reassurance that the proposed experiment couldn't wipe us out But even these impacts differ in one possibly important way

from those that would occur in the Brookhaven accelerator. When a fast particle crashes onto the Moon's surface, it hits a nucleus that is almost at rest, and gives it a "kick" or recoil. The resultant strangelets, produced as debris in the collision, would share this recoil motion, and therefore be sent hurtling through the lunar material. In contrast, the accelerator experiments involve symmetrical collisions, where two particles approach each
(124)
other "head on." There is then no recoil: the strangelets have no net motion and therefore might stand more chance of grabbing ambient material.

Since the experiment would generate conditions that have never happened naturally, the only reassurance came from two theoretical arguments. First, even if strangelets could exist, theorists thought it unlikely that they would form in these violent collisions: it seemed more likely that the debris would disperse in the aftermath of the collision, rather than reassembling into a single lump. Second, if strangelets form, theorists would expect them to have a positive electric charge. On the other hand, to trigger runaway growth the strangelets would have to be negatively charged (so that they would attract, rather than repel, positively charged atomic nuclei in their surroundings).

The best theoretical guesses are therefore reassuring. Sheldon Glashow, a theorist, and Richard Wilson, an expert on energy and environmental issues, succinctly summarised the situation like this: "If strangelets exist (which is conceivable), and if they form reasonably stable lumps (which is unlikely), and if they are negatively charged (although the theory strongly favours positive charges), and if tiny strangelets can be created at the [Brookhaven] Retatdvistic Heavy Ion Collider (which is exceedingly unlikely), then there might just be a problem. A new-born strangelet could engulf atomic nuclei, growing relentlessly and ultimately consuming the entire Earth. The word 'unlikely,' however many times it is repeated, just isn't enough to assuage our fears of this total disaster."

What Risks Are Acceptable?

The accelerator experiments didn't *give* me any sleepless nights. Indeed, I don't know of any physicist who betrayed the
(125)
slightest anxiety about them. However, these attitudes are little more than subjective assessments, based on some knowledge of the relevant science. The theoretical arguments depend on probabilities rather than certainties, as Glashow and Wilson spell out clearly. There is no evidence that exactly the same conditions have ever occurred naturally. We cannot be absolutely sure that strangelets couldn't lead to a runaway disaster.

The Brookhaven report (and a parallel effort by scientists from the biggest European accelerator, CERN, in Geneva) were presented as reassuring. However, even if one accepted their reasoning completely, the level of confidence they offered hardly seemed enough. They estimated that if the experiment were run for ten years, the risk of a catastrophe was no more than one in fifty million. These might seem impressive odds: a chance of disaster smaller than the chance of winning the UK's national lottery with a single ticket, which is about one in fourteen million. However, if the downside is destruction of the world's population, and the benefit is only to "pure" science, this isn't good enough. The natural way to measure die gravity of a threat is to multiply its probability by the number of people at risk, to calculate the "expected number" of deaths. The entire world's population would be at risk, so the experts were telling us that the expected number of human deaths (in that technical sense of "expected") could be as high as 120 (the number obtained by taking the world's population to be six billion and dividing by fifty million).

Obviously, nobody would argue in favour of doing a physics experiment knowing that its "fallout" could kill up to 120 people. This is not, of course, quite what we were told in this case: we were told instead that there could be up to one chance in fifty million of killing six billion people . Is this prospect any more acceptable? Most of us would, I think, still be uneasy. We are more tolerant of risks that we expose ourselves to voluntar-

(126)

ily, or if we see some compensating benefit. Neither of these conditions pertains here (except for those physicists who are actually interested in what might be learnt from the experiment).

My Cambridge colleague Adrian Kent has emphasised a second factor: the finality and completeness of the extinction that this scenario would entail. It would deprive us of an expectation—important to most of us—that some biological or cultural legacy will survive our deaths; it would dash the hope that our lives and our works may be part of some continuing progress. It would, worse still, foreclose the existence of a (perhaps far larger) total number of people in all future generations. Wiping out all the world's people (and indeed destroying not just humans but the entire biosphere) could therefore be deemed far more than six billion times worse than the death of one person, So perhaps we should set an even more stringent threshold on the possible risk before sanctioning such experiments.

Philosophers have long debated how to balance the rights and interests of "possible people," who might have some future existence, against those of people who actually exist. For some, like Schopenhauer, the painless elimination of the world would not rate as an evil at all. But most would resonate more with Jonathan Schell's response: "While it is true that extinction cannot be felt by those whose fate it is—the unborn, who would stay unborn—the same cannot be said, of course, for extinction's alternative, survival. If we shut the unborn out of life, they will never have the chance to lament their fate, but if we let them into life they will have abundant opportunity to be glad that they were born instead of having been prenatally severed from existence by us. What we must desire first of all is that people be born, for their own sakes, and not for any other reason. Everything else —our wish to serve the future generations by preparing a decent world for them to

live in, and our wish to lead a decent life ourselves in a common world made
(127)
secure by the safety of the future generations—flows from this commitment. Life comes first, the rest is secondary."

Who Should Decide?

No decision to go ahead with an experiment with a conceivable "Doomsday downside" should be made unless the general public (or a representative group of them) is satisfied that the risk is below what they collectively regard as an acceptable threshold. The theorists in this episode seemed to have aimed to reassure the public about a concern that they considered unreasonable, rather than to make an objective analysis. The public is entitled to more safeguards than that. It isn't good enough to make a slapdash estimate of even the tiniest risk of destroying the world.
Francesco Calogero is one of the few who have addressed this issue thoughtfully. He is not only a physicist, but also a long-time activist for arms control, and a former general secretary of the Pugwash conferences. He expresses his concerns like this: "I am somewhat disturbed by what I perceive to be the lack of candour in discussing these matters... Many, indeed most [of those with whom I have had private discussion and exchanged messages] seem more concerned with the public relations impact of what they, or others, say or write, than in making sure that the facts are presented with complete scientific objectivity."
How should society guard against being unknowingly exposed to a not-quite-zero risk of an event with an almost infinite downside? Calogero suggests that no experiment that could conceivably carry such risks should be approved without a prior exercise, of a kind, familiar from risk analyses in other contexts, involving a "Red Team" of experts (which would not
(128)
include any of the group actually proposing the experiment) that would play devil's advocate, trying to think of the worst that might happen, and a "Blue Team" that would try to think of antidotes or counter-arguments.
When the purpose is to probe conditions where the physics is "extreme" and very poorly known, it is hard to rule out anything completely. Can we ever be sure enough of our reasoning to offer reassurance with the confidence level of a million, a billion, or even a trillion to one? Theoretical arguments can seldom offer adequate comfort at this level: they can never be firmer than the assumptions on which they are based, and only recklessly overconfident theorists would stake odds of a billion to one on the validity of their assumptions.
Even if a believable number could be assigned to the probability of a catastrophic outcome, the question remains: How low would the alleged risk have to be before we would give our informed consent to these experiments? There is no specific countervailing benefit to the rest of us, so the level would surely be lower than the experimenters might willingly accept on their own behalf. (It would also be far

lower than the risk of nuclear devastation that citizens might have accepted during the Cold War, based on their personal assessment of what was at stake.) Some would argue that one chance in fifty million was low enough, because that is below the chance that within the next year an asteroid large enough to cause global devastation will strike Earth. (This is like arguing that the extra carcinogenic effect of artificial radiation is acceptable if it does not do more than double the risk from natural radiation.) But even this limit doesn't seem stringent enough. We may become resigned to a natural risk (like asteroids or natural pollutants) that we cannot do much about, but that doesn't mean that we should acquiesce in an extra avoidable risk of the same magnitude. Indeed, efforts are made to reduce risks far below that

(129)

level whenever we can. That is why, for example, it is worth some effort to ameliorate the risk of asteroid impact.

UK government guidelines on radiation hazards deem it unacceptable that even the limited group of workers in a nuclear power station should risk more than one chance in one hundred thousand per year of dying through the effects of radiation exposure. If this very risk-averse criterion were applied to the accelerator experiment, taking the world's population as being at risk but accepting an equally stringent maximum number of deaths, we would require an assurance that the chance of catastrophe was below one in a thousand trillion. If equal weight were attached to the lives of all potential people who might ever exist—a philosophically controversial stance, of course—then it could even be argued that the tolerable risk was up to a million times lower still.

[2010 sprach ich Rees auf seine Passagen über Risiken des Cern an. Er reagierte sehr ungehalten: Sein Buch werde immer aus dem Zusammenhang gerissen zitiert; es gebe keinerlei Gefahr am Cern.]