

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository:<https://orca.cardiff.ac.uk/id/eprint/146739/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Shackel, Nicholas 2022. Uncertainty phobia and epistemic forbearance in a pandemic. Royal Institute of Philosophy Supplement 92 , pp. 271-291. 10.1017/S1358246122000248

Publishers page: <https://doi.org/10.1017/S1358246122000248>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies. See <http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



Uncertainty Phobia and Epistemic Forbearance in a Pandemic

Nicholas Shackel December 2021

Forthcoming in *Values and Virtues for a Challenging World*. Eds. Jefferson, A. *et al.*
Cambridge: Cambridge University Press.

In this chapter¹ I show how challenges to our ability to tame the uncertainty of a pandemic leaves us vulnerable to uncertainty phobia. This is because not all the uncertainty that matters can be tamed by our knowledge of the relevant probabilities, contrary to what many believe. We are vulnerable because unrelievable wild uncertainty is a hard burden to bear, especially so when we must act in the face of it.

The source of unrelievable wild uncertainty is that the nature of probability distributions matters for whether knowledge of them tames uncertainty. It matters because a warrant for the taming is provided by two theorems, but this warrant applies only to some kinds of probability distribution. Essentially, this is because the theorems are about what happens at a mathematical limit but real life never reaches the limit. Consequently, the warrant depends on how quickly the random processes producing the uncertainty converge towards their limit. If they are governed by one class of probability distributions, they converge quickly enough to possess the warrant. If they are governed by another class of probability distributions, they converge towards their limit too slowly and so do not possess that warrant. The random processes of pandemics involve the slow kind .

Faced with such a burden, as we are in a pandemic, we are tempted to retreat into uncertainty phobia, leading to fixed definite opinions, precisely when the exercise of sound judgement to determine our responses requires our opinions to be hedged and mobile. Coping with a pandemic requires us to bear the burden of unrelievable wild uncertainty rather than give in to the temptation of uncertainty phobia. Pandemics require the virtue of epistemic forbearance.

1. Confidence: certainty, uncertainty, evidence and stakes

When we act on the basis of a belief what we do depends on our confidence in the proposition believed.² If we are certain, we usually act without hesitation: if we are not, we take precautions. Our confidence in a proposition should vary with the balance of evidence. Strong evidence in favour leads to certainty in its truth, strong evidence against leads to certainty in its falsehood. As evidence weakens, certainty weakens into uncertainty. Sketching this as a graph, this variation in our confidence would look something like figure 1.

¹ My thanks to Jon Webber for very helpful comments and suggestions. This paper was written during my tenure of the 2021-22 Mind Association Major Research Fellowship, for which award I am very grateful.

² Here, a proposition is what is believed, or what is asserted by an assertion, rather than a proposal. For example, if I say or believe the bridge is safe, what I say or believe is the proposition that the bridge is safe.

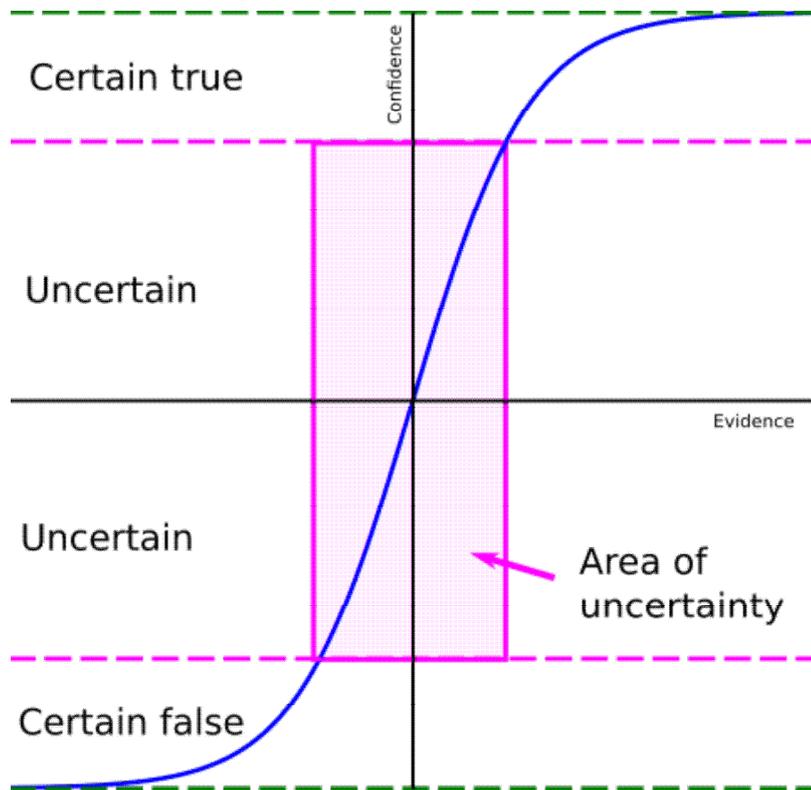


Figure 1 Mundane confidence. x-axis for the balance of the evidence for and against the proposition, i.e. balance positive is support for its truth, balance negative is support for its falsehood. y-axis for confidence: positive is confidence in truth, negative is confidence in falsehood.

Certainty has its own range, from fairly certain to absolutely certain, as does uncertainty. The boundaries are vague of course, and the box of uncertainty is just a rough indication of the extent of uncertainty's relation to evidence. But this is a useful model nonetheless. (There are no units on the axes here because these graphs are intended only to illustrate the features of how confidence varies with evidence.)

Bearing in mind that the confidence I am speaking of is the confidence to act on a belief, it is clear that confidence does not depend only on the evidence. In mundane cases, where the practical consequences of our belief being mistaken are not severe, it makes sense for our confidence to increase quite quickly with the evidence. For example, if the badness of being late is mild then, being fairly certain from memory that the bus departs at 5.10pm and so we will get it if we leave now is reasonable (even if in fact it departs at 5.05pm and so we miss it). We will call such examples low stakes cases. For high stakes cases we need stronger evidence to be certain and the relation of our confidence to evidence looks more like figure 2.

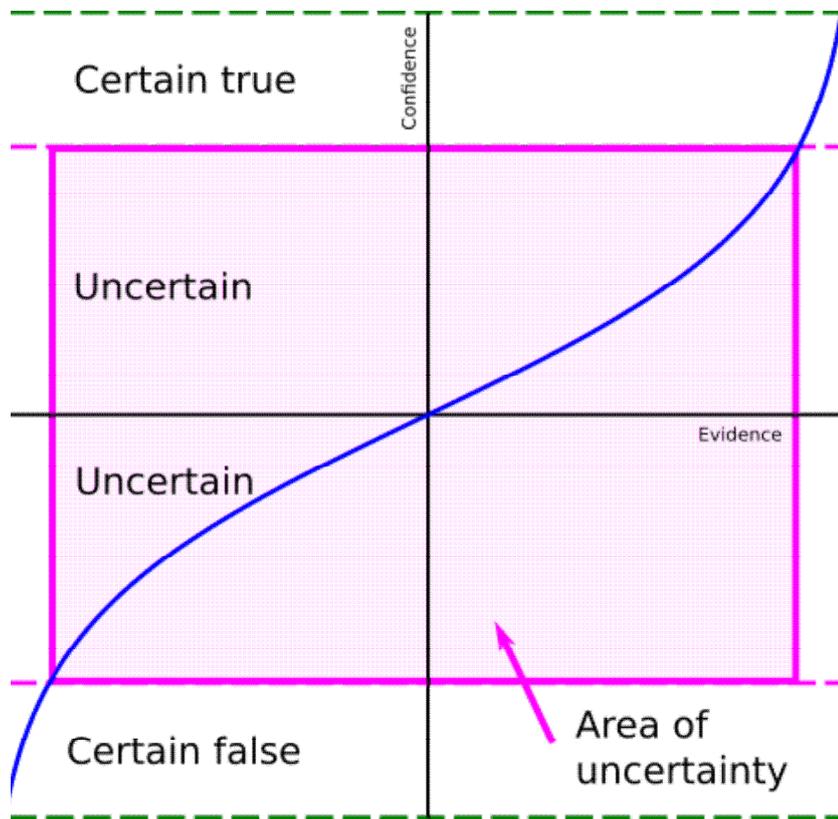


Figure 2 High stakes confidence.

Note how much greater the area of uncertainty is for high stakes cases. When our reliance on a belief has bad consequences if we are mistaken, then we should not be certain without stronger evidence. If missing the bus means missing the train which means missing the plane, for example, instead of relying on the evidence of my memory of the timetable, I should get stronger evidence by checking the various timetables.³

2. Taming uncertainty

The main tool we have developed to tame uncertainty is the mathematical theory of probability. Our hunger for certainty about how things are can be satiated instead by certainty in the probabilities of how they might be. We bear the uncertainty by being able to plan on probabilities. Confidence in probability reassures us and thereby suffices for acting despite uncertainty whilst giving good guidance for how to act with the precaution required by uncertainty.

When it comes to sophisticated applications to complex problems, knowing the relevant probability need not be as simple as knowing a specific numerical value. Probability functions map entire sets of possibilities onto probabilities and knowing such a probability function—knowing the probability distribution over those possibilities—greatly increases our power to tame uncertainty. There are many classes of probability distribution and each member of a class is picked out by few parameters. For example, each Normal distribution is distinguished from all the others by its mean and variance.

It is often as important to know the class of probability distribution as to know specific numerical probabilities. Even when one does not know the parameters that are needed to derive the numerical probabilities, knowing the class of distribution allows knowing various qualitative features of the uncertainty faced, features which may be as important

³ These kinds of rational influences of practical stakes are well known from the contextualism versus invariantism for knowledge debate. They are used to motivate contextualism and challenge invariantism, although here I am not taking a side in that debate (but see Shackel, 2011, MS-a).

as knowing the exact probabilities in planning what to do and responding to a developing situation.

Furthermore, knowing the class of distribution allows us to know how good an estimate of the relevant parameters the available data give. So we can know how good our estimate of the probabilities is despite a paucity of available data. For example, if we know that the distribution of car crashes is a Poisson distribution (Nicholson and Wong, 1993), we may be able to get good estimates, and know just how good those estimates are, for a particular road with little traffic and little crash data.

Consequently, the power of probability to tame uncertainty goes well beyond numerical probabilities. Taming uncertainty is multifaceted, from knowing probabilities through knowing means and variances, knowing probability distributions, to knowing how well calibrated the predictions on which we base our actions are likely to be.

3. Tolerable and dangerous risks

In general, the uncertainties that most interest us are unknown facts and events which may bring benefits and may bring harms. A risk is a possibility of a harm and taking a risk is doing something that may bring a harm. A *tolerable risk* is one for which the possible harm is predictably endurable. A *dangerous risk* is one for which the possible harm is unpredictably unendurable. An *intolerable risk*, of predictably unendurable harm, is almost always stupid to take (the exceptions, if there are any, being where they are accompanied by the possibility of enormous benefits), so we won't bother thinking about those here. What is endurable, and therefore what is tolerable or dangerous, depends to some extent on context and feasibility, of course. What is tolerable or dangerous will also vary with who is doing the choosing. For example, what may be a dangerous risk for an individual may be a tolerable risk for a government, just because societies can endure despite the loss of some of their constituent individuals.

How, then, do we know which risks are tolerable and which dangerous? Much of life confronts us with this question, since much of life consists in taking opportunities for the sake of their possible benefit and doing so despite their risks. A pandemic confronts us with this question relentlessly.

The issue turns on predictability. Where we have certainty, predictability follows. Where we have uncertainty, there we have difficulty. There are usually many different uncertainties that we need to consider. The harm we risk for each opportunity need not be singular but may extend over a range of bad things. If we drive to the beach, the possible harms run from negligible to devastating damage to the car and ourselves, yet together these form a tolerable risk. So the phrase 'predictably endurable' does not mean no possibility of a harm great enough to be unendurable. It means the range of harms that we expect are endurable. Driving to the beach is a tolerable risk because devastating damage to the car and ourselves is sufficiently unlikely to be outwith the range we expect. In general, what we want to be able to do is predict that range.

Here is where the taming power of probability may enter. If we know the probability of harms, then we can predict that range. We can replace our uncertainty over the harm with certainty over probabilities of harms. We can then examine the tolerability of the risk in various ways: for example, we can consider the ratio of chances of the range of endurable harms versus unendurable harms; if the harms themselves are quantifiable we can calculate the mean (the mathematical expectation), variance, skewness etc., and can give error bounds; and so on. In this way we can predict mathematically whether the harm is endurable or not.

Or at least, so it appears. The assumption lying behind this appearance is that good knowledge of probability tames uncertainty, by which I mean allows us to calculate probabilities that are close to true probabilities and to make reliable predictions.

Unfortunately, whether this assumption is true varies with the class of probability distributions involved. Here is not the place to rehearse my formal arguments on this point.⁴ Here is the place to simply report their results.

In short, uncertainty is born wild. Good knowledge of what I call tame probabilities suffices to tame uncertainty and good knowledge of what I call wild probabilities leaves uncertainty wild. Uncertainty involving wild probabilities can be entirely tameable or may instead take a very long time and a great deal of data to tame.

The borderline between tame and wild probabilities is usually the borderline between thin and fat-tailed probability distributions. What this means is illustrated in figure 3.

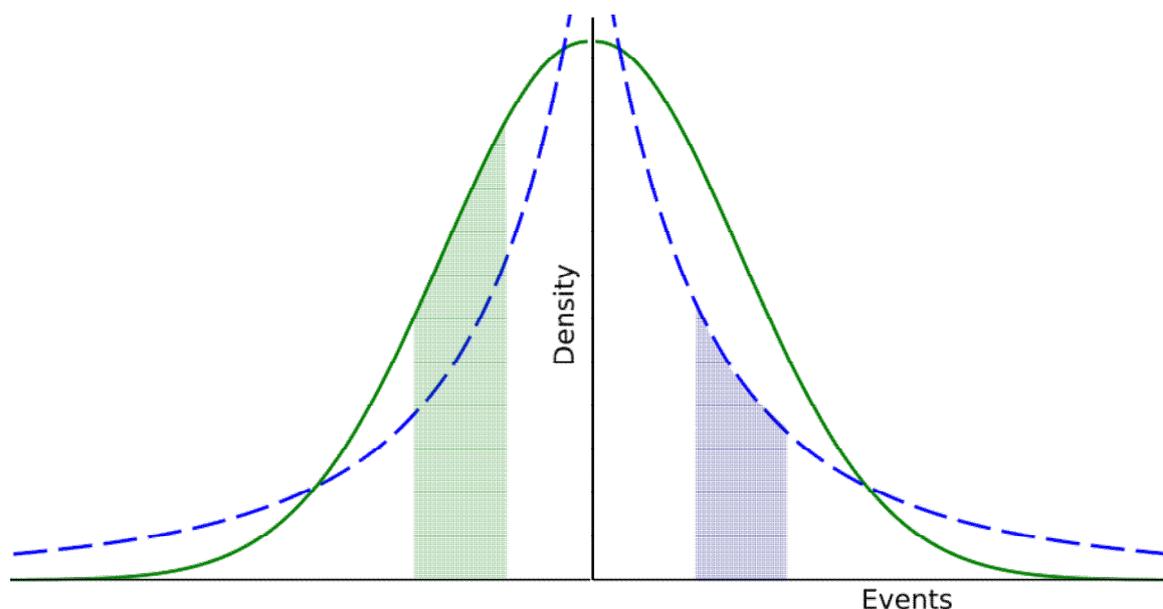


Figure 3 Probability density curves showing a thin and a fat-tailed distribution.

Here, the probability that an event is within a given range is the area under the curve (see shading for examples). The solid line distribution is called thin-tailed because its tails on either side get very close to zero quickly. The dashed line distribution is called fat-tailed because the tail takes a long time to get similarly close to zero. This means that extreme events are extremely unlikely for thin-tailed distributions but not for fat-tailed ones..

The distinguishing feature of a random process governed by a fat-tailed distribution is that events from that tail will dominate the cumulative result of repeated events from that random process, whereas for thin-tailed distributions it is events from the centre of the distribution that dominate. For example, if wealth has a fat-tailed distribution and we select at random a large enough sample of people, it is almost certain that nearly all of the total wealth is contributed by a single very wealthy person (e.g. a billionaire) than by a lot of wealthy people (e.g. millionaires). By contrast, height has a thin-tailed distribution and if we select at random a large sample of people, it is almost certain that nearly all of the total height is contributed by many people close to the mean rather than by a few extremely tall people.

If we have a process involving only thin-tailed probabilities, we can presume that good knowledge of the probability tames uncertainty. If the process involves fat-tailed probabilities, the uncertainty remains wild. These presumptions are not exceptionless, but

⁴ I provide these formal arguments in my forthcoming paper, 'Wild Uncertainty in a Pandemic' (Shackel, MS-b).

an exception must be proved, not assumed. The reason for this is to do with the warrant from a pair of theorems, which is explained in section 5.

4. Taming Uncertainty in the Covid-19 Pandemic

We use epidemic models to attempt to tame the uncertainty of a pandemic. They have been widely used during the Covid-19 pandemic (e.g. Anastassopoulou *et al*, 2020; Calafiore *et al*, 2020; Cooper *et al*, 2020; Ferguson *et al*, 2020; Ferrari *et al*, 2021; Martínez, 2021; Shringi *et al*, 2021). In this section and the next I will give a very much simplified description of why, at the beginning of a pandemic and for some considerable time, epidemic models leave uncertainty untamed.⁵

The fundamental issue is that both in estimating the parameters of epidemic models and in then using those models to make predictions, we rely on two theorems, the central limit theorem and the law of large numbers. The power of these theorems is that they apply to any underlying probability distribution,⁶ whether thin or fat-tailed, and they even apply when we don't know what those underlying probabilities distributions are. It is on this power that we rely for a warrant that the epidemic models tame the uncertainty.

The central limit theorem assures us that the features of our observed samples are close to the features of the whole population, and thereby give us good inputs for our model. The law of large numbers tells us that the future observed outcomes will be close to our predictions. These theorems are therefore essential to justifying any claim that the probabilities of future eventualities determined by an epidemic model (such as the probabilities of numbers of hospitalizations next month) will be close to the truth. Without these theorems, we have no good reason to believe that these predictions will be reliable.

Consequently, these two theorems are needed for our epidemic models to tame uncertainty. Let us call their warrant to this effect *the warrant from theory*.⁷

Our attempts to tame the uncertainty of a pandemic involve complex concatenations of models. Estimates of pandemic parameters can themselves rely on models and the use of those models relies on those estimates, and this mutual feedback may involve a number of different epidemic models. There are also many other kinds of models involved in coping with a pandemic, which concern the effect on our myriad normal activities when large numbers of people are falling sick. For example, health services use models to plan priorities, beds, staffing and supplies. These other models take as inputs the outputs of the various epidemic models, thereby adding further concatenations into the mix. The extent of such concatenations of models makes it very hard to disentangle and trace the routes of reliance on our two theorems. We have no audits tracing where we rely on the warrant from theory and where we do not. In the absence of such audits, our use of epidemic models to attempt to tame the uncertainty of a pandemic depends quite generally on the warrant from theory.

5. Untameable uncertainty in the Covid-19 Pandemic

Unfortunately the warrant from theory, whilst satisfactory for thin-tailed random processes, can fail for fat-tailed ones. The source of the problem is that, technically, the two theorems are about what happens *at the limit*, i.e. as the number of observations tends

⁵ For the explanations of what I here can only report, see Shackel, MS-b

⁶ For the former theorem to apply, a probability distribution must have both mean and variance, for the latter it must have a mean.

⁷ There are other warrants on which we rely as well, of course, such as the empirical science justifying the epidemic models. Their failure would pose a different problem from my concern here, since absent such warranted models, we would lack good knowledge of the probabilities and so we would know that we lack the basic tool to attempt to tame the uncertainty. So for our purposes we can assume all those warrants are in place.

to infinity. In real life, it can matter how quickly the accumulation of outcomes converges to the limit, i.e. how many observations you need to be within a specified distance of the limit. Thin-tailed random processes converge quickly, fat-tailed random processes converge slowly, sometimes very slowly indeed. This makes a huge difference to the sample size needed for an estimate of a parameter to be within a specified distance of its true value. For example, it can take only 30 observations of events governed by a Normal distribution (which is always thin-tailed) for our estimate to be highly likely to be that close. To get that close with a Pareto distribution⁸ (which is always fat-tailed) it takes 100,000,000,000 (one hundred thousand million) observations (see Taleb, 2020, p.40).

So where a random process involves fat-tailed probabilities, even if we know which class of fat-tailed probability distributions are involved, the warrant from theory fails because convergence is too slow. It can be so slow that we simply cannot get a large enough sample for the central limit theorem or the law of large numbers to apply. Even where over time we may accumulate large enough samples for them to apply, it may still take a very long time before we have achieved that accumulation needed for the warrant to be in place.

Here, then, is why the presumption must be that thin-tailed distributions are tame probabilities and fat-tailed distributions are wild probabilities. In general, and provably so, thin-tailed distributions converge quickly enough to have the warrant from theory whilst fat-tailed distributions do not converge quickly enough and do not have the warrant. To take an exception, therefore, requires fulfilling a burden of proof. For example, needing to rely on a specific fat-tailed distribution being an exception to this general rule requires being able to show that its convergence to the limit is sufficiently atypical for the class of distributions to which it belongs. Sufficiency here is determined by the purposes for which one needs to rely on it.⁹

Pandemic random processes involve a number of crucial fat-tailed distributions. The number of deaths caused by a pandemic is one good example. An analysis of 72 pandemics from the past 2500 years, with estimated deaths for each one scaled to the current world population, shows that the cumulative effect of these pandemics is dominated by a few extreme events: the Plague of Justinian (541–549 AD) and the Black Death (1346-1353 AD) each claimed the equivalent of more than two billion lives; the five next largest pandemics each claimed the equivalent of more than 100 million lives (Cirillo and Taleb, 2020, pp. 608-9). The characteristic of a fat-tailed distribution is precisely that the cumulative effect is dominated by a few extreme events from the distribution's fat tail. This analysis shows that deaths from pandemics form an extremely fat-tailed probability distribution (Cirillo and Taleb, 2020, p. 606).

Another good example is a critical input to epidemic models. The reproduction rate (also called the reproduction number) is the number of secondary infections per infected person, i.e. how many people an infectious person infects during the time they are ill. The reproduction rate is the mean of the random variable, P , which for each infected person takes the value of the number of people that person infects.

Someone is called a 'superspreader' if their value for P is high, although different authors give different thresholds. Recalling that fat-tails mean single events produce most of the cumulative effect, the existence of superspreaders is itself an indicator of a fat-tailed distribution. We have very good evidence that the transmission of Covid-19 has

⁸ This particular Pareto distribution generates outcomes fitting the Pareto 80/20 principle: see below.

⁹ For example, the class of log-normal distributions is at the borderline. Those with low variance can be thin-tailed but those with high variance are fat-tailed. We might be in the position to prove which we had, which might then allow us to treat it as an exception.

been heavily determined by superspreading. According to one study, 60–75% of cases infect nobody, while 10–20% of cases cause 80% of all secondary infections (Chen *et al*, 2021; see also Hasan *et al*, 2020; Lau *et al*, 2020; Sun *et al*, 2021).

This estimate fits the Pareto 80/20 principle, named after the economist who first drew it to our attention (Pareto, 1896). The class of fat-tailed probability distributions called Pareto are so-called precisely because random processes with such distributions exhibit this pattern. When we see a random process exhibiting such a pattern, the underlying random variables are probably fat-tailed. Hence, the underlying random variable P for Covid-19 probably has a fat-tailed distribution. Recent research supports this conclusion. Wong and Collins have shown how where there is superspreading, the probability distribution of P is fat-tailed, and conclude that

combine[d] empirical observations of SARS-CoV and SARS-CoV-2 transmission and extreme value statistics...show that the distribution of secondary cases is consistent with being fat-tailed, implying that large superspreading events are extremal, yet probable, occurrences. (2020, p.29416)

Estimating the reproduction rate is, then, a matter of estimating the mean of a fat-tailed random variable. This is very hard to do well. We have already seen that it can require an enormous sample for our estimate to be as accurate a sample size of 30 would give for the mean of a thin-tailed distribution. The second problem is that we will systematically underestimate the mean because the overwhelming majority of samples from systematic random sampling will not include those rare superspreaders at all.¹⁰

A third good example of a fat-tailed distribution involved in the spread of Covid-19 concerns the network of human acquaintances. The human world consists of many clusters of mutually acquainted people, some of whom are acquainted with people in other clusters. Pandemics can spread easily within a cluster, any member of which can then spread the infection to another cluster by infecting an acquaintance of theirs in that other cluster. Someone who is very well connected, which is to say, has many acquaintances, can therefore spread the infection to many clusters.

The number of acquaintances of a person is a random variable. If that random variable is fat-tailed the network will be what is called a small-world network, which in this case would mean that although most humans are not acquaintances, a path from one person to another consisting entirely of acquaintances has, on average, a fairly small number of people in it. A small enough mean path length implies that the random variable is fat-tailed.

We know that the network of human acquaintances is a small-world network (Milgram, 1967; Collins and Chow, 1998) with a remarkably short mean path length. Milgram's originating experiment (Milgram, 1967) found 5 or less steps of acquaintance sufficed within the US. and Watts and Strogatz estimated the global mean path length to be 6 (1998). Consequently this evidence shows human acquaintance is probably governed by a fat-tailed random variable.

So the paths of pandemic transmission are significantly constituted by the network of human acquaintances. This is why the more complex stochastic epidemic models attempt to include the random processes of human acquaintance, movement and meetings that ground the causal transmission of a pandemic. The network is a small world with a fat-tailed random variable governing the number of acquaintances. Watts and Strogatz have shown that

¹⁰ For the technical details of this, see Shackel, MS-b.

infectious diseases are predicted to spread much more easily and quickly in a small world. (1998, p.442)

Hence there are network features grounding the causal transmission of pandemics that have fat-tailed distributions and these are parameters for the more complex stochastic models.

Our epidemic models of Covid-19 thus depend on inputs and parameters that are derived from random processes governed by fat-tailed probability distributions. The probabilities are therefore presumptively wild. In the absence of proof that they are exceptions to the general rule of fat-tailed distributions, we would require huge samples to overcome their slow convergence to the limit before our estimates attained the needed accuracy for reliable model outputs. Since we do not have such enormous samples, the warrant from theory fails. The wild uncertainty of a pandemic remains.

Could we ever, given enough time, accumulate enough data to finally tame this uncertainty? The answer to this question is unclear. It may depend on the specific fat-tailed probabilities involved. Some fat-tailed distributions have no variance and even no mean and for them no warrant from theory is ever directly available.¹¹ Any uncertainty governed by such distributions may be completely untameable!

Our need for the warrant from theory is strongest at the beginning of a pandemic, because then our data is most limited. It continues strongly for some considerable time. First, and unavoidably, because of the slowness of convergence of any fat-tailed random processes involved and the consequent inaccuracy of estimation. Second, as the pathogen itself evolves, the parameters of the pandemic random process that we are trying to estimate may not be stationary phenomena. Similarly, our responses that attempt to constrain the pandemic change the environment within which it operates and this can produce a moving target for our estimates. All such non-stationarity weakens the extent to which the accumulation of data over time strengthens the quality of our estimates.

We need to be aware here of hindsight bias. Eventually, we will be able to fit our models to many years' worth of data and at that point our models will seem able to predict the course of that historical pandemic. But this does not mean that we could have built these models and tamed our uncertainty *when we needed* to during that pandemic.¹²

In conclusion, then, Covid-19, in common with pandemics in general, has exhibited fat-tailed random processes and we have no audit tracing the routes for the warrant from theory, so we have no audit of its failure. In the absence of that audit, the failure is quite general. So initially and for some considerable time, the Covid-19 pandemic faced us with untamed uncertainty.

6. Taming failures during the Covid-19 pandemic

At first slowly, and then quickly, we became aware of the dangerous risks of Covid-19. We saw repeated reversals in government policy as governments attempted to 'follow the science'. When we looked at the science, we saw persistent wide divergences in estimates and predictions coming from the models of different teams of experts and persistent failures of model predictions.

For example, in August 2020 Ioannidis et al pointed out that on 27 March 2020 'brilliant scientists expected 100,000,000 cases accruing within 4 weeks in the USA' (Ioannidis *et al*, 2020, p.1). By contrast, the US Centre for Disease Control reports total

¹¹ I say not directly because there are some technical things we can do when there is an upper bound on a random variable which may ameliorate (but cannot get around) this problem.

¹² An illustration of the difficulty here is from the research on the stock market showing models well fitted to even many past years quickly fail to predict the future, and when a model is then built to countenance the new failures, it too fails for its future. (Fama, 1970; Summers, 1986; Bernstein, 1992; Mandelbrot and Hudson, 2004)

cumulative positive specimens from 1st March to 25th April 2020 as 702,814 (United States Centre for Disease Control, 2020, p.3). Of course, there were more infections than positive specimens, but this fact cannot account for such an enormous disparity.

Similarly, the team at Imperial College London predicted in March 2020 that there would be roughly 1,500,000 deaths across the UK and the USA by the middle of June 2020 (Ferguson *et al*, 2020).¹³ In fact the number of deaths across the whole world by that date was less than one-third of that number ((Ioannidis *et al* 2020, p.1). Ioannidis *et al* found instances of the same model going wrong in both directions: ‘the model initially over predicted enormously, and then it under predicted’ (2020, p.4). Additionally,

even for short-term forecasting when the epidemic wave waned, models presented confusingly diverse predictions with huge uncertainty. (2020, p.3)

Ioannidis *et al* conclude that

despite involving many excellent modellers, best intentions, and highly sophisticated tools, forecasting efforts have largely failed (Ioannidis *et al*, 2020, p.1)

Subsequent research has confirmed this impression of unreliability. James *et al* note ‘A proliferation of models, often diverging widely in their projections’ (2021, p.379). Poor calibration has been observed. Poor calibration means (in this case) that predictions diverge from actuality and do so with frequency lying outside predicted probabilistic error bounds. For example, predictions of 1000 deaths per day with a 95% confidence interval of 100 would be poorly calibrated if actual deaths were outside the range 900-1100 more than 5% of the time. Gnanvi *et al* (2021) report the poor calibration of the Covid-19 pandemic models they studied. 25% were outside their 95% confidence intervals for numbers of cases when only 5% should have been. Endo *et al* conclude that ‘calibration of the [mortality risk] models were poor’ (Endo *et al*, 2021, p.1). See also Eker, 2020; Holmdahl and Buckee, 2020 and for a general overview see Cepelewicz, 2021.

Evidently, such failures of prediction mean that calculations of probabilities and predictions using these models did not tame the uncertainty of Covid-19, at least initially and for some considerable time. The epidemic models used were well grounded in current methodologies and literature and, whilst no one would claim they are flawless, their warrant is commensurate with scientific warrants in medicine in general. The failure of those models therefore shows that the conclusion of section 5, that the Covid-19 pandemic faced us with untamed uncertainty, was manifest during the first six months of the Covid-19 pandemic.¹⁴

7. Uncertainty Phobia

Uncertainty is burdensome and it is especially burdensome in high stakes cases, when distinguishing tolerable and dangerous risks is very important. It is partly its burden that drives us to eliminate it by seeking better evidence. Yet sometimes no better evidence is available and still we must act. The burden of such unrelievable uncertainty can drive us instead to an irrationality I shall call *uncertainty phobia*. Instead of bearing the uncertainty we may end up responding like figure 4:

¹³ It is tricky to extract exactly the cumulative predicted figure from that Imperial College report (Ferguson *et al*, 2020), but mid-June is about the peak daily deaths in those predictions that were for total 500,000 UK deaths and 2,200,000 US deaths in that wave (assuming no mitigation). Since mid-June is halfway through that wave, this means predicted deaths of roughly 1,500,000 between the UK and the US by that point.

¹⁴ For avoidance of doubt, I am not arguing that epidemic models have no use in managing a pandemic. Rather, when they leave uncertainty wild, we need to think harder about just what help they can yet provide. For more on this, see Shackel MS.

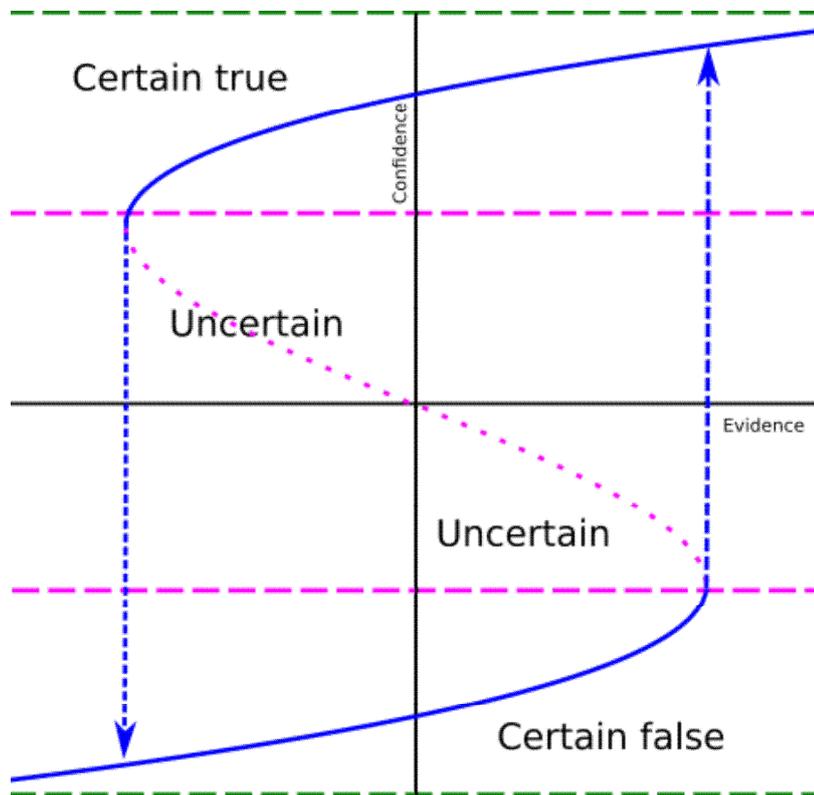


Figure 4 *Uncertainty phobic confidence.*

The dotted diagonal of the S-shape may be steep or shallow. What is critical is that this fold in our response removes uncertainty by making the range of uncertain confidences unavailable. If the evidence changes against our belief, we may continue to be certain even when the evidence turns negative, until it reaches a point at which we suddenly switch from being certain true to certain false (and vice versa). Consequently, there is no area of uncertainty because our confidence can never take a place on the dotted diagonal.

If we fall into uncertainty phobia over a question, we will be certain even when we should not be. Which way we will be certain will be path dependent. Whichever side of a question we were on as the stakes rose, and however weakly we were on that side, we will end up certain on that side.

This also means that uncertainty phobia will amplify the effects of even very small irrationalities already in place. I doubt we have the ability to avoid small irrationalities, such as neglecting entirely some weak but uncongenial evidence. Were such an irrationality to place us very slightly on one side of a question when we should have been slightly on the other, that small irrationality determines where we end up fixed.

The upshot here is that anyone driven into uncertainty phobia on any questions will be stubbornly certain, unreasonably so, on those questions. How exactly we would sustain and rationalise our stubborn certainties is an open question. There are various well-known kinds of cognitive error that would suffice. One well known example is confirmation bias: our tendency to notice and remember evidence that seems to confirm the belief we already hold at the expense of evidence that counts against it. But really, the problem is far worse than that, since in uncertainty phobia, even whilst we may be taking in the balance of the evidence accurately, we have almost entirely lost our graded sensitivity to it.

What has driven us here is that the height of the stakes makes distinguishing tolerable and dangerous risks very important and the unrelievable uncertainty has removed our ability to do that rationally. As a result, we have evaded the zone of uncertainty altogether.

These practicalities are in the driving seat and cognition is, for the time being, their passenger.

There may be times when this is perfectly rational. For example, imagine that you must leap over a yawning chasm to save your life, and we know that you are more likely to succeed if you act with certainty that you will succeed. In such a case, it may be practically wise to become (however temporarily) uncertainty phobic so that you can acquire the certainty you need to succeed, despite this being theoretically irrational. That we have this ability may even be an evolved tendency.

That being said, a lot of the time uncertainty phobia is badly irrational. When we, both individually and together, are addressing sustained wavering evidence bearing on prolonged high stakes activities requiring good judgement to distinguish tolerable and dangerous risks, uncertainty phobia is going to be unwise just because it is theoretically irrational. In such cases, I think we can safely say that uncertainty phobia is an epistemic vice.

8. Polarisation in Response to Covid-19 Untameable Uncertainty

That our response to high-stakes unrelievable uncertainty can be uncertainty phobic makes an empirical prediction: when stakes are high and when the evidence is wavering around for a long time whilst action must be taken, we will see a polarisation of opinion on many of the relevant factual claims and consequent polarisation on what should be done. Such a polarisation is simply the upshot of widespread uncertainty phobia. Many people are stubbornly certain on all these questions and the polarisation between their individual certainties is the amplification of even the slightest original disagreement.

The conclusion of section 5, in short, is that pandemics in general and Covid-19 in particular face us with sustained, unrelievable, wild uncertainty. Of course, for any particular one of us, that the uncertainty is unrelieved need not depend on the fact that the uncertainty is wild. For example, in a case where tame probabilities tame the uncertainty, we may happen not to know of the taming. We may, however, learn of the taming and may see its manifestation in the management of a situation, and in that way have our uncertainty relieved. The problem in a pandemic is that (at least initially and for some considerable time) the uncertainty is not simply unrelieved but is unrelievable just because it is wild. In that case, any claim that the models are taming the pandemic uncertainty is not true. In fact, it is propaganda. The failure to tame will become evident soon enough and that failure itself may drive many people to uncertainty phobia when they realise they have been deceived.

So the prediction made by uncertainty phobia applies to the Covid-19 pandemic. It seems to me that the prediction is satisfied to a significant degree by what we have seen during the pandemic. We have lived through prolonged unrelievable uncertainty and on many questions of fact and policy opinions have become polarised. This suggests that a lot of us, perhaps all of us, have some tendency to uncertainty phobia. To what extent we might be uncertainty phobic is, of course, a question to be investigated by empirical scientists rather than by philosophers like myself.

9. The Virtue of Epistemic Forbearance

Thus, at least at the beginning and for some considerable time, a pandemic confronts us with wild uncertainty. Consequently, we have to give up the belief that our models can generally tame the uncertainty. We must accept that pandemics face us with unpredictably unendurable harms. The risks we face are dangerous, not tolerable.

It may be that eventually we accumulate enough data and we know enough about the specific random processes of a pandemic that our models begin to give us sufficiently

accurate outputs to tame the uncertainty. And that would be a good thing. But we must be careful not to jump the gun. We will be tempted to do exactly that.

The harms of a pandemic are very high stakes and their being unpredictably unendurable is very frightening. In such a situation uncertainty phobia is a standing temptation. Uncertainty phobia predicts widespread polarization of opinion and we have seen just such a polarization of opinion on almost all the questions of what we should be doing about Covid-19. Whether it is on masks, travel restrictions, border closings, or vaccinations, the prevailing opinions are fixed certainties on either side of each question. So we certainly can fall into this temptation.

Uncertainty phobia is not always a bad thing. But facing a pandemic is not like leaping over a yawning chasm. Similarly, polarization of opinion is not always a bad thing. There are moral questions on which there are quite rightly sharp disagreements (even though the fact of sharp disagreement should give us some pause). Yet this is not the situation with Covid-19. There are disagreements about the priorities of the various values bearing on what we should do, but not such as to justify the polarisation we have seen.

We are instead faced with needing individual and social responses, responses which must properly countenance the varying untamed uncertainties pandemics pose. Because the uncertainties cannot be tamed, these responses cannot be calculated. They must rely on sound, sober and considered judgement. Far from our opinions being fixed certainties, they should instead be mobile uncertainties.

So, the nature of the uncertainty of pandemics, being wild and unrelievable for a long time, threatens our ability to sustain what rationality requires, namely, uncertain belief. Initially, and for a considerable time, we must simply avoid the temptation of uncertainty phobia and bear the wild uncertainty that faces us. I am not saying that is easy to do. Nevertheless, it is what we must do.

This requires virtue. We need *honesty* with ourselves, to keep in mind the balance of evidence (as best as we can ascertain it). We need *discrimination*, to distinguish the temptation of uncertainty phobia from a justified choice. Our hunger for certainty is justified by the risk being dangerous, but we must distinguish that hunger from the proper motive for factual opinion and see that it is a temptation to certainty rather than a justification for certainty. We can identify that certainty would be a vice here by distinguishing the balance of evidence we have from the balance of evidence needed for certainty given the stakes we face. This in turn requires *epistemic sensitivity* to the degree of confidence warranted by the evidence given the stakes. Finally, so long as the wild uncertainty is unrelieved, so long must we sustain this honesty, discrimination and sensitivity. The virtue that avoids uncertainty phobia by combining and sustaining these virtues is a kind of patience and fortitude that we may call *epistemic forbearance*.

10. References

- Cleo Anastassopoulou *et al* 2020. Data-Based Analysis, Modelling and Forecasting of the Covid-19 Outbreak. *PLOS ONE*, 15 (3), pp. e0230405.
<https://doi.org/10.1371/journal.pone.0230405>
- Peter L. Bernstein 1992. *Capital Ideas : The Improbable Origins of Modern Wall Street*. New York: Free Press : Maxwell Macmillan Canada.
- Giuseppe C. Calafiore *et al* 2020. A Time-Varying SIRD Model for the Covid-19 Contagion in Italy. *Annu Rev Control*, 50 pp. 361–372.
<https://doi.org/10.1016/j.arcontrol.2020.10.005>
- Jordana Cepelewicz 2021. The Hard Lessons of Modeling the Coronavirus Pandemic. *Quanta*,

- Paul Z. Chen *et al* 2021. Understanding Why Superspreading Drives the Covid-19 Pandemic but Not the H1n1 Pandemic. *The Lancet Infectious Diseases*, 21 (9), pp. 1203–1204. [https://doi.org/10.1016/S1473-3099\(21\)00406-0](https://doi.org/10.1016/S1473-3099(21)00406-0)
- Pasquale Cirillo & Nassim Nicholas Taleb 2020. Tail Risk of Contagious Diseases. *Nature Physics*, 16 (6), pp. 606–613. <https://doi.org/10.1038/s41567-020-0921-x>
- James J. Collins & Carson C. Chow 1998. It's a Small World. *Nature*, 393 (6684), pp. 409–410. <https://doi.org/10.1038/30835>
- Ian Cooper *et al* 2020. A Sir Model Assumption for the Spread of Covid-19 in Different Communities. *Chaos, solitons, and fractals*, 139 pp. 110057. <https://doi.org/10.1016/j.chaos.2020.110057>
- Sibel Eker 2020. Validity and Usefulness of Covid-19 Models. *Humanities and Social Sciences Communications*, 7 (1), pp. 54. <https://doi.org/10.1057/s41599-020-00553-4>
- Hideki Endo *et al* 2021. Conventional Risk Prediction Models Fail to Accurately Predict Mortality Risk among Patients with Coronavirus Disease 2019 in Intensive Care Units: A Difficult Time to Assess Clinical Severity and Quality of Care. *Journal of Intensive Care*, 9 (1), pp. 42. <https://doi.org/10.1186/s40560-021-00557-5>
- Eugene F. Fama 1970. Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance*, 25 (2), pp. 383–417. <https://doi.org/10.2307/2325486>
- Neil M Ferguson *et al* 2020. Impact of Non-Pharmaceutical Interventions to Reduce Covid19 Mortality and Healthcare Demand. 16 March 2020 Imperial College London pp. 1–20. doi.org/10.25561/77482
- Luisa Ferrari *et al* 2021. Modelling Provincial Covid-19 Epidemic Data Using an Adjusted Time-Dependent Sird Model. *International Journal of Environmental Research and Public Health*, 18 (12), pp. 6563. <https://www.mdpi.com/1660-4601/18/12/6563>
- Janyce Eunice Gnanvi *et al* 2021. On the Reliability of Predictions on Covid-19 Dynamics: A Systematic and Critical Review of Modelling Techniques. *Infectious Disease Modelling*, 6 pp. 258–272. <https://doi.org/https://doi.org/10.1016/j.idm.2020.12.008>
- Agus Hasan *et al* 2020. Superspreading in Early Transmissions of Covid-19 in Indonesia. *Scientific reports*, 10 (1), pp. 22386. <https://doi.org/10.1038/s41598-020-79352-5>
- Inga Holmdahl & Caroline Buckee 2020. Wrong but Useful — What Covid-19 Epidemiologic Models Can and Cannot Tell Us. *New England Journal of Medicine*, 383 (4), pp. 303–305. <https://doi.org/10.1056/NEJMp2016822>
- John P. A. Ioannidis *et al* 2020. Forecasting for Covid-19 Has Failed. *Int J Forecast*, <https://doi.org/10.1016/j.ijforecast.2020.08.004>
- Lyndon P. James *et al* 2021. The Use and Misuse of Mathematical Modeling for Infectious Disease Policymaking: Lessons for the Covid-19 Pandemic. *Medical Decision Making*, 41 (4), pp. 379–385. <https://doi.org/10.1177/0272989X21990391>
- Max S. Y. Lau *et al* 2020. Characterizing Superspreading Events and Age-Specific Infectiousness of Sars-Cov-2 Transmission in Georgia, USA. *Proceedings of the National Academy of Sciences of the United States of America*, 117 (36), pp. 22430–22435. <https://doi.org/10.1073/pnas.2011802117>
- Benoit B. Mandelbrot & Richard L. Hudson 2004. *The Misbehaviour of Markets : A Fractal View of Risk, Ruin and Reward*. London: Profile.
- Vicente Martínez 2021. A Modified Sird Model to Study the Evolution of the Covid-19 Pandemic in Spain. *Symmetry*, 13 (4), pp. 723. <https://www.mdpi.com/2073-8994/13/4/723>
- Stanley Milgram 1967. The Small World Problem. *Psychology Today*, 2 pp. 60–67.
- Alan Nicholson & Yiik-Diew Wong 1993. Are Accidents Poisson Distributed? A Statistical Test. *Accident Analysis & Prevention*, 25 (1), pp. 91–97. [https://doi.org/10.1016/0001-4575\(93\)90100-B](https://doi.org/10.1016/0001-4575(93)90100-B)

- Vilfredo Pareto 1896. *Cours D'economie Politique*. Lausanne: Rouge
- Nicholas Shackel 2011. Explaining Variation in Knowledge by Full Belief. *Proceedings of the 34th Ludwig Wittgenstein Symposium*. Eds Jaeger, C. and Loffler, W. Kirchberg am Wechsel, Austria: Ontos Verlag pp 95–102.
<https://doi.org/http://dx.doi.org/10.1515/9783110329018.95>
- Nicholas Shackel MS-a. Stakes and Full Belief.
- Nicholas Shackel MS-b. Wild Uncertainty and Pandemics.
- Sakshi Shringi *et al* 2021. Modified Sird Model for Covid-19 Spread Prediction for Northern and Southern States of India. *Chaos, Solitons & Fractals*, 148 pp. 111039. <https://doi.org/10.1016/j.chaos.2021.111039>
- Lawrence H. Summers 1986. Does the Stock Market Rationally Reflect Fundamental Values? *The Journal of Finance*, 41 (3), pp. 591–601.
<https://doi.org/https://doi.org/10.1111/j.1540-6261.1986.tb04519.x>
- Kaiyuan Sun *et al* 2021. Transmission Heterogeneities, Kinetics, and Controllability of Sars-Cov-2. *Science*, 371 (6526), <https://doi.org/10.1126/science.abe2424>
- Nicholas Nassim Taleb 2020. *Statistical Consequences of Fat Tails: Real World Preasymptotics, Epistemology, and Applications*. Scribe Media.
- United States Centre for Disease Control 2020. Covidview-05-01-2020.
<https://www.cdc.gov/coronavirus/2019-ncov/covid-data/pdf/covidview-05-01-2020.pdf>
- Duncan J. Watts & Steven H. Strogatz 1998. Collective Dynamics of ‘Small-World’ Networks. *Nature*, 393 (6684), pp. 440–442. <https://doi.org/10.1038/30918>
- Felix Wong & James J. Collins 2020. Evidence That Coronavirus Superspreading Is Fat-Tailed. *Proceedings of the National Academy of Sciences*, 117 (47), pp. 29416–29418. <https://www.pnas.org/content/pnas/117/47/29416.full.pdf>