

On Single Point Forecasts for Fat-Tailed Variables

Nassim Nicholas Taleb^{*†}, Yaneer Bar-Yam[‡], and Pasquale Cirillo^{§¶}

^{*}Universa Investments, [†]Tandon School of Engineering, New York University

[‡]New England Complex Systems Institute

[§]Institute for the Future, University of Nicosia, [¶]S-T-A-T-S GmbH, Switzerland

Accepted, *International Journal of Forecasting*

Abstract—We discuss common errors and fallacies when using naive "evidence based" empiricism and point forecasts for fat-tailed variables, as well as the insufficiency of using naive first-order scientific methods for tail risk management.

We use the COVID-19 pandemic as the background for the discussion and as an example of a phenomenon characterized by a multiplicative nature, and what mitigating policies must result from the statistical properties and associated risks. In doing so, we also respond to the points raised by Ioannidis et al.(2020)

July 27, 2020. We thank Pierre Pinson and Spyros Makridakis for helping organize this discussion, and John Ioannidis for his gracious engagement.

MAIN STATEMENTS

- (i)– Forecasting single variables in fat-tailed domains is in violation of both common sense and probability theory.
- (ii)– Pandemics are extremely fat-tailed events, with potentially destructive tail risk. Any model ignoring this is necessarily flawed.
- (iii)– Science is not about making single points predictions but about understanding properties (which can *sometimes* be tested by single point estimates and predictions).
- (iv)– Sound risk management is concerned with extremes, tails and their full properties, and not with averages, the bulk of a distribution or naive estimates.
- (v)– Naive fortune-cookie evidentiary methods fail to work under both risk management and fat tails, because the absence of raw evidence can play a large role in the properties.
- (vi)– There are feedback mechanisms between forecast and reaction that affects the validity of some predictions.
- (vii)– Individuals risks fail to translate into systemic risks under multiplicative processes.
- (viii)– One should never treat the "costs" of mitigation without taking into account the costs of the disease, and in some cases naive cost-benefit analyses fail (for sure when statistical averages are nonconvergent or invalid for tail risk purposes).




Fig. 1. A high variance Lognormal distributions. 85% of observations fall below the mean. Half the observations fall below 13% of the mean. The lognormal has milder tails than the Pareto which has been shown to represent pandemics.

- (ix) Historically, in the aftermath of the Great Plague, economies were less fragile to pandemics, equipped to factor-in effective mechanisms of containment (quarantines) in their operating costs. It is more cogent to blame overoptimization than reaction to disease.

The article is organized at three levels. First, we make general comments around the nine points in the Main Statements, explaining how single point forecasts is an unscientific simplification incompatible with processes with richer properties. Next we go deeper into the technical arguments. Finally we address specific points in Ioannidis et al. [16] and answer their arguments concerning our piece.

COMMENTARY

Both forecasters and their critics are wrong: At the onset of the COVID-19 pandemic, many research groups and agencies produced single point "forecasts" for the pandemic—most relied upon trivial logistic regressions, or upon the compartmental SIR model, sometimes supplemented with cellular automata, or with agent-based models assuming various social rules and behaviors. Apparently, the prevailing idea is that *producing a single numerical estimate* is how science is done, and how science-informed decision-making ought to be

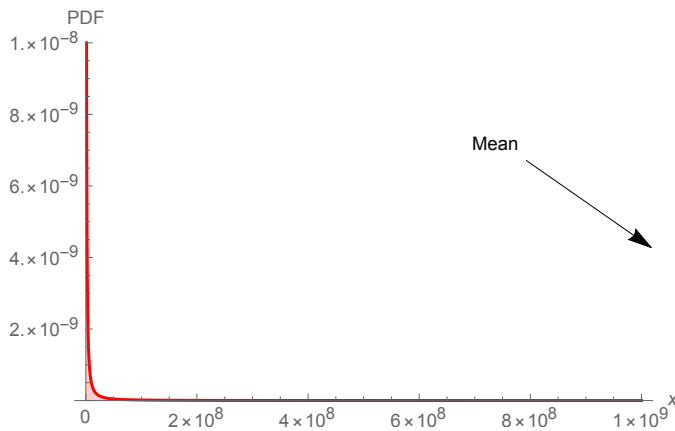


Fig. 2. A Pareto distribution with a tail similar to that of the pandemics. Makes no sense to forecast a single point. The "mean" is so far away you almost never observe it. You need to forecast things other than the mean. And most of the density is where there is noise.

done: bean counters producing precise numbers. And always within a narrowly considered set of options identified by the researchers.

Well, no. That is not how "science is done", at least in this domain, and that is not how informed decision-making should develop.

Furthermore, subsequently and ironically, many criticized the plethora of predictions produced, because these did not play out (no surprise there). This is also wrong, because both forecasters (who missed) and their critics (complaining) were wrong. Indeed, forecasters would have been wrong anyway, even if they had got their predictions right. In fact, as we will clarify throughout this article, 1) in some domains (i.e. under fat tails) naive forecasts are poor descriptors of a system (hence highly unscientific), even when they might appear reasonable; 2) for some functions (risk management related), or some classes of exposures (systemic ones), these forecasts are extremely misplaced.

Statistical attributes of pandemics: Using tools from extreme value theory (EVT), Cirillo and Taleb [6] have recently shown that pandemic deaths are *patently* fat-tailed¹—a fact some people like Benoit Mandelbrot (or one of the authors, in *The Black Swan* [32]) had already guessed, but never formally investigated. Even more, the estimated tail parameter α is smaller than 1, suggesting an apparently infinite risk [6], in line with destructive events like wars [4], [5], and the so-called "dismal" theorem [39]. Pandemics do therefore represent a source of existential risk. The implication is that much of what takes place in the bulk of the distribution is just noise, according to "the tail wags the dog" effect [6], [33]. And one should never forecast, pontificate, or theorise from noise! Under fat tails, all relevant and vital information lies in fact in

¹A non-negative continuous random variable X has a fat-tailed distribution, if its survival function $S(x) = P(X \geq x)$ is regularly varying, formally $S(x) = L(x)x^{-\alpha}$, where $L(x)$ is a slowly varying function, for which $\lim_{x \rightarrow \infty} \frac{L(tx)}{L(x)} = 1$ for $t > 0$ [7], [9], [10]. The parameter α is known as the tail parameter, and it governs the fatness of the tail (the smaller α the fatter the tail) and the existence of moments ($E[X^p] < \infty$ if and only if $\alpha > p$).

the tails themselves (hence in the extremes), which can show remarkably stable properties.

Remark 1: Observed events vs Observed Properties

Random variables with unstable (and uninformative) sample moments may still have extremely stable and informative tail properties, centrally useful for robust inference and risk taking. Furthermore, these reveal evidence.

*This is the central problem with the misunderstanding of **The Black Swan** [32]: some events may have stable and well-known properties, yet they do not lend themselves to prediction.*

Fortune-cookie evidentiary methods: In the early stages of the COVID-19 pandemic, scholars like Ioannidis [14] suggested that one should wait for "more evidence" before acting with respect to that pandemic, claiming that "we are making decisions without reliable data".

Firstly, there seems to be some probabilistic confusion, leading towards the so-called delay fallacy [13]: "if we wait we will know more about X, hence no decision about X should be made now."

In front of potentially fat-tailed random variables, more evidence is not necessarily needed. Extra (usually imprecise) observations, especially when coming from the bulk of the distribution, will not guarantee extra knowledge. Extremes are rare by definition, and when they manifest themselves it is often too late to intervene. Sufficient—and solid—evidence, in particular for risk management purposes, is already available *in the tail properties themselves*. An existential risk needs to be killed in the egg, when it is still cheap to do so. Events of the last few months have shown that waiting for better data has generated substantial delays, causing thousands of deaths and serious economic consequences.

Secondly, unreliable data²—or any source of serious uncertainty—should, under some conditions, make us follow the "paranoid" route. More uncertainty in a system makes precautionary decisions more obvious. If you are uncertain about the skills of the pilot, you get off the plane *when it is still possible to do so*. If there is an asteroid headed for earth, should we wait for it to arrive to see what the impact will be? We might counter that there were asteroids in the past that had devastating impacts, and besides we can calculate the physics. The logical fallacy runs deeper: "We did not see this particular asteroid yet" misses the very nature of the power of science to generalize (and classify), and the power of actions to possibly change the outcome of events. Similarly, if we had a hurricane headed for Florida, a statement like "We have not seen this hurricane yet, perhaps it will not be like the other hurricanes!" misses the essential role of risk management: to take preventive actions, not to complain ex

²Many of those complaining about the quality of data and asking for more evidence before taking action, even in extremely risky situations, rarely treat the inputs of their predictive models as imprecise [4], [38], stressing them, and performing serious robustness checks of their claims.

post. And if people take action boarding up windows, and evacuating, the claim "look it was not so devastating", that someone might afterwards make, should be considered closer to a lunatic conspiracy fringe than scientific discourse.

By definition, evidence follows and never precedes rare impactful events. Waiting for the accident before putting the seat belt on, or evidence of fire before buying insurance would make the perpetrator exit the gene pool. Ancestral wisdom has numerous versions such as *Cineri nunc medicina datur* (one does not give remedies to the dead), or the famous saying by Seneca *Serum est cavendi tempus in mediis malis* (you don't wait for peril to run its course to start defending yourself).

However, just as there are frivolous lawsuits there are frivolous risk claims and, as we will see further down, we limit these precautionary considerations to a precise class of fat tailed multiplicative processes –when there is systemic risk.

Remark 2: Fundamental Risk Asymmetry

For matters of survival, particularly when systemic, and in the presence of multiplicative processes (like a pandemic), we require "evidence of no harm" rather than "evidence of harm."

TECHNICAL COMMENTS

The Law of Large Numbers (LLN) and Evidence: In order to leave the domain of ancient divination (or modern anecdote) and enter proper empirical science, forecasting must abide by both evidentiary and probabilistic rigor. Any forecasting activity about the mean (or the parameter) of a phenomenon requires the working of the law of large numbers (LLN), guaranteeing the convergence of the sample mean at a known rate, when the number n of observations increases. This is surely well-known and established, except that some are not aware that, even if the theory remains the same, the actual story changes under fat tails.

Even in front of the most well-behaved and non-erratic random phenomenon, if one claimed fitness or non-fitness of a forecasting ability on the basis of a single observation ($n = 1$), he or she would be rightly accused of unscientific claim. Unfortunately, with fat-tailed variables that " $n = 1$ " error can be made with $n = 10^6$. In the case of events like pandemics, even larger $n \rightarrow \infty$ can still be anecdotal.

Remark 3: LLN and speed of convergence

Fat-tailed random variables with tail exponent $\alpha \leq 1$ are simply not forecastable. They do not obey the LLN, as their theoretical mean is not defined, so there is nothing the sample mean can converge to. But we can still understand several useful tail properties.

And even for random variables with $1 < \alpha \leq 2$, the LLN can be extremely slow, requiring an often unavailable number of observations to produce somehow reliable forecasts.

As a matter of fact, owing to preasymptotic properties, a conservative heuristic is to consider variables with $\alpha \leq 2.5$ as not forecastable in practice. Their sample mean will be too unstable and will require way too much data for forecasts to be reliable in a reasonable amount of time. Notice in fact that 10^{14} observations are needed for the sample mean of a Pareto "80/20", with $\alpha \approx 1.13$, to emulate the gains in reliability of the sample average of a 30-data-points sample from a Normal distribution [33].

Assuming significance and reliability with a low n is an insult to everything we have learned since Bernoulli, or perhaps even Cardano.

Also notice that discussing the optimality of any alarm system [3], [18], [31] trying to perform predictions on averages would prove meaningless under extremely fat-tails, i.e. when $\alpha \leq 2$, that is when the LLN works very slowly or does not work. In fact, even when the expected value is well-defined (i.e. $1 < \alpha < 2$), the non-existence of the variance would affect all the relevant quantities for the verification of optimality [8], from the size of the alarm to the number of correct and false alarms, from the probability of detection of catastrophes to the chance of undetected events. For all these quantities, the naive sample estimates commonly used would prove misleading. A solution could be the implementation on EVT-based approaches, possibly with the additional tools of [5] or [21], but at this stage nothing similar exists, to the best of our knowledge.

For this and other reasons specified later, the application of a non-naive precautionary principle [23] appears to be the viable solution in front of potentially existential risks.

Science is about understanding properties, not forecasting single outcomes: Figures 1 and 3 show the extent of the problem of forecasting the average (and so other quantities) under fat tails. Most of the information is away from the center of the distribution. The most likely observations are far from the true mean of the phenomenon and very large samples are needed for reliable estimation. In the lognormal case of Figure 1, 85% of all observations fall below the mean; half the observations even fall below 13% of the mean. In the Paretian situation of Figure 3, mimicking the distribution of pandemic deaths, the situation gets even worse: the mean is so far away that we will almost never observe it. It is therefore preferable to look at other quantities, like for example the tail exponent.

In some situations of fast-acting LLN, as (sometimes) in physics, properties can be revealed by single predictive experiments. But it is a fallacy to assume that a single predictive experiment can actually validate any theory; it is rather a single tail event that can falsify a theory.

Sometimes, as recently shown on the *International Journal of Forecasting* by one of the authors [34], a forecaster may find a single quantity that is actually forecastable, say the survival function. For n observations a tail survival function has an error of $o(\frac{1}{n})$, even when tail moments are not tractable, which is why many predict binary outcomes—as with the "superforecasting" masquerade. In [33], it is shown how—paradoxically—the more intractable the higher moments of the variable, the more tractable the survival function becomes. Metrics such as the Brier score are well adapted to binary

survival functions, though not to the corresponding random variables. That is why survival functions are essentially useless for risk management purposes. In insurance, for instance, one never uses survival functions for hedging, but rather expected shortfalls—binary functions are reserved to (illegal) gambling.³

We do not observe properties of empirical distributions:

A commentator (Andrew Gelman) [12] wrote "The sad truth, I'm afraid, is that Taleb is right: point forecasts are close to useless, and distributional forecasts are really hard."

The problem is actually worse. In fact, distributional forecasts are more than hard—and often uninformative. Building so-called empirical distributions by survival functions does not reveal tail properties since it will necessarily be censored and miss tail observations—those that under very fat tails (say $\alpha \leq 2$) harbor not most, but literally all of the properties [33]. In other words, probabilities are thin-tailed (since they are bounded by 0 and 1) but the corresponding payoff is not, so small errors in probability translate into large changes in payoffs. However, as further discussed in [6], the tail parameters are themselves thin-tailed distributed, hence reveal their properties rather rapidly. Simply, tail parameters extrapolate—while survival functions don't—and methods to measure the tail are quite potent.⁴

Uncertainty goes one way; errors in growth rates induce biases and massive fat tails for the quantity of interest: Consider the simple model

$$X_t = X_0 e^{r(t-t_0)},$$

where X_t represents the quantity of interest (say the number of fatalities in pandemics) between periods t_0 and t ,

$$r = \frac{1}{(t-t_0)} \int_{t_0}^t r_s ds$$

and r_s is an instantaneous rate.

Using the histograms of r and X , Figure 3 shows something fundamental: a well-behaved distribution, that of r , may lead to an untractable one, that of X ; furthermore, the more volatile r , the more downward-biased your observation of the mean of X .

Implication: one cannot naively translate between the rate of growth r and X_T , because errors in r could be small (but surely not zero), but their impact will be explosive on X , because of exponentiation.

Simply, if r is exponentially distributed (or part of that family), X will be power law. The tail α is a direct function of the variance: the higher the variance of r , the thicker the tail of X .

³The main problem is that the conditional expectation is not convergent: $\lim_{K \rightarrow \infty} \frac{1}{K} \mathbb{E}(X|X > K) > 1$, see [34] for a lengthy discussion.

⁴This also relates to the superforecasters masquerade mentioned earlier: building survival functions for tail assessments via sports-like "tournaments" as in [37], instead of using more rigorous approaches like EVT, is simply wrong and violates elementary probability theory.

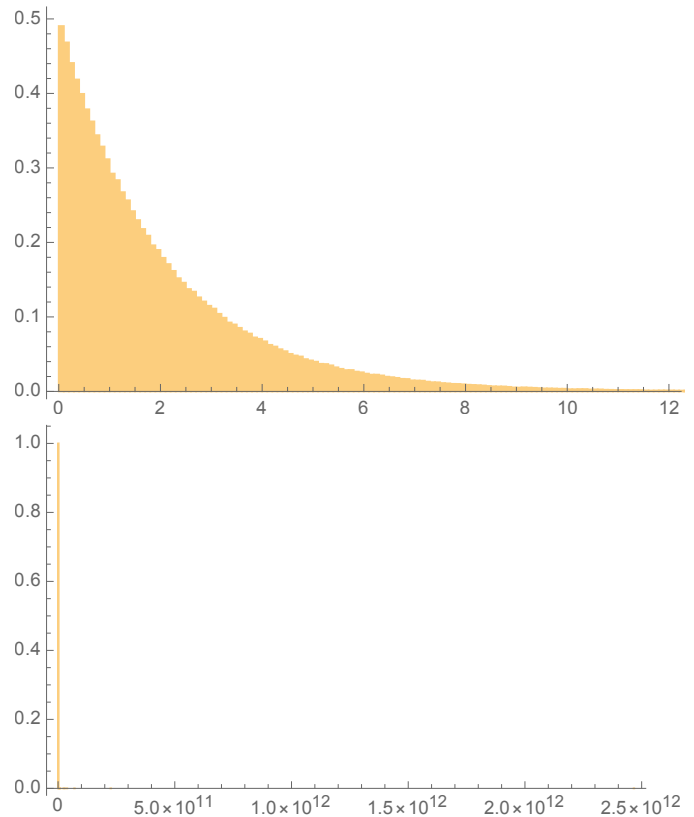


Fig. 3. Above, a histogram of 10^6 realizations of r , from an exponential distribution with sole parameter $\lambda = \frac{1}{2}$. Below, that of $X = e^r$. We can see the difference between the two distributions. The sample kurtosis are 9 and 10^6 respectively (in fact it is infinite for the second) —all values for the second one are dominated by a single large deviations.

Remark 4: Errors in Exponential Growth

- 1) Errors in growth rates of a disease increase the fatness of tails in the distribution of fatalities.
- 2) Errors in growth rates translate, on balance, into higher expected casualties.

We note that in the context of dynamical systems an exponential dynamics is defined as chaotic [19]. While the study of chaos often considers systems with fixed parameters and variable initial conditions, the same sensitivities arise due to variations in parameters; in this case, the value of contagion rate (R) and the social behaviors that affect it. Indeed this means that by changing human behavior, the dynamics can be strongly affected, thus allowing for the opening of opportunities for extinction.⁵

⁵One of the authors has shown [25] that with increasing global transportation there is a phase transition to global extinction with probability 1. This indicates that historical distributions don't account for the severity or frequency of current or future extreme events because the fat tailed distributions themselves are coalescing to unit probability extreme events over shorter time intervals due to global changes in societal behaviors. During this process the probability distributions for events in any time interval becomes progressively more weighted to large scale events. Thus historical decadal or century intervals between pandemics are inadequate descriptors of current risk.

Never cross a river that is 4 feet deep on average : Risk management (or policy making) should focus on tail properties and not on the body of probability distributions. For instance, The Netherlands have a policy of building and calibrating their dams and dykes not on the average height of the sea level, but on the extremes, and not only on the historical ones, but also on those one can expect by modelling the tail using EVT, via semi-parametric approaches [7], [9].

Science is not about safety: Science is a procedure to update knowledge; and it can be wrong provided it produces interesting discussions that lead to more discoveries. But real life is not an experiment. If we used a p-value of .01 or other methods of statistical comfort for airplane safety, few pilots and flight attendants would still be alive. For matters that have systemic effects and/or entail survival, the asymmetry is even more pronounced.

Forecasts can result in adjustments that make forecasts less accurate: It is obvious that if forecasts lead to adjustments, and responses that affect the studied phenomenon, then one can no longer judge these forecasts on their subsequent accuracy. Yet the point does not seem to be part of the standard discourse on COVID-19.

By various mechanisms, including what is known as Goodhart's law [30], a forecast can become a target that is gamed by participants—see also the Lucas' critique applying the point more generally to dynamical systems. In that sense a forecast can be a warning of the style "if you do not act, these are the costs" ⁶.

More generally, any game theoretical framework has an interplay of information and expectation that causes forecasts to become self-canceling. The entire apparatus of efficient markets—and modern economics—is based on such self-canceling aspect of prediction, under both rational expectations and an arbitrage-free world.

REMARKS SPECIFIC TO IOANNIDIS ET AL.

Systemic risks vs individual risks: A fundamental problem, in both [15] and [16], lies in ignoring scaling: systemic risks do not resemble (even qualitatively) individual risks. The macro and the micro-properties of contagious events, given their infective multiplicative nature, don't map directly onto one another.

Ioannidis et al. [15] write: "the average daily risk of dying from coronavirus for a person <65 years old is equivalent to the risk of dying driving a distance of 13 to 101 miles by car per day during that COVID-19 fatality season in 17 of the 24 hotbeds (...) For many hotbeds, the risk of death is in the same level roughly as dying from a car accident during daily commute."

Even if Ioannidis et al.'s computation were to hold true for one individual (it does not), conditionally on an excess of 10^3 of such individuals dying, the probability that the cause of death is COVID-19 and not a car accident converges to 1. When you die of a contagious disease, people around you are

⁶For instance Dr. Fauci's warning that the number of (verified) infections could reach 100K per day (*New York Times*, June 30, 2020) should not be interpreted as a forecast to be judged according to its accuracy; rather a signal about what could happen should one avoid taking action.

at risk of contagion, and they can then infect other people, in a cascading effect. It is quite elementary: car accidents are not contagious, while COVID-19 is. You cannot conflate the two objects: one is additive in the aggregate, the other is multiplicative. In [33], it has been shown that this is a severe error, leading to macroscopic blunders⁷.

Remark 5: Additive vs. Multiplicative Risks (Scaling of Probabilities)

Under multiplicative effects the risks for a collective do not scale up from the risks of an individual. Trivially, systemic risks can be extreme, where the individual ones are low, or vice-versa.

Trade-offs and Ergodicity: One could say: panic saves lives, but at what economic price? Let us put aside ethical arguments, and answer it, ignoring for a moment the value of human life.

The fact is that some classes of (systemic) risks require being killed in the egg, also from an economic point of view. The good news is that there are not so many—but pandemics as we said fall squarely within the category.

The "dismal" theorem [39] mentioned earlier tells us that it is an error to use trade-off analysis under existential risk. There have been many proofs of similar arguments on grounds of ergodicity, well-known by insurance companies since Cramèr: simply, you cannot use naive B-school costs-benefit analyses for Russian roulette, because of the presence of an absorption barrier [32]. But one should not blame Ioannidis et al. [15] for this error in reasoning: it has been shown to be unfortunately prevalent in the decision-science literature [24].

Remark 6: Ruin Problems

Traditional cost-benefit analysis fails to apply to situations where statistical averages are unreliable, if not invalid.

Moreover, it is not correct to assume, more or less implicitly, that a disease brings no or little costs, while mitigation is burdensome. There are indeed severe nonlinearities at play.

First of all, risk is beyond the simple and direct disease-specific mortality rate. In fact, letting the disease run above a certain threshold would compound its effect (in an explosive manner), because of the saturation of services, causing for example the displacement of other patients, many in potentially critical conditions; something that we have seen happening in the Region of Lombardy in Italy, in New York City, and elsewhere for several weeks during the spring of 2020 [29]. Furthermore, for survivors the illness itself represents a large economic drain, be it only from lost working hours, not counting the costs of hospitalization. And for every severe infection, there is an unspecified number of morbidities, with unknown (but definitely larger than zero) additional mortality

⁷Note that this is also a typical example of "size fallacy," in which different risky events are compared just on the basis of their probabilities of occurrence, without caring about their different nature [13].

and long term costs for the health system [1], [29], as it has been the case for other diseases like SARS [22].⁸

Remark 7: False Dichotomies

One should not treat the economy and the disease as separate independent items, particularly by viewing a naive trade-off between economic costs and pandemic mitigation.

Moreover, never underestimate consumers' (nonlinear) behavior. When risks are visible (and a pandemic definitely is), people tend to modify their behavior, rationally or not, also switching to alternatives, with nonlinear effects on the businesses concerned [26]. This is the reason why the airline industry in the U.S. manages to have fewer than 1 fatal crash in 25×10^6 flights (and aims at an even more favorable ratio). One may claim that it is irrational to spend so much of our resources mitigating plane crashes, but airline companies know that, in case of fewer checks and efforts, consumers would then probably switch to other companies, if not directly to other types of transportation.

Take the hospitality industry. Unless there is once again comfort on the part of the public, restaurants and hotels will be unprofitable. The rule of thumb in NYC is that a drop of 15% in revenues is sufficient to make a restaurant shutter permanently; there has been a large drop in restaurant attendance in Sweden where the state did not enforce lockdowns, owing to a high rate of voluntary self-isolation [17].

The United States (and many other countries worldwide) have spent trillions of dollars on sophisticated weaponry in the past decades, to counter *uncertain* threats. It would be a good idea to question these expenditures first, before doubting the spending to stave off certain pandemics.

Likewise, it would be a good idea to question first the excessive burden on Western economies, particularly the U.S., of measures taken to ensure workplace and transportation safety which, we saw, are driven by the legal system and the tort mechanisms.

Remark 8: Domain Dependence

It is not rational to worry about pandemic costs (extremely fat-tailed exposure), while not also questioning other sizable insurance-style expenditures for transportation and workers safety.

It is therefore incorrect to claim that it is the authorities'

⁸**Geronticide:** This discussion does not even cover the ethical discussion of trade-offs and their inapplicability in some domains, perhaps the most central discussion. At what price will you kill your parents/grandparents? A million dollars? Ten million? A billion? Furthermore, the fact that older people are more vulnerable to the disease brings considerations of geronticide (senicide): one misses that the silver rule [32] commands treating older generations under a moral liability, as one wishes to be treated by the next generation. Letting the disease run through older generations violates the interdicts on geronticide and intergenerational obligations. The fact that your parents did not sacrifice their own parents creates an obligation to not sacrifice them; your children will spare you in turn, under the same rule.

response to the pandemic that caused unemployment in the transportation or hospitality industries. As a matter of fact, the arguments proposed by two of the authors [23], last January 2020, were aimed at lowering the economic effect of the pandemic: prevention is orders of magnitude cheaper than the cure—recall that *sed prior est sanitas quam sit curatio morbi*.

We note that many comments of the type "the pandemic has caused *only* 640K fatalities" (as of July 25, 2020) simply ignore the fact that, in practically every location subject to the pandemic, there has been local or governmental action to mitigate it—we do not consider the counterfactual of "what if" because it is not visible.

Remark 9: Economic Fragility

The argument in [32] is that we live in an over-optimized environment, in which a slight drop in sales or a change in consumer preferences may cause wild interlocking industry collapses. This nonlinearity is similar to "large a movie theater with a very small door at the times of fire."

It is more cogent to blame the over-optimized economic structure than the general reaction to the disease.

Early Mitigation and Economic History: We note here that while the Great Plague took place in the fourteenth century, quarantines were enforced five centuries later as economies understood they could not afford recurrences. Between the Habsburg and the Ottoman Empire, there were *lazarettos* along the border, and every active Mediterranean port enforced quarantines for travelers along the expanded silk road, while pilgrim routes were subjected to similar measures. For instance, in the 1830s, in *the Count of Monte Christo*, a traveler from Paris to Ioanina (where Prof. Ioannidis was previously located), had to spend four days in quarantine to get there, while there was no particular threat of disease. In fact, the novelist was underestimating, for historian records show mandatory nine days for ordinary travelers and fifteen days for merchants according to [28], [27]. Economies adapted to early mitigation throughout the centuries preceding our era. Furthermore, the Ottoman Empire has ready *lazarettos* for additional quarantining along specified stations at the first signs of a pandemic.

Mitigation has another effect: to delay and temporize, while we can understand the properties of the disease. While initial treatments are under high opacity, later treatments allow for gains of collective experience⁹.

Selection Bias and Class of Events: In [16], the authors erroneously maintain that choosing tail events as done by [6] is "selection bias". Actually, the standard technique there used is the exact opposite of selection bias: in EVT, one purposely focuses on extremes to derive properties that influence the outcomes, especially from a risk management point of view. One could more reasonably argue that the data in [6] do not

⁹Aside from considerations of geronticide, when the costs of the Swedish experiment are finally told, one of the factors will be the early loss of life when later (current) medical practice would have saved them even before a vaccine.

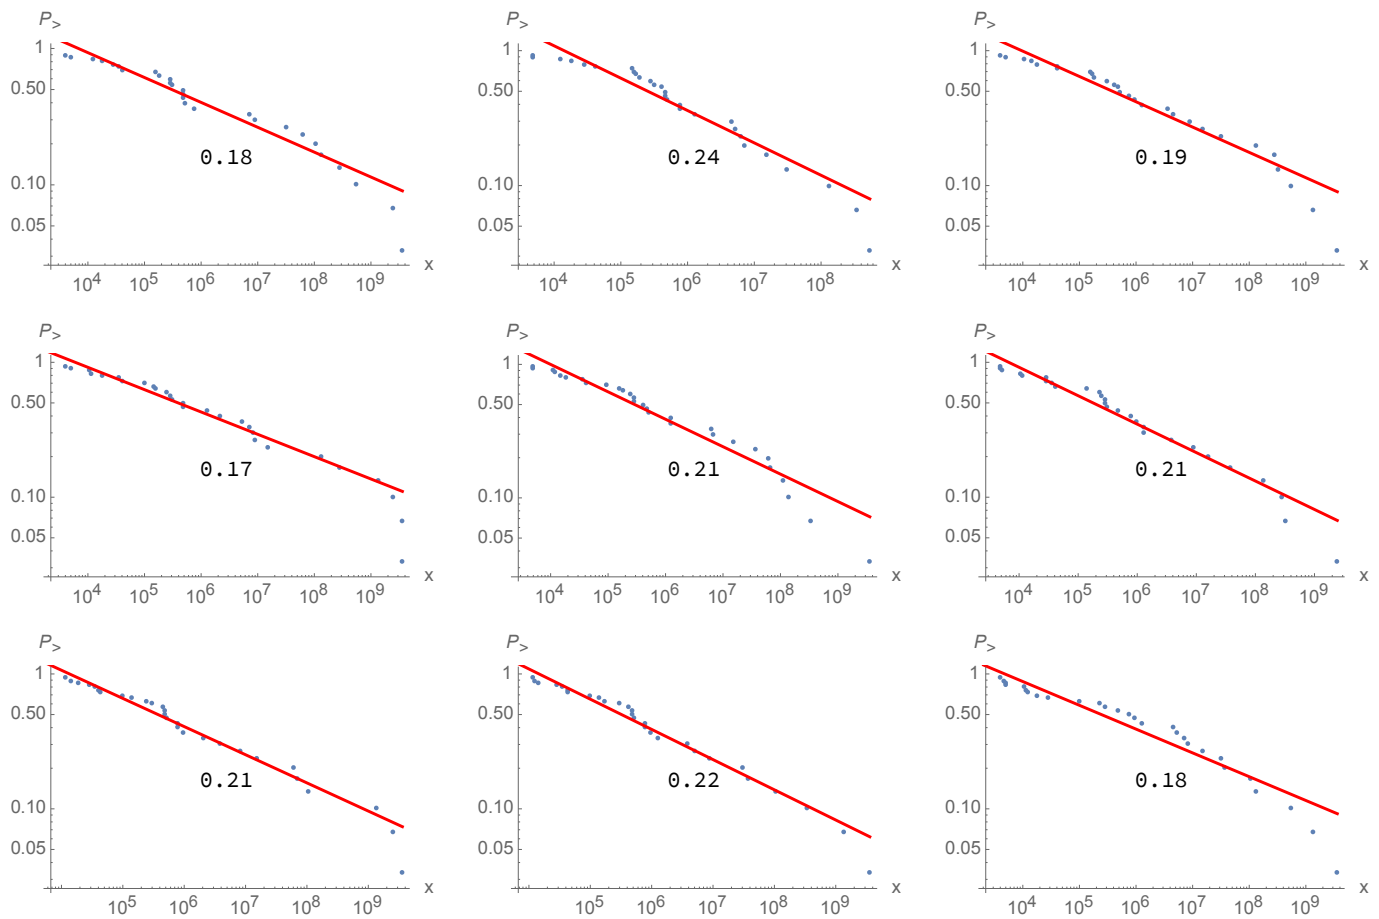


Fig. 4. Zipf plots (log-log plots of the empirical survival function $P_{>}$) for nine random selections of 30 out of the 72 pandemics in [6]. The number in the center represents the naive (OLS) estimate of the tail parameter α , readable as the absolute slope of the red negative line. The values of α appear to be stable notwithstanding the sampling, signalling the robustness of the approach and the inconsistency of the "selection bias" critique. The values are also in line with the more rigorous EVT-based findings of [6].

contain all the extremes, but, by jackknifing and bootstrapping the data, the authors actually show the robustness of their results to variations and holes in historical observations: the tail index α is consistently lower than 1. In Figure 4 a simple illustration is given, showing that one can be quite radical in dealing with the uncertainty in pandemic fatalities, and still find out that the findings of [6] hold true.

When the authors in [16] state that "Tens of millions of outbreaks with a couple deaths must have happened throughout time," to support their selection bias claim against [6], they seem to overlook the fact that the analysis deals with pandemics and not with a single stermutation. The class of events under considerations in [6] is precisely defined as "pandemics with fatalities in excess of $1K$," and their dataset likely contains most (if not all) of them. Worrying about many missing observations in the left tail of the distribution of pandemic deaths is thus misplaced.

Conditional information: One may be entitled to ask: as we get to know the disease, do the tails get thinner? Early in the game one must rely on conditional information, but as our knowledge of the disease progresses, shouldn't we be allowed to ignore tails?

Alas, no. The scale of the pandemic might change, but

the tail properties will remain invariant. Furthermore, there is an additional paradox. If one does not take the pandemic seriously, it will likely run wild (particularly under the connectivity of the modern world, several orders of magnitude higher than in the past [2]). And diseases mutate, increasing or decreasing in both lethality and contagiousness. The argument would therefore resemble the following: "we have not observed many plane crashes lately, let's relax our safety measures".

Finally, we conclude this section with an encouraging point: fat tails do not make the world more complicated and do not cause frivolous worries, to the contrary. Understanding them actually reduces costs of reaction because they tell us what to target—and when to do so. Because network models tend to follow certain patterns to generate large tail events [2], [11], in front of contagious diseases wisdom in action is to kill the exponential growth in the egg via three central measures 1) reducing super-spreader events; 2) monitoring and reducing mobility for those coming from far-away places (via quarantines); 3) looking for cheap measures with large payoffs in terms of the reduction of the multiplicative effects (e.g. face

masks¹⁰). Anything that "demultiplies the multiplicative" helps [35].

Drastic shotgun measures such as lockdowns are the price of avoiding early traveler quarantines and border monitoring; they can be *–temporarily and cum grano salis* –of help, especially in the very early stages of the new contagious disease, when uncertainty is maximal, to help isolating and tracing the infections, and also buying some time for understanding the disease and the way it spreads. Indeed such drastic and painful measures can carry long-lasting damages to the system, not counting an excessive price in terms of personal freedoms.

But they are the price of not having a good coordinated tail risk management in place –to repeat, border monitoring and control of superspreader events being the very first such measures. And lockdowns are the costs of ignoring arguments such as increased connectivity in our environment and conflating additive and multiplicative risks.

To conclude, as the trader lore transmitted by generations of operators goes, "if you must panic, it pays to panic early."

The Ottoman Empire integrated Byzantine knowledge accumulated since at least the Plague of Justinian; it is sad to see ancient cultures more risk-conscious, better learners from history, and economically more effective than modern governments. They avoided modern "evidence based" reductions that, as we saw, are insulting to both science and wisdom. And, had it not been for such a collective ancestral risk-awareness and understanding of asymmetry, we doubt that many of us would be here today.

Now, what did we learn from the pandemic? That an intelligent application of the precautionary principle [23] consists in formulating decisions that are wise in both foresight and hindsight. Here again, this is ancient: it maps to Aristotle's *phronesis* as presented in his *Nichomachean Ethics*.

REFERENCES

- [1] M. Ackermann, S.E. Verleden, M. Kuehnel, A. Haverich, T. Welte, F. Laenger, A. Vanstapel, C. Werlein, H. Stark, A. Tzankov, W.W. Li, V.W. Li, S.J. Mentzer, D. Jonigk (2020). Pulmonary Vascular Endothelialitis, Thrombosis, and Angiogenesis in Covid-19. *New England Journal of Medicine* 383, 120-128.
- [2] R. Albert and A.-L. Barabasi (2002). Statistical mechanics of complex networks. *Reviews of Modern Physics* 74: 47.
- [3] M.A. Amaral-Turkman, K.F. Turkman (1990). Optimal alarm systems for autoregressive process; a Bayesian approach. *Computational Statistics and Data Analysis* 19, 307-314.
- [4] P. Cirillo, N.N. Taleb (2016). On the statistical properties and tail risk of violent conflicts. *Physica A: Statistical Mechanics and its Applications* 452, 29-45.
- [5] P. Cirillo, N.N. Taleb (2016). Expected shortfall estimation for apparently infinite-mean models of operational risk. *Quantitative Finance* 16, 1485-1494.
- [6] P. Cirillo, N.N. Taleb (2020). Tail risk of contagious diseases. *Nature Physics* 16, 606-613.
- [7] L. de Haan, A. Ferreira (2006). *Extreme Value Theory: An Introduction*. Springer.
- [8] J. de Maré (1980). Optimal prediction of catastrophes with application to Gaussian process. *Annals of Probability* 8, 841-850.
- [9] P. Embrechts, C. Klüppelberg, T. Mikosch (2003). *Modelling Extremal Events*. Springer.
- [10] M. Falk, J. Hüslér J, R. D. Reiss R-D (2004). *Laws of small numbers: extremes and rare events*, Birkhäuser.
- [11] U. Garibaldi, E. Scalas (2010). *Finitary Probabilistic Methods in Econophysics*. Cambridge: Cambridge University Press.
- [12] A. Gelman (2020). Some forecasting for COVID-19 has failed: a discussion of Taleb and Ioannidis et al.. Available online at <https://statmodeling.stat.columbia.edu/2020/06/17/some-forecasting-for-covid-19-has-failed-a-discussion-of-taleb-and-ioannidis-et-al>
- [13] S.O. Hansson (2004). Fallacies of Risk. *Journal of Risk Research* 7, 353-360.
- [14] J.P.A. Ioannidis (2020). A fiasco in the making? As the coronavirus pandemic takes hold, we are making decisions without reliable data. *Stat March 17*, <https://www.statnews.com/2020/03/17/a-fiasco-in-the-making-as-the-coronavirus-pandemic-takes-hold-we-are-making-decisions-without-reliable-data/>.
- [15] J.P.A. Ioannidis, C. Axfors, D.G. Contopoulos-Ioannidis (2000a). Population-level COVID-19 mortality risk for non-elderly individuals overall and for non-elderly individuals without underlying diseases in pandemic epicenters. *Environmental Research* 188, 109890.
- [16] J.P.A. Ioannidis, S. Cripps, M.A. Tanner (2020b). Forecasting for COVID-19 has failed. *International Institute of Forecasters*, available at <https://forecasters.org/blog/2020/06/14/forecasting-for-covid-19-has-failed>
- [17] S. C. L. Kamerlin, Peter M Kasson (2020) Managing COVID-19 spread with voluntary public-health measures: Sweden as a case study for pandemic control, *Clinical Infectious Diseases*, <https://doi.org/10.1093/cid/ciaa864>
- [18] G. Lindgren (1975). Prediction for a random time point. *Annals of Probability* 3, 412-433.
- [19] E.N. Lorenz (1963). Deterministic non-periodic flow. *Journal of the Atmospheric Sciences*. 20(2): 130-141.
- [20] R. Lucas (1976). Econometric Policy Evaluation: A Critique. In K. Brunner, A. Meltzer (eds.). *The Phillips Curve and Labor Markets*. Carnegie-Rochester Conference Series on Public Policy 1, 19-46. American Elsevier.
- [21] J. Nešlehová, P. Embrechts, V. Chavez-Demoulin (2006). Infinite-mean models and the LDA for operational risk. *Journal of Operational Risk* 1, 3-25.
- [22] J.C. Ngai, F.W. Ko, S.W. To, M. Tong, D.S. Hui DS (2010). The long-term impact of severe acute respiratory syndrome on pulmonary function, exercise capacity and health status. *Respirology* 15, 543-550.
- [23] J. Norman, Y. Bar-Yam, N.N. Taleb (2020). Systemic Risk of Pandemic via Novel Pathogens - Coronavirus: A Note. *New England Complex Systems Institute*.
- [24] O. Peters, M. Gell-Mann (2016). Evaluating gambles using dynamics. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 26, 023103.
- [25] E. M. Rauch and Y. Bar-Yam (2006) Long-range interaction and evolutionary stability in a predator-prey system, *Physical Review E* 73: 020903.
- [26] N.L. Rose (1992). Fear of flying? Economic analysis of airline safety. *Journal of Economic Perspectives*, 6(2), 75-94.
- [27] A Roberts (2017) in Nukhet Varlik (ed.). *Plague and Contagion in the Islamic Mediterranean*. Arc Humanities Press.
- [28] G Sariyildiz and O Daglar Macar (2017) in Nukhet Varlik (ed.). *Plague and Contagion in the Islamic Mediterranean*. Arc Humanities Press.
- [29] K. Søreide, J. Hallet, J.B. Matthews, A.A. Schnitzbauer, P.D. Line, P. Lai, J. Otero, D. Callegaro, S.G. Warner, N.N. Baxter, C. Teh, J. Ng-Kamstra, J.G. Meara, L. Hagander, L. Lorenzon (2020). Immediate and long-term impact of the COVID-19 pandemic on delivery of surgical services. *The British Journal of Surgery*, <https://doi.org/10.1002/bjs.11670>.
- [30] M. Strathern (1997). Improving ratings: Audit in the British University system. *European Review* 5, 305-321.
- [31] A. Svensson, R. Lindquist, G. Lindgren (1996). Optimal prediction of catastrophes in autoregressive moving average processes. *Journal of Time Series Analysis* 17, 511-531.
- [32] N.N. Taleb (2001-2018). *Incerto: Fooled by Randomness, The Black Swan, The Bed of Procrustes, Antifragile, and Skin in the Game*. Penguin
- [33] N.N. Taleb (2020a). *Statistical Consequences of Fat Tails*. STEM Academic Press.
- [34] N. N. Taleb (2020b). On the statistical differences between binary forecasts and real-world payoffs. *International Journal of Forecasting*, in press <https://doi.org/10.1016/j.ijforecast.2019.12.004>.

¹⁰Most of the trillions spent could have been saved if authorities understood the double nonlinearities in face masks: 1) the compounding effect of both parties having protection, 2) the nonlinearity of the dose response with disproportional drop in the probability of infection from a reduction in viral load [35].

- [35] N.N. Taleb (2020c) The Masks Masquerade. <https://medium.com/incerto/the-masks-masquerade-7de897b517b7>
- [36] N.N. Taleb, P. Cirillo (2019). The Decline of Violent Conflict: What do the data really say? In A. Toje, N.V.S. Bård, eds. *The Causes of Peace: What We Know Now*. Nobel Symposium Proceedings. Norwegian Nobel Institute, 57-85.
- [37] P.E. Tetlock, D. Gardner (2016). *Superforecasting: The art and science of prediction*. Random House.
- [38] R. Viertl (1995). *Statistical Methods for Non-Precise Data*. CRC Press.
- [39] M.L. Weitzman (2009). On modeling and interpreting the economics of catastrophic climate change. *Review of Economics and Statistics* 1, 1-19.