

Don't Throw the Baby out with the Pandemic: A Comment on “Naive Probabilism” by Harry Crane

Lucas Mentch*

Date February 23, 2021

Abstract

In early March 2020, just as the COVID-19 pandemic was beginning to unfold in the United States, Harry Crane published a thought provoking article entitled, “Naive Probabilism” in which he is critical of what he sees as an exclusive dependence on probabilistic reasoning for decision making, particularly in high-stakes settings. Much of the pandemic-related damage, he argues, is the result of “blind reliance on naive probabilism and the denial of basic common sense.” Here I push back on this perspective, arguing instead that what Crane perceives as a “denial of basic common sense ” can be better seen as a failure to properly understand uncertainty and estimate the true costs associated with various actions. Pointing to examples of poorly executed probabilistic reasoning – even in extreme cases – is not sufficient grounds for denying the merits of probabilistic reasoning itself; to do so is the epitome of throwing the baby out with the bathwater – or, in this case, with the pandemic.

1 Introduction

Last year was unlike any other ever experienced for virtually every living American. As COVID-19 began to spread through the world, fear of the novel coronavirus in the US was on the rise. In late February and early March, respectively, two prominent academics – Cass Sunstein (Harvard Law) and John Ioannidis (Stanford Medicine) – wrote critically and dismissively of such fears (Ioannidis 2020; Sunstein 2020). In expressing their views, Sunstein and Ioannidis relied heavily on basic probability concepts, ultimately arguing that the protective measures being taken on both a personal and societal level were not necessarily justified based on the available evidence at the time.

In response, Harry Crane (Rutgers Statistics) wrote critically of these views, coining the idea of “naive probabilism”, which he defines as the “*(naive) view, held by many technocrats and academics, that all rational thought boils down to probability calculations*” Crane (2020). “*The 2019-2020 coronavirus outbreak (COVID-19)*”, he argues, “*is a living example of the dire consequences of such probabilistic naiveté.*” In his abstract, he goes on to say “*The major damage caused by the spread of COVID-19 is attributable to a failure to act and a refusal to acknowledge what was in plain sight. This shared negligence stems from a blind reliance on naive probabilism and the denial of basic common sense by global and local leaders, and many in the general public.*”

*University of Pittsburgh, Department of Statistics

In place of this kind of “probabilistic reasoning”, Crane argues for a kind of “possibilistic reasoning” in high-risk settings. In his words, “*probability must be neglected to secure survival.*” We should forget how improbable such extreme outcomes may be and instead “*prepare for as many as are feasible*”.

As I write this now, with the benefit of nearly a year’s worth of hindsight, it’s easy to see the wisdom in Crane’s thinking rather than that of Sunstein and Ioannidis. Given the unmistakably devastating global impact of COVID-19, it’s difficult if not impossible to argue that we should have better prepared for this possibility.

But Crane’s critique, in my view, goes too far. In this comment, I’ll argue, in essence, that early covid-related failures can be more easily seen as a failed *application* of probabilistic reasoning, rather than a failure of probabilistic reasoning itself. Later in the paper, I’ll argue that Crane’s critique of naive probabilism is itself naive and that the kind of “possibilistic reasoning” that he seems to advocate could lead to disastrous outcomes if implemented on a societal level.

Disclaimer: A few days before publicly posting this comment, I had a short conversation with Harry about his paper. Based on that conversation, he has since made clear (publicly) that my interpretations below of some of his recommendations – in particular, his comment “*At a societal level, do everything possible to reduce the spread.*” – were not what he intended them to be. However, because I had already written out my responses prior to that conversation and because I suspect other readers may have had similar interpretations, I’m electing to leave them in this comment in the interest of transparency. Nonetheless, readers should be aware upfront that many of the things I argue *against*, especially in later sections, are not necessarily the things Harry intended to argue *for*.

2 Points of Agreement

Let me open with an anecdote:

Very early on in the pandemic – likely within a week or so of when Crane’s article first appeared – I vividly recall a conversation I was having with an older gentleman who was also a former cancer patient. After the obligatory, “wow, this is crazy” part of the conversation, he asked if I was doing anything differently in everyday life on account of the virus. My wife and I had two young children at the time, and so I explained some of the precautions we were taking in running everyday errands – generally only one of us would go into stores and we would do so while wearing a mask and gloves (yes, gloves seemed reasonable at the time). Though this conversation took place over the phone, I could sense a raised eyebrow. As I recall, he then responded to me with something like the following: *What are you so worried about? They already said you don’t have to do any of that stuff.* I explained my reasoning nonetheless and urged him to take similar precautions. I also predicted (correctly) that the “science” might well begin to change once mask supply began to return to normal levels.

2.1 Reasoning with Possibilities

I open with this story because (i) it’s always fun to say “I called it” and (ii) I think it lends support to the kind of decision-making Crane seems to advocate in such situations. In deciding to wear a mask at the time, I wasn’t making any kind of precise probability

calculation. My reasoning was much simpler: this virus *could* be really dangerous, taking these precautions *could* reduce my chances of getting it, these precautions come at a relatively low cost, therefore I'll just take them. I was reasoning in terms of possibilities, not probabilities. To use Crane's words, I was "*thinking about plausibility and possibility long before precisely quantifying probabilities of any kind.*" I was doing what most would simply call "erring on the side of caution".

Even so, I don't see this simplified form of decision making as a rejection of probabilistic reasoning; the next several sections press deeper on this point.

3 Possibilities as Probabilities

Suppose a friend wants to bet you \$1 that the next roll of a die will be a '1' – if the die lands on a '1', you lose \$1, if it lands on any other number, you win \$1. You take that bet, right? But now suppose your friend is a billionaire and they don't want to bet you \$1, they want to bet you your entire net worth. Do you still take it?

I'll admit that at present, I count very few billionaires among my circle of friends¹ so I can't say for sure what I would do. Nonetheless, I can't imagine I would actually take that second bet. But why? The probability of me winning didn't change – I'm the same (heavy) favorite in both settings. So why would I take the bet for \$1 but turn down the bet for my entire net worth?

Crane (2020) gives the obvious answer: "*The gamble is a simple mathematical problem, but the gambler is a complex person with a life outside of gambling, a need to maintain a bankroll, pay a mortgage, raise a family, etc.*" This is a very simple and intuitive point: objective costs aside, losing the second bet comes at a tremendous *personal cost*. I don't just lose some substantial amount of money, I might also lose my home and the ability to feed my children etc. By contrast, the *personal gain* I get from winning becomes far less appealing. Relatively speaking, if I win, my life gets *a little* better; if I lose, my life gets *a lot* worse.

But here is where I suspect Crane and I may part ways. In making the decision to not accept the second bet, I did not, in my view, abandon probabilistic reasoning. I still relied on a classic decision theoretic framework, I simply took into account all costs – both objective and personal. I did not simply say that there is a *possibility* of a very negative outcome and do everything in my power to avoid it.

To be even more explicit, I'm still taking probability into account. In this particular case, I wouldn't take the bet. But could you tilt the odds so heavily in my favor that I would? Almost certainly. Could you tilt the odds so heavily in my favor that I would take the bet at age 80? Maybe not.

Probabilities in a Pandemic

How does all of this fit into the current topic of mitigating a pandemic? For starters, to (over-) simplify things we could imagine our initial COVID-19 response options were limited to only two: "do" or "do not" take preventative action – this takes the place of "do" or "do not" take the bet in the example above. Here though (depending on how we want to frame the problem) the cost-benefit imbalance is reversed: we might imagine a relatively small cost associated with taking action, but an enormous benefit to preventing

¹Please feel free to reach out if you're a billionaire in need of a friend.

a pandemic. Because of this large imbalance, the probability of a pandemic happening must be very small in order for us *not* to take action. Put differently, if there is even a small chance of a global pandemic, we ought to take action to try and prevent it.

Here, Crane (2020) is quite explicit about his disagreement with this kind of formulation. The costs – both objective and personal – are simply so large in a pandemic, he seems to argue, that it no longer makes sense to think in terms of probabilities. We should instead think only in terms of possibilities. Formalizing this slightly, if we let p denote the probability of a global pandemic, Crane (2020) seems to suggest that if $p > 0$, then that alone is sufficient to warrant action (possibly extreme – “*At a societal level, do everything possible to reduce the spread.*”) to prevent the spread of the virus.

Decision-Making Under Uncertainty

There remains an elephant in the room with respect to my above formulation of Crane’s position – one major difference that seems very important to the case laid out in Crane (2020). In my description above, I’m implicitly assuming we know what p is and therefore we need only determine the costs of various actions in order to decide what to do. In reality, we certainly do not and so if we want to remain within the decision-making framework above, we have to estimate it. Crane (2020) makes the following statement in regard to this issue:

“Probability calculations are very precise, and for that reason alone they are also of very limited use. Outside of gambling, financial applications, and some physical and engineering problems—and even these are limited—mathematical probability is of little direct use for reasoning with uncertainty.”

Here I would disagree quite strongly: I see no reason why probability calculations need to be very precise in order to be useful.

Let’s return briefly to the die-rolling bets discussed at the beginning of this section. Suppose I didn’t know how many sides the die has – maybe it has the standard 6, maybe it has 10, maybe it has 20. I’m not sure that anything changes – I still take the \$1 bet; I still reject the net worth bet. Suppose I know the die has at least 1000 sides (I win at least 99.9% of the time) – there’s still a lot of uncertainty but regardless, I’m much more likely to take that bet regardless of whatever the true probability happens to be.

There are two important points I think ought to be stressed here:

- I’m intentionally using the language of “very precise” – I’m assuming that some minimal notion of accuracy/precision exists. This is obviously necessary in order for probability estimates to be of *some* value, I just don’t think that high precision is necessary in all settings (though it’s obviously desirable).
- Certainly there are situations where you do need a high degree of precision in order to have any confidence that you are making an optimal decision – imagine, for example, any setting in which two outcomes have nearly the same (true, unobserved) probability, but where the costs are very different. It seems to me, however, that Crane (2020) is making the very different argument that high precision is *always* necessary.

4 Was that a circle?

After reading the previous section, readers might wonder whether I just went on a long-winded detour only to arrive back at the same kind of reasoning being argued for in Crane (2020). Crane is arguing that the cost of a global pandemic is so high that we shouldn't be concerned about the precise probability of it occurring – if it's a possibility, we must act. I'm arguing that in situations with substantially lopsided costs, I'm quite happy to utilize probabilistic reasoning even with somewhat imprecise probability estimates – we don't necessarily need to wait for additional data in order to act. So maybe what I'm calling "probabilistic reasoning with uncertainty" is simply the same as Crane's "possibilistic reasoning."

But I think this is not the case for at least two reasons:

1. **In reality, there are more than two possible options.** In the simple framework outlined above, there are two choices we can make in the face of a potential pandemic: act to stop it or don't. But in reality, there is a near continuous range of options in between and, crucially, each have very different costs associated with them. Encouraging people to wash their hands and stay well-hydrated is one low-cost option. Encouraging social distancing and mask wearing is a bit higher-cost, but still reasonable in the eyes of many. Government mandated shutdowns of schools and non-essential businesses feels much higher. Government mandated stay-at-home orders or even barricading people in their homes are higher still.

If we're willing to accept then that some of these actions come at a higher societal cost, then it would seem reasonable to take low-cost actions when the risk for a pandemic is low and (possibly) accept higher-cost actions at higher risks.

But under the framework advocated in Crane (2020), there are no such things as higher risks or larger probabilities. There are only possibilities. And if there is *any* possibility of a pandemic – any $p > 0$ – then *any* actions are justified. I will return to this point in the section below.

2. **I don't accept the premise that choosing to act to try and prevent the spread is always – especially in the long term – necessarily a relatively low-cost action.** I realize that this may sound a bit absurd on the surface – no, I'm not arguing that short-term mask wearing or social distancing is somehow "as costly" as a global pandemic. Bear with me and let me begin with another anecdote.

A few years ago, I was in my office at the University of Pittsburgh late one afternoon and the fire alarm went off. People began to grumble and slowly shuffle outside, eventually congregating on the walkway in front of the building. A few minutes later, a police officer showed up and headed inside. Then a fire truck showed up. Then another. Then another. Everyone's reaction was exactly the same: *Wait ... is there an actual fire?*²

There's a reason why fire alarms, at least in the United States, are met with sighs and eye rolls rather than fear and panic. Most of us have suffered through so many of these alarms that were either drills (no danger) or the result of small accidental

²As I recall, I believe there was a small construction-related fire on the roof, though it didn't appear anyone was in any significant danger.

triggers³ (no real danger) that we've become conditioned to associate the sound with frustration and inconvenience rather than genuine alarm. That's why it's common on college campuses to see people – including and especially graduate students, postdocs, and faculty – turn off their light, put on headphones, and keep working rather than go outside for a few minutes.

Let's compare the possible outcomes: on one hand, you have to stop working and walk outside for a few minutes; on the other, you get a few more minutes of work done before burning alive. I'll go out on a limb and guess that most people would prefer the former. So why do they so often try to stay?

Crane could argue that I'm making exactly his point here – those people *should* be fleeing the building in panic. They're incorrectly applying probabilistic reasoning where they shouldn't. Maybe when you hear the fire alarm go off, your instinct should be to immediately jump through the nearest window. Use some expensive lab equipment to smash it open if you have to. Will you die if you don't? Probably not, but, in Crane's words, "*probability must be neglected to secure survival.*"

I would argue that a more reasonable approach in this scenario would be to look around, assess the situation, and proceed to exit the building in your normal fashion unless there is evidence to suggest otherwise. If you begin walking down the stairs and see a fire raging on the floors below, smashing a nearby window may begin to look like a very reasonable option. But immediately resorting to this kind of extreme action without any evidence of immediate danger does not seem reasonable.

To be clear, I'm not arguing that we should get rid of fire alarms. I'm not sure any reasonable person would argue that they are not important, potentially life-saving devices. But I think it is reasonable to worry about their overuse. Once people have begun to experience an onslaught of needless alarms, the alarm itself becomes a less effective tool for saving lives in the rare circumstance where a potentially deadly fire is raging. To return to Point 1 above, sounding the alarm anytime there is even the *possibility* of a fire or removing fire alarms from a building are not the only two options. Put differently, if we had the option to install a different alarm system that had the same (high) probability of going off for dangerous fires while simultaneously minimizing the number of needless alarms, wouldn't that second system be preferable?

This final point leads well into the final point described in the following section: in these high-risk scenarios, who gets to decide whether there's a possibility of an extreme outcome and who decides what actions would be effective in preventing it, or at least mitigating the risk?

5 The Dangers of Possibilistic Reasoning

"Forget how 'improbable' any of them are, prepare for as many as are feasible, within the reasonable limits of time, money and resources. At an individual level, store extra water, non-perishable food, bullets (dry powder), cash ("dry powder"). At a societal level,

³e.g. Someone leaving something in the microwave too long.

do everything possible to reduce the spread."⁴

The final point I want to make is also perhaps the most obvious – what distinguishes the kind of “possibilistic” reasoning for which Crane argues from that of flat out fear-mongering? If implemented on a societal level, what’s to stop any given politician from enacting whatever measure(s) s/he deems necessary in order to prevent a supposed threat? It’s difficult to argue, after all, that anything has a “true” probability of zero. Do we really want to hand powerful authorities the keys to whatever extreme measures they like because of some threat they perceive, regardless of how plausible that threat may be?

In discussing the pandemic in particular, it’s impossible to argue about the appropriateness of our early response (or lack thereof) without at least a basic sense of what role government and authorities more generally ought to play.

Suppose, for example, a patient presents at a local hospital with an unusual combination of symptoms. It could be that the patient has a strange presentation of a common ailment; it could be that they represent patient 0 of the next pandemic. Should we close local businesses? Shut down the city? Or should we take a more common approach, which might include wearing masks and gloves around that patient and keeping them in isolation until we get more information.

I suspect a common response to this might be something like, “*ok, but in the case of COVID-19, there were many more signs of a possible pandemic. It’s not fair to compare the two situations.*” But in reality, we *must* compare those situations. Without question, both are indications of a possible upcoming pandemic. But to argue that one is more likely than the other to be the start of such a pandemic is to implicitly assign probabilities.

Finally, along similar lines, who should get to decide what actions should be taken to reduce the spread? Do we need definitive evidence that those actions will actually reduce the spread, or is mere scientific speculation enough to warrant forced action? Should we take “high-risk” actions? This would seemingly be implied by Crane’s suggestion of “everything possible”, but I would suggest that these are difficult questions without simple answers. It’s not always obvious whether a particular action will reduce spread or whether those actions potentially present additional dangers.

Consider, for example, the vaccines produced during 2020. Contrary to popular belief, these do not appear to have been difficult to create. Moderna, whose vaccine was just recently approved in mid-December 2020, claims to have developed that vaccine in just two days almost a year prior in January 2020 (Neilson 2020). Given our experience in manufacturing effective vaccines, it seems reasonable to have thought at the time that this vaccine held substantial promise at reducing the spread. Should we have immediately begun a mass vaccination campaign before the virus even reached the shores of the US? Before you roll your eyes at this suggestion, remember that hundreds of thousands of people died waiting for its approval.

In reality, we have no choice but to weigh the costs and benefits of taking various actions, even in high-risk settings. We need an effective means of evaluating risk vs reward and probabilities – even if taken as only loose estimates, should serve an important role in deciding which actions are most appropriate. The all-in, “assume the worst regardless of how plausible it may be”, “*do everything possible*” approach is almost never appropriate.

⁴Again, readers should refer to the disclaimer at the end of Section 1. The points below were written prior to my conversation with Harry and I think they’re important points to make, and so I’m electing to leave them in.

6 Discussion

In dealing with the fallout from COVID-19, it's certainly worthwhile to sit back and evaluate how we got here. Countless baffling mistakes were made early on by numerous political and scientific authorities. I certainly agree with Crane that waiting to take any action whatsoever until more data comes in is certainly not the best approach when you're staring down the very real possibility of a pandemic. If you hear a train coming, you needn't try to calculate the exact probability that it will kill you before stepping off the tracks. But the near-opposite approach that Crane (2020) seems to advocate of taking any and all preventative measures – regardless of how extreme they may be relative to the plausibility of the threat – isn't the best approach either.

Final Thoughts & Acknowledgements

I'm very grateful that Harry Crane took the time to write his initial paper on this topic; I think it's great that statisticians are stepping in and offering up their thoughts, opinions, and perspectives. Along those lines, I'd like to thank both Harry Crane and Ryan Martin for creating a platform like researchers.one where this kind of open debate and discussion is not only possible, but welcomed.

References

- Harry Crane (2020). Naive Probabilism *Researchers.One*, <https://researchers.one/articles/20.03.00003v1>
- John 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, 2020, <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/>
- Cass R. Sunstein (2020). The Cognitive Bias That Makes Us Panic About Coronavirus: Feeling anxious? Blame “probability neglect.” *Bloomberg Opinion*. February 28, 2020, <https://www.bloomberg.com/opinion/articles/2020-02-28/coronavirus-panic-caused-by-probability-neglect>
- Susie Neilson, Andrew Dunn, and Aria Bendix (2020). Moderna's groundbreaking coronavirus vaccine was designed in just 2 days. *Business Insider*. December 19, 2020, <https://www.businessinsider.com/moderna-designed-coronavirus-vaccine-in-2-days-2020-11>