

Hallucinations... Real And Artificial, And What They Mean For Investing And AI Today

By Vineer Bhansali | October 15th, 2026

The following article was published [here](#) on forbes.com

Last week I finished my longest ultramarathon so far – the 220-mile aptly named “The Mammoth”, in the Eastern Sierras ([here](#)). Held at an average altitude of about 8,000 feet, this “run” passes through some of the most beautiful places in California, including two climbs to the summit of Mammoth mountain (11,053 feet), a favorite Southern California skiing destination in the winter. Since a good portion of the course is very steep (with total altitude gain of some 31,000 feet), and the rest soft and pumice covered on rolling fire roads, almost everyone (including the winners) have to hike a good 70-80% of the “run”. With my older knees, I knew I would have to power hike over 90% of the terrain, which would make covering the course very slow (~four days and nights), which in turn would mean less time to sleep, which would mean...sleep deprivation induced hallucinations. I signed up for this event earlier this year as a birthday present to myself, romanticizing the time to unplug and just think for what would become four days and four nights of non-stop forward movement. I knew sleep deprivation would be an issue, but little did I anticipate what this would teach me about interpretation of fact and fiction works in our brains; and why being grounded in reality is a human superpower so important in investing and portfolio management.

Warning: It is well established that sleep deprivation can be very detrimental for health, which is why world records for sleep deprivation are no longer accepted by the Guinness World Records (for a fascinating history read [this](#)). But there are also some lessons for the current era of AI driven investing that I only realized, after only four days of minimal sleep. Now that the experiment is complete, I don't think there will be a repeat for me.


The first night was relatively uneventful. I had budgeted two hours of sleep to get ahead of the sleep deprivation. However, the body, all jacked up from movement, was not able to relax enough to fall asleep. As I started to move the next morning, the sunlight and views kept me wide awake and moving as I got the boost from the normal circadian rhythm. By the second night and 40 hours in, the exhaustion started to build, which, surprisingly, makes sleeping harder, not easier. By the way, in order to prioritize sleep I also sacrificed caffeine intake. As the second 12-hour night wore on, I attempted a couple of 10-15 minute “dirt-naps” on the side of the trail in an emergency bivy. But the cold and the remote location made even this hard to pull off. By the third night I had covered some 140 miles (averaging a little over 50 miles a day, or two marathons each day for a total of about six

back to back so far). A long, monotonous section that climbs a few thousand feet in the cold Sierra night made perfect conditions for moving while “micro-sleeping”, where one goes in and out of sleep while still moving. This is when the brush started to take magical digital properties reflected by my headlamp, and rocks and pebbles started to take the form of objects of clothing with clearly written, but meaningless text on them. I could read the text written on the stones but not interpret the meaning. Not unlike what current, state-of-the-art LLMs do when they predict the next word that should appear in a string by statistical association. Coherent sounding but not always relevant or meaningful. More on this below. Other runners off in the distance seemed like they were taking fast elevators up and down on the trails as their headlamps lit up the undulating trails, and there were distinct buildings that would disappear as I got close to them. And yes, there was even some music, I think... Unlike dreams, one remembers hallucinations clearly since they are in a partially wakeful state. I only realized that I was hiking up and down in small bursts of microsleep, with the narrow mission to keep moving forward guiding my actions, when I walked into a cold creek knee deep that literally woke me up. Wet and cold, I got about 30 minutes of real sleep at the next aid station and when the sunlight arrived, voila, the hallucinations were gone! Sleep was the reset the hallucinating mind needed. Coherence was returned from mere rest and enabling the thinking mind to be back in charge.

(For the record, I finished the race in 144th place, and since finishing was really the only goal for me, I was happy to be in the 70% of starters who finished. It would have been worse to go through all the hallucinations etc. and not get to the finish line, since that would have meant having to come back again next year...no thanks!)

When I researched the causes of mental hallucinations, one “theory” appealed to me that I will now explain. We can think of the brain as weighting “real” information that comes from the surroundings, along with the predictions it is making for what comes next. This is because prediction and minimizing surprise relative to the prediction seems to be the mechanism by which we ensure survival. Predicting and verifying with data is known as the “Bayesian Brain” concept, where the brain starts with a “prior” hypothesis based on its accumulated experience, and then updates this with a “likelihood” function based on the facts that it is observing. Through this updating it converts the prior into a “posterior”, which is the data driven updated belief distribution. The weighting factor between reality and internally generated predictions are inversely proportional to the noise in the signal. The inherent noisiness of the predictions it makes and confirms or discards makes hallucinations non-existent when the brain is well rested and “ground truth” like sunlight can help in creating concrete images on the brain. This is because in the bright light and with a clear mind, the uncertainty in what we are observing is negligible relative to the noise in the imaginary signals. This is probably the reason why the hallucinations tended to disappear after even a small amount of sleep and when it was bright. On the other hand, when sleep deprived and in the dark, there was no ground truth from concrete visual

images, so the noisy signals generated by the brain itself got a relatively high weighting in making the brain believe that it was seeing something consistent -- in other words the made-up predictions are matched by made up visual images to create consistency and minimize surprise. The posterior was basically “made-up” with erroneous or non-existent data. This is when the hallucinations took the place of reality. The simple solution to making the hallucinations and fake reality go away? Take a nap and get some light into the brain to drown out the random noise by a dose of reality! Fortunately, humans have a way to ground predictions with empirical data to make sure that they don’t run away with their self-consistent but imaginary facts.



AI rendition of sagebrush with random nonsense text

This episode gave me some insight into the challenges of the increasing use of artificial intelligence models becoming popular today in investing. I have three recent episodes where highly sophisticated AI models have presumably hallucinated and made up completely cogent and internally self-consistent facts that turned out to be false. The first example is where I was working on the weighting of a duration neutral yield curve trade and the AI driven model confidently provided all the analysis, but with one fatal error that was deeply embedded in the results and only prior domain knowledge could uncover it, where it used \$100,000 as contract size for each CBOT 2-year futures contract, rather than the correct \$200,000 per contract. A factor of two mistake that it did not catch until prompted. In the second case, our trading team asked for a quote on the unwind of a cleared interest

rate swap, and the trading bot of the market maker (a very well-known player) unknowingly switched signs where they effectively posted a gain to us instead of a loss (which we promptly pointed out to them to fix and correct as this was the ethical thing to do). In the third case, relayed to us by a client, a well-known AI system confidently predicted that the interest rate futures markets were predicting three Fed rate “cuts” instead of hikes, a sign flip error. It goes without saying that each of these three examples can be potentially catastrophic for a trader who puts real money at risk. In each of the three cases, when the error was pointed out, the AI system readily admitted that it had made a factual error, but it took a human to point it out and turn its confident prediction on its head. Of course, I wish we could just make the AI system take a nap and correct its hallucinations upon awakening, but unfortunately that approach does not seem to work, yet.

Given the inability of many (most) AI systems being used today in investment management don't have a “ground truth” validation mechanism, my opinion on the usefulness of AI tools is quite simple – while AI might be able to compute with speed and efficiency, for the time being the judgment and decision making will still heavily rely on humans who have a reference point with which to judge right from wrong so as to avoid catastrophic errors. When the euphoria around efficiency and speed that AI has enabled wears off, the realization is likely to arise that in domains where accuracy can mean life or death, or large profits and losses, human judgement and quality control remains still remains the gold standard because of the ability to fact check and correct course.

I am more or less recovered from the self-induced bout of hallucinations from my four-day sleep deprivation experiment in the Sierras, but I am more than ever convinced that when the dust settles, AI will never fully replace humans where expert skills and domain knowledge are of paramount importance. In the meantime, blind use of sloppy AI will create massive opportunities for those who are able to realize the blunders that can arise from not having a quality control mechanisms in place. And I hope is that during this period, AI hallucinations don't drive financial markets off the proverbial cliff.

Important Disclosures

Vineer Bhansali, Ph.D. is the Founder and Chief Investment Officer of LongTail Alpha, LLC, an SEC-registered investment adviser and a CFTC-registered CTA and CPO. Any opinions or views expressed by Dr. Bhansali are solely those of Dr. Bhansali and do not necessarily reflect the opinions or views of LongTail Alpha, LLC or any of its affiliates (collectively, "LongTail Alpha"), or any other associated persons of LongTail Alpha. You should not treat any opinion expressed by Dr. Bhansali as investment advice or as a recommendation to make an investment in any particular investment strategy or investment product. Dr. Bhansali's opinions and commentaries are based upon information he considers credible, but which may not constitute research by LongTail Alpha. Dr. Bhansali does not warrant the completeness or accuracy of the information upon which his opinions or commentaries are based.

This publication is for illustrative and informational purposes only and does not represent an offer or solicitation with respect to the purchase or sale of any particular security, strategy or investment product. Past performance is not indicative of future results.

Different types of investments involve varying degrees of risk, including possible loss of the principal amount invested. Therefore, it should not be assumed that future performance of any specific investment or investment strategy, or any non-investment related content, will be profitable or prove successful. Nothing contained herein is intended to predict the performance of any investment.