

The Complex Markets Hypothesis

In which I hypothesise that markets are subjective, uncertain, complex, stochastic, adaptive, fractal, reflexive ... - really any clever sounding adjective you like – just not efficient.

Allen Farrington

@allenf32 / allen-f@hotmail.co.uk



photo by [skeeze](#), via Pixabay

Around a month ago, Nic Carter asked me to have a look at a final draft of [his article on the basics of the Efficient Markets Hypothesis](#). Dancing around the edges of Bitcoin Twitter as I am prone to do, I immediately grasped both the need for and the point of such an article; the question of whether the upcoming ‘halving’ is ‘priced in’ or not had “*become a source of great rancor and debate,*” as Nic wrote. For the uninitiated, ‘the halving’ is the reduction of the bitcoin block reward from 12.5 bitcoin to 6.25, expected around May 2020. Nic set himself the task of explaining the EMH more or less from scratch, in such a way that the explanation would naturally lend itself towards insight on questions of Bitcoin’s market behaviour.

I think he did a great job and the article is well worth reading. But I couldn’t help thinking as I went through it that, basically, I didn’t believe this stuff the first time around, and it all seemed strangely incongruous in a setting explicitly involving Bitcoin, what with the tendency of serious thinkers in this space to treat highly mathematised mainstream / neoclassical financial economics with something between suspicion and disdain.

To be completely clear, this is in no way a ‘rebuttal’ to Nic. He articulated the EMH very well, but didn’t defend it. That wasn’t the point of his article at all. He watered down the presentation at several points by saying (quite helpful) things like:

“I do not believe in the ‘strong form’ of the EMH. No finance professional I know does. It is generally a straw man,”

and,

“Interestingly, by caveating the EMH, we have stumbled on an alternative conception entirely. The model I have described here somewhat resembles Andrew Lo’s adaptive market hypothesis. Indeed, while I am very happy to maintain that most (liquid) markets are efficient, most of the time, the adaptive market model far more closely captures my views on the markets than any of the generic EMH formulations.”

One passage in particular stuck out to me:

“Referring to it as a model makes it very clear that it’s just an abstraction of the world, a description of the way markets should (and generally do) work, but by no means an iron law. It’s just a useful way to think about markets.”

This is where I’m not so sure. Yes, it’s an abstraction, and no, it’s not an iron law. But I don’t think it’s a terribly good abstraction, and I think the reason is that it subtly contradicts and elides what *are*, in fact, iron laws, or as close to iron laws as can be found in economics. It’s a useful way to think about markets, to a point, but I want to explore what I think is a more useful way.

My argument will go through the following propositions, which serve as headings for their own sub-sections of discussion: value is subjective; uncertainty is not risk; economic complexity resists equilibria; markets aggregate prices, not information; and, markets tend to leverage efficiency.

I will conclude with some additional commentary on Andrew Lo’s *Adaptive Markets Hypothesis* and Benoit Mandelbrot’s interpretation of fractal geometry in financial markets, simply because, of all the reading around this topic that was thrown up by Nic’s article, these two were by far the most intriguing. I didn’t want to do either an injustice by bending their arguments too far to make them fit my own, but I think that they can be very fruitfully analysed with the conceptual tools we will have developed by the conclusion of the essay. I will also occasionally invoke the concepts of ‘reflexivity’, as articulated by George Soros in [The Alchemy of Finance](#), and several concepts popularised and articulated by Nassim Taleb, such as ‘skin in the game’ and ‘robustness’.

This might seem like an excessive coverage list just to offer a counter to the claim that markets are ‘efficient’ – which seems pretty reasonable in and of itself. If it is at all reassuring to the reader before diving in, I don’t think my thesis has five intimidating-sounding propositions, so much as one quite simple idea, from which many related propositions can be shown to follow. I think that, fundamentally, the efficient markets hypothesis is contradicted by the implications of value being subjective, and that some basic elements of complex systems are helpful, in places, to nudge the reasoning along. This essay is an attempt to tease these implications out.

Value is Subjective

You shouldn’t compare apples and oranges, except that sometimes you have to, like when you are hungry. If apples and oranges are the same price, you need to make a decision that simply cannot be mathematised. You either like apples more than oranges, or vice versa. And actually, even this may not be true. Maybe you know full well you like oranges, but you just feel like an apple today, or you need apples for a pie recipe for which oranges would be *très gauche*. This reasoning is readily extended in all directions; which is *objectively* better, a novel by Dickens or Austen? A hardback or an ebook by either, or anybody? And what about the higher order capital goods that go into producing apples, oranges, novels, Kindles, and the like? Clearly they are ‘worth’ only whatever their buyer subjectively

assesses as likely to be a worthwhile investment given the (again) subjective valuations of others as to the worth of apples, oranges, novels, and whatnot ...

This is all fine and dandy; readily understood since the marginal revolution of Menger, Jevons, and Walras in the 1870s rigorously refuted cost and labour theories of value. As Menger put it in his magisterial *Principles of Economics*,

“Value is thus nothing inherent in goods, no property of them, nor an independent thing existing by itself. It is a judgment economizing men make about the importance of the goods at their disposal for the maintenance of their lives and well-being. Hence value does not exist outside the consciousness of men.”

Fair enough. But the first seductive trappings of the EMH come from the rarely articulated assumption that such essential subjectivity is erased in financial markets because the goods in the market are defined only in terms of cash flows. There may not be a scientific answer as to whether apples are better than oranges, but surely \$10 is better than \$5? And surely \$10 now is better than \$10 in the future? But what about \$5 now or \$10 in the future?

There are (at least) two reasons this reductionism is misleading. The first comes from the mainstream neoclassical treatment of temporal discounting, which is to assume that only exponential discounting can possibly be “optimal”. The widespread prevalence of alternative approaches – hyperbolic discounting, for example - is then usually treated via behavioural economics, as a deviation from optimality that is evidence of irrational cognitive biases.

This has been challenged by a recently published preprint paper by Alex Adamou, Yonatan Berman, Diomides Mavroyiannis, and Ole Peters, entitled *The Microfoundations of Discounting* ([arXiv link here](#)) arguing that the single assumption of an individual aiming to optimise the growth rate of her wealth can generate different discounting regimes that are optimal relative to the conditions by which her wealth grows in the first place. This in turn rests on the relationship between her current wealth and the payments that may be received. Sometimes this the discounting that pops out is exponential, sometimes hyperbolic, sometimes something else entirely. It depends on her circumstances.

I would editorialise here that an underlying cause of confusion is that people value time itself, and, naturally, do so subjectively. It may be fair enough to say that they typically want to use their time as efficiently as possible - or grow their wealth the fastest - but this is rather vacuous in isolation. Padding it out with circumstantial information immediately runs into the fact that everybody’s circumstances are different. As Adamou said on Twitter shortly after the paper’s first release, *not many 90-year olds play the stock market*. It’s funny because it’s true.

And it is easy to see how this result can be used as a wedge to pry open a conceptual can of worms. In financial markets, there are far more variables to compare than just the discount rate – and if we can’t even assess an objective discount rate, we really are in trouble! In choosing between financial assets we are choosing between non-deterministic streams of future cash flows, as well as (maybe – who knows?) desiring to preserve some initial capital value.

Assume these cash flows are ‘risky’, in the sense that we can assign probabilities to their space of outcomes. In the following section, we will see that really the cash flows are not ‘risky’, but ‘uncertain’, which makes this problem even worse – but we can stick with ‘risky’ for now as it works well enough to make the point. There can be no objective answer because different market participants could easily have different risk preferences, exposure preferences, liquidity needs, timeframes, and so on.

Timeframes are worth dwelling on for a second longer (there’s ‘time’ again) because this points to an ill-definition in my hasty setup of the problem: to what space of outcomes are we assigning probabilities, exactly? Financial markets do not have an end-point, so this makes no sense on the face of it. If we amend it by suggesting (obviously ludicrously) that the probabilities are well-defined for

every interval's end-point, *forever*, then we invite the obvious criticism that different participants may care about different sequences of intervals. Particularly if their different discount rates (which we admitted they must have) have a different effect on how far in the future cash flows have to come to be discounted back to a value that is negligible in the present. Once again, people value *time itself subjectively*.

In the readily understood language employed just above, market participants almost certainly have different *circumstances* to one another, from which different subjective valuations will naturally emerge. What seems to you like a stupidly low price at which to sell an asset might be ideal for the seller because they are facing a margin call elsewhere in their portfolio (see Nic's cited example of what blew up LTCM despite it being a 'rational bet'), or because they hold too much of this asset for their liking and want to rebalance their exposure. Or perhaps some price might seem stupidly high to buy, but the buyer has a funding gap so large that they need to invest in something that has a non-zero probability of appreciating by that much. If you *need* to double your money, then the 'risk-free asset' is infinitely risky. There is no right answer, because value is subjective.

Uncertainty is not Risk

'Risk' characterises a nondeterministic system for which the space of possible outcomes can be assigned probabilities. Expected values are meaningful and hence prices, if they exist in such a system, lend themselves to effective hedging. 'Uncertainty' characterises a nondeterministic system for which probabilities *cannot* be assigned to the space of outcomes. Uncertain outcomes cannot be hedged. This distinction in economics is usually credited to Frank Knight and his wonderful 1921 book, [*Risk, Uncertainty, and Profit*](#). In the introduction, Knight writes,

"It will appear that a measurable uncertainty, or "risk" proper, as we shall use the term, is so far different from an unmeasurable one that it is not in effect an uncertainty at all. We shall accordingly restrict the term "uncertainty" to cases of the non-quantitative type. It is this "true" uncertainty, and not risk, as has been argued, which forms the basis of a valid theory of profit and accounts for the divergence between actual and theoretical competition."

Keynes is often also credited an excellent exposition,

"By "uncertain" knowledge, let me explain, I do not mean merely to distinguish what is known for certain from what is only probable. The game of roulette is not subject, in this sense, to uncertainty... Or, again, the expectation of life is only slightly uncertain. Even the weather is only moderately uncertain. The sense in which I am using the term is that in which the prospect of a European war is uncertain, or the price of copper and the rate of interest twenty years hence, or the obsolescence of a new invention, or the position of private wealth owners in the social system in 1970. About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know. Nevertheless, the necessity for action and for decision compels us as practical men to do our best to overlook this awkward fact and to behave exactly as we should if we had behind us a good Benthamite calculation of a series of prospective advantages and disadvantages, each multiplied by its appropriate probability, waiting to be summed."

The conclusion of the Keynes passage is particularly insightful as it gets at why it is so important to be clear on the difference, which otherwise might seem like little more than semantics: people need to act. They will strive for a basis to treat uncertainty as if it were risk so as to tackle it more easily, but however successful they are or are not, they must act nonetheless.

The Knight extract hints at the direction of the book's argument, which I will summarise here: that profit is the essence of competitive uncertainty. Were there no uncertainty, but merely quantifiable risk in patterns of production and consumption, competition would drive all prices to a stable and commoditised equilibrium. In financial vocabulary, we would say there would be no such thing as a sustainable competitive advantage. The cost of capital would be the risk-free rate, as would all returns on capital, meaning profit is minimised. In aggregate, profit would function merely as a kind of force pulling all economic activity to this precise point of strong attraction.

But of course, uncertainty is very real, as Keynes' quote makes delightfully clear. I would argue, in fact, that in the economic realm it is a direct consequence of subjective value; in engaging in pursuing profit, you are guessing what others will value. As Knight later writes,

“With uncertainty present, doing things, the actual execution of activity, becomes in a real sense a secondary part of life; the primary problem or function is deciding what to do and how to do it.”

So far I have danced around the key word and concept here, so as to try to let the reader arrive at it herself, but this ‘deciding what to do and how to do it’, and ‘pursuing profit’, we call *entrepreneurship*. In a world with uncertainty, the role of the entrepreneur is to shoulder the uncertainty of untried combinations of capital, the success of which will ultimately be dependent on the subjective valuations of others. This is not something that can be calculated or mathematised, as any entrepreneur (or VC) will tell you. As Ross Emmett noted in his [centennial review of Risk, Uncertainty, and Profit](#), it is no coincidence that the word ‘judgment’ appears on average every two pages in the book.

There are two points about the *process* of entrepreneurship that I believe ought to be explored further, and which lead us to Soros and Taleb: you can't just *imagine* starting a business; you have to actually do it in order to learn anything. And, in order to do it, you have to expose yourself to your own successes and failures. Your experiment changes the system in which you are experimenting, and you will inevitably have a stake in the experiment's result.

This is fertile ground in which to plant Soros' theory of reflexivity. As briefly as possible, and certainly not doing it justice, Soros believes that financial markets are fundamentally resistant to truly scientific analysis because they can only be fully understood in such a way that acknowledges the fact that thinking about the system influences the system. He writes that the scientific method:

“is clearly not applicable to reflexive situations because even if all the observable facts are identical, the prevailing views of the participants are liable to be different when an experiment is repeated. The very fact that an experiment has been conducted is liable to change the perceptions of the participants. Yet, without testing, generalisations cannot be falsified.”

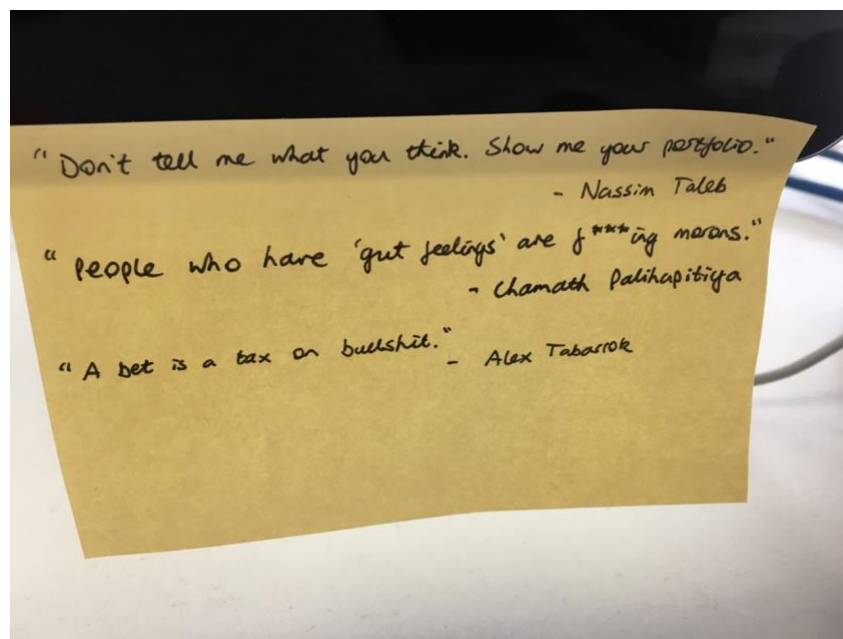
All potential entrepreneurial activity is uncertain (by definition) but the fact of engaging in it crystallises the knowledge of its success or failure. The subjective valuations on which its success depends are revealed by the experiment, and you can't repeat the experiment pretending you don't now know this information. Alternatively, this can be conceived of in terms of the difference between thinking and acting, or talking and doing. In a reflexive environment, you can't say *what would have happened had you done something*, because, had you done it, you would have changed the circumstances that lead to you now claiming you would have done it. As Yogi Berra (allegedly) said, *“in theory there is no difference between theory and practice, but in practice, there is.”*¹

We can also now invoke ‘skin in the game’, a [phrase of dubious origin](#), but nowadays associated primarily with Nassim Taleb, and expounded in his 2017 [book of the same name](#). Again, not doing it justice (he did write a whole book about this) Taleb believes that people ought to have

¹ and as Amy Adams [vigorously proclaims](#) in *Talladega Nights*, quite the treatise on risk and uncertainty by the way, and with a [criminally underrated soundtrack](#), *“Ricky Bobby is not a thinker. Ricky Bobby is a driver.”*

equal exposure to the potential upsides and downsides of their decisions; ‘ought to’ in both a moral sense of deserving the outcome, but also in the sense of optimal system design, in that such an arrangement encourages people to behave the most prudently out of all possible incentive schemes. It readily applies here in that braving the wild uncertainties of entrepreneurship requires capital – it requires a stake on which the entrepreneur *might* get the upside of profit, but *might* get the downside of loss. I say ‘might’ because you cannot possibly know the odds of such a wager. It relies not on risk, but on uncertainty. As Taleb writes, “*entrepreneurs are heroes in our society. They fail for the rest of us.*”

The combined appreciation of ‘judgment’ and ‘skin in the game’ is key to understanding what entrepreneurs are actually *doing*. They do not merely throw capital into a combinatorial vacuum; they are intuiting the wants and needs of potential customers. And as I simply cannot resist the opportunity to employ, possibly my favourite quote from any economist, ever: as Alex Tabarrok says, *a bet is a tax on bullshit. Or, don’t talk; do.*



my desk at work. It's good to keep these things in mind.

That same Emmett review of Knight’s book noted that the very concept of *Knightian uncertainty* re-emerged in the public consciousness around a decade ago due to two events: the role ironically played by financial risk instruments in the financial crisis, which neoclassical economists had up until that point insisted would reduce uncertainty in markets (search for “Raghuram Rajan Jackson Hole” if unfamiliar); and Taleb publishing the bestseller [The Black Swan](#).² Writes Emmett,

“Taleb did not suggest that uncertainty could be handled by risk markets. Instead, he made a very Knightian argument: since you cannot protect yourself entirely against uncertainty, you should build robustness into your personal life, your company, your economic theory, and even the institutions of

² Although, naturally, Taleb deplores the concept of ‘Knightian Uncertainty’. What I believe Taleb truly objects to, however, is how the concept has come to be used, rather than anything Knight himself believed. Economists often invoke ‘Knightian Uncertainty’ as a sleight of hand to demarcate some corner of reality, and imply that everywhere else is merely ‘risky’ and can be modelled. This is nonsense. In real life, everything is uncertain, or, as [Joseph Walker](#) succinctly put, “*Taleb’s problem with Knightian uncertainty is that there’s no such thing as non-Knightian uncertainty.*” I think Knight would almost certainly agree, as would anybody who has actually *read* Knight, instead of employing his name in the course of *macro-bullshitting*, as Taleb would put it.

your society, to withstand uncertainty and avoid tragic results. These actions imply costs that may limit other aspects of your business, and even your openness to new opportunity.”

But enough about entrepreneurship, what about financial markets? Well, financial markets are readily understood as one degree removed from entrepreneurship. With adequate mental flexibility, you can think of them as markets for fractions of entrepreneurial activity. Entrepreneurship-by-proxy, we might say. If you want, you can use them to mimic the uncertainty profile of an entrepreneur: your ‘portfolio’ could be 100% the equity of the company you wish you founded. Or 200%, with leverage, if you are really gung-ho! But most people think precisely the opposite way: markets present the opportunity to tame the rabid uncertainty of entrepreneurship in isolation, and skim some portion of its aggregate benefit.

There is an additional complication. The fact of such markets usually being liquid enough to enable widespread ownership creates the incentive to think not about the underlying entrepreneurship at all, but only about the expectations of other market participants – to ignore the fundamentals and consider only the valuation. There are shades of Soros’ reflexivity here. The market depends to some extent on the thinking of those participating in the market *about* the market. This is sometimes called a Keynesian beauty contest, after Keynes’ analogy of judging a beauty contest not on the basis of who you think is most beautiful, but on the basis of who you think others will think is the most beautiful. But if everybody is doing that, then you really need to judge on the basis of who you think others will think others will think is the most beautiful, and so on. Unlike the entrepreneur, who must only worry about the subjective valuation of his potential customers, participants in financial markets must worry, in addition, about the subjective valuation *of this subjective valuation* by other market participants. There is often grumbling at this point that this represents ‘speculation’ as opposed to ‘investment’, and I certainly buy the idea that over time periods long enough to reflect real economic activity allowed for by the investments, such concerns will make less and less of a difference. As Benjamin Graham famously said, “*in the short run, the market is a voting machine, but in the long run, it is a weighing machine*”. But the voting still happens. It is clearly real and needs to be accounted for. Risk is once again useless. Uncertainty abounds.

This range of possibilities is intriguing and points to a deeper understanding of what financial markets really *are*: the aim of a great deal of finance is to grapple with the totality of uncertainty inherent in entrepreneurial activity – equally well understood as ‘investment of capital’, given the need for a ‘stake’ – by partitioning it into different exposures that can sensibly be described as *relatively* more or less ‘risky’. The aim of doing so is generally to minimise the cost of capital going towards real investment by tailoring the packaging of uncertainty to the ‘risk profiles’ of those willing to invest, as balanced by escalating transaction costs if this process becomes too fine-grained.

This is the essence of a capital structure: the more senior the capital claim, the better defined the probability space of outcomes for that instrument. Uncertainty in aggregate cannot be altered, nor can its influence be completely removed from individual instruments, but exposure to uncertainty can be unevenly parcelled out amongst instruments.

This suggests a far more sophisticated understanding of ‘the risk/reward trade-off’ and ‘the equity premium’ than is generally accepted in the realm of modern portfolio theory, and, by extension, the EMH: bonds are likely to get a lower return than stocks not because they are less ‘risky’ (which in that context is even more questionably interpreted as ‘less volatile’) but because they are engineered to be less uncertain. The burden of uncertainty is deliberately shifted from debt to equity. You don’t get a ‘higher reward’ for taking on the ‘risk/volatility’ of equities; you deliberately expose yourself to the uncertain *possibility* of a greater reward in exchange for accepting an uncertain possibility of a greater loss.

It is worth pondering for a second that this is arguably why the ‘equity risk premium’ even exists (and why neoclassical economists are so confused about it, while financial professionals are not in the slightest) – if there really were no uncertainty in investment and every enterprise – and hence

every financial instrument linked to it – had a calculable risk profile, then price discrepancies derivable from expectation values could be arbitrated away. There would be no equity risk premium – nor a risk premium of any kind on any asset. Everything would be priced correctly and volatility would be zero. That volatility is *never* zero clearly invalidates this idea. I suggest that the distinction between risk and uncertainty provides at least part of the explanation: unless, by remarkable coincidence, every market participant’s opportunity costs (of exposure, liquidity, time, etc.) and perception of uncertainty (of fundamentals, other’ perceptions of fundamentals, others’ perception of others’ perceptions, etc.) is all identical, and remains so over a period of time, price-altering trading will occur.

(I will note in passing that this commentary is merely intended to provide the intuition that something is amiss with the ‘equity premium puzzle’. It is incomplete as an explanation. In the later section on leverage efficiency, I will cover Peters’ and Adamou’s more formal proof of the puzzle’s non-puzzliness.)

An important concept to appreciate in the context of uncertainty is that of ‘heuristics’. This is an important loose-end to tie up before moving on from uncertainty, along with one more, ‘randomness and unpredictability’, which I cover shortly. This is quite a simple idea that originates with Herbert Simon and has been taken up with force more recently by Gerd Gigerenzer of the Max Planck Institute for Human Development and more obliquely by Taleb. Simon’s framing began by assuming that individuals do not in fact have perfect information, nor the resources to compute perfectly optimal decisions. Given these constraints, Simon proposed that individuals demonstrate *bounded rationality*; they will be as rational as they can given the information and resources they actually have. This probably sounds straightforward enough – perhaps tautological – but notice it flies in the face of behavioural economics, which tends to cover for neoclassical economics by saying, effectively, that since information and competition are perfect, risk is always defined and the optimal decision can always be calculated, but the reason people don’t do so is that they are hopelessly irrational. I have always thought this is quite silly on the face of it, but it is clearly also seductive. Anybody reading the likes of *Thinking, Fast and Slow* immediately gets the intellectual rush of thinking everybody is stupid except him.

Bounded rationality encourages the development of ‘heuristics’, which the reader may recall behavioural economists railing against. A heuristic is effectively a rule of thumb for dealing with an uncertain environment that you are pretty sure will work even if you can’t explain why, precisely. The classic example is that of a dog and a frisbee, or an outfielder in baseball catching a flyball: the outfielder *could* solve enough differential equations to calculate the spot the ball will land, but the dog certainly can’t. And it turns out that neither do: in real life, they adjust their running speed and direction such that the angle at which they see the frisbee or ball stays constant. And it works. No equations required. Yippee.³

The implied simplicity of heuristics has subtle mathematical importance, also. A more technical way of specifying this is to say that they have very few parameters – discrete, independent information inputs to the decision procedure - ideally they could even have zero. In a purely risky environment (if such a thing exists, which, in real life, it almost certainly does not) a decision procedure ought to have *as many* parameters as are needed to capture the underlying probability distribution. But the more uncertainty you add to such an environment, the more dangerous this becomes, essentially because what you are doing is fine-tuning your model to an environment that simply no longer exists.

³ Farrington’s Heuristic is another good example, which I made up right now in preparing version 1.2 – if a writer is discussing risk, uncertainty, knowledge, and the like, if he refers to Gödel’s Incompleteness Theorems, and if he is not obviously joking, then everything else he says can immediately be dismissed because he is a bullshitting charlatan enamoured by [cargo cult math](#). This heuristic has only one binary parameter – ‘is he joking?’ – and so is highly robust.

In case anybody cares, the theorems are **NOT** about ‘knowledge’: they are about provability within first order formal logical theories strong enough to model the arithmetic of the natural numbers. This is quite a specific mathematical thing that bears no relation whatsoever to epistemology or metaphysics. Also, there are two of them, which turns out to be important if you understand what the first one says.

Eventually you will get an unforeseen fluctuation so large that your overfitted model gives you a truly awful suggestion. Heuristics are *robust* to such circumstances in light of having very few parameters to begin with. Think back to the outfielder: imagine he solves all the necessary fluid dynamical equations, taking account of the fly ball's mass, velocity, and rotation, the air's viscosity, the turbulence generated, and so on. If there is then a gust of wind, he's screwed. His calculation will be completely wrong. But if he embraces the heuristic of *just looking at the damn ball* this won't matter!

Interested readers are encouraged to peruse Gigerenzer's recent(ish) work on the use of [heuristics in finance](#), and [their abuse in behavioural economics](#), which recently got a [shout out in Bloomberg](#), or [this video](#) which is a great introduction to Gigerenzer's ideas, as well as their connection to Taleb's more informal thinking on the same topic. (also, note the number of times Gigerenzer uses the word 'complex'. It is no coincidence that this is *a lot*, as we shall shortly see):

The video linked to is around an hour and a half, so the reader need not take such a detour now, but I would encourage it at some point, as both Gigerenzer and Taleb are excellent. My favourite excerpt comes around the 19-minute mark, when Gigerenzer recalls that Harry Markowitz – considered the founder of modern portfolio theory – didn't actually use any Nobel-prize winning modern portfolio theory for his own retirement portfolio; he used the zero-parameter $1/n$ approach. If one were being especially mean-spirited, one might say that he didn't want his own bullshit to be taxed. And as it turns out, in order for the Markowitz many-many-many parameter approach to investing to consistently outperform $1/n$, you would need around 500 years of data to finetune the parameters. Of course, you also need the market to *not change at all* in that time. Good luck with that.

Since markets feature multitudes of interrelated uncertainties, it is reasonable to expect participants to interact with them not with the perfect rationality of provably optimal behaviour, but with the bounded rationality of heuristics, which are selected on the basis of judgment, intuition, creativity, etc. Basically, people mostly are not stupid. And if they are, they have skin in the game, so they get punished, and possibly wiped out.

A kind of nice, conceptual corollary to 'risk is not uncertainty' is, 'unpredictability is not randomness'. There can be unpredictable events that are not random, and randomness that is not unpredictable. The difference essentially comes down to 'causation'. Think of Keynes' example of the obsolescence of a new invention. This is 'unpredictable' not because it is subject to an extremely complicated probability density function, but because the path of causation that would lead to such a situation involves too much uncertainty to coherently grasp. Or think of the bitcoin mining process. The time series of the first non-zero character in the hash of every block is certifiably random, but it is not unpredictably random. It is the result of a highly coordinated and purposeful effort. It doesn't spring forth from beta decay. Because we understand the causal process by which this time series emerges, we can predict this randomness very effectively.

A key building block of the EMH is the 'random walk hypothesis': the idea that you can 'prove' using statistical methods that stock prices follow 'random walks' – a kind of well-defined and genuinely random mathematical behaviour. But you can do no such thing. You can prove that they are *indistinguishable* from random walks, but that is really just saying you can use a statistical test to prove that some data can pass a statistical test. If you *understand* what *causes* price movements, you will arrive at no such nonsense as claiming that the moves are, themselves, random. They very probably *look* random because they are fundamentally unpredictable from the data. And they are fundamentally unpredictable from the data because they derive from the incalculable interplay of millions of market participants' subjective assessments of the at-root uncertain process of entrepreneurship.

None of this is based on randomness, nor 'risk', nor 'luck'. It is based on the unknown and unknowable profit that results from intuiting the results of untried and unrepeatable experiments and backing one's intuition with skin in the game.

Before moving on, I think it is worth tying all of this to where it is more tangibly sensible, lest the reader not quite know what to do with it all. A big deal was made recently about Netflix being by far the best performing US mid-to-large-cap stock of the 2010s. Netflix is useful as an example because

of the scale of its success, but note the following argument does not depend on scale at all. While you could craft an explanation as complicated as you like, I think saying, *streaming is better than cable*, pretty much does it, once added to all the circumstantial factors to do with the competitive and technological environment. Now imagine an investor in 2010 whose thesis was that streaming is better than cable and would likely win in the long run, who surveyed the competitive environment, and decided Netflix would be a good investment. Is their outperformance over the next 10 years ‘luck’? Was all the ‘information’ ‘in the price’ in 2010? Would the CAPM tell you what the price *should* have been? Did the stock go for a nice little random walk to the moon?

This is clearly an insane interpretation. Consider the alternative: The investor better intuited the subjective values of future consumers than did the average market participant. Very likely she justified this on the basis of a heuristic or two. She staked capital on this bet - which was not risky and random, but uncertain and unpredictable - and exposed herself to a payoff that turned out to be huge, *because she was right!* To the peddlers of the EMH, rational expectations, perfect information, and the like, this obviously sensible interpretation is utterly heretical.

Economic Complexity Resists Equilibria

The link between profit and entrepreneurship can be tugged at ever-so-slightly further, and invites a brief detour into the basics of complex systems. The argument goes more or less as follows: the discussion on uncertainty needn’t be interpreted as a call to abandon mathematical analysis altogether – just the sloppy mathematics of risk and randomness that has effectively no connection to the real world. There is an alternative mathematical approach, however, which directly addresses and contradicts the standard neoclassical formalism.

The starting point is Israel Kirzner, widely considered one of the foremost scholars of entrepreneurship, and his book, *Competition and Entrepreneurship*. One of Kirzner’s theses is a positive argument that has roughly two parts, as follows: first, entrepreneurship is by its nature non-exclusionary. It is a price discrepancy between the costs of available factors of production and the revenues to be gained by employing them in a particular way – or, profit. In other words, it is perfectly competitive. It does not rely on any privileged position with respect to access to assets; The assets are presumed to be available on the market. They are just not yet employed in that way, but could be, with capital that is presumably homogeneous. Anybody could do so. The only barrier is that of the willingness to judge and stake on uncertainty. He writes,

“The entrepreneur’s activity is essentially competitive. And thus competition is inherent in the nature of the entrepreneurial market process. Or, to put it the other way around, entrepreneurship is inherent in the competitive market process.”

This notion of what ‘competition’ really means is highly antithetical to the neoclassical usage. In fact, it is more or less the exact opposite. Rather than meaning something like, *tending towards abnormal profit and hence away from equilibrium*, the neoclassicals mean, *tending towards equilibrium and hence away from abnormal profit*. Kirzner bemoans this,

“Clearly, if a state of affairs is to be labelled competitive, and if this label is to bear any relation to the layman’s use of the term, the term must mean either a state of affairs from which competitive activity (in the layman’s sense) is to be expected or a state of affairs that is the consequence of competitive activity ... [Yet] competition, to the equilibrium price theorist, turned out to refer to a state of affairs into which so many competing participants have already entered that no room remains for additional entry (or other modification of existing market conditions). The most unfortunate aspect of this use of the term ‘competition’; is of course that, by referring to the situation in which no room

remains for further steps in the competitive market process, the word has come to be understood as the very opposite of the kind of activity of which that process consists. Thus, as we shall discover, any real-world departure from equilibrium conditions came to be stamped as the opposite of ‘competitive’ and hence, by simple extension, as actually ‘monopolistic’.”

I’d note in passing the delightful similarity in the concluding thought of this extract to the argument of Peter Thiel’s [Zero to One](#), considered by many a kind of spiritual bible for – you guessed it – *entrepreneurship*. Anyway ...

Kirzner’s second positive argument is that correcting this conceptual blunder leads one to realise that a realistic description of competitive markets would be not as constantly at equilibrium, but rather as constantly *out of equilibrium*. And that’s really all we need to move on to complex systems.

Complex systems are commonly associated with the Santa Fe Institute, and popularised by W. Mitchell Waldrop’s fantastic popular science book, [Complexity](#). Waldrop focuses, for the most part, on one of the SFI’s first ever workshops, held between a group of physicists and economists in 1987. The proceedings of the workshop are fantastic, have aged very well, and seem to your author cheap relative to his subjective valuation of them at [~\\$70 in paperback or ~\\$140 in hardback](#). My thinking here comes from the very first paper of the workshop, W. Brian Arthur’s now somewhat infamous work on increasing returns. To get a sense of what I mean by ‘infamous’, consider the following from Waldrop:

“Arthur had convinced himself that increasing returns pointed the way to the future for economics, a future in which he and his colleagues would work alongside the physicists and the biologists to understand the messiness, the upheaval, and the spontaneous self-organisation of the world. He’d convinced himself that increasing returns could be the foundation for a new and very different kind of economic science.

Unfortunately, however, he hadn’t much luck convincing anybody else. Outside of his immediate circle at Stanford, most economists thought his ideas were – strange. Journal editors were telling him that this increasing-returns stuff ‘wasn’t economics.’ In seminars, a good fraction of the audience reacted with outrage: how dare he suggest that the economy was not in equilibrium!”

Readers can probably sense where this is going.

Arthur’s paper at the workshop, *Self-Reinforcing Mechanisms in Economics*, is a breath of fresh air if you have ever slogged through the incessant [cargo cult math](#) of neoclassical financial economics (as I had to in researching this essay – thanks a lot, Nic!) It is frankly just all so sensible! Okay, so there are a few differential equations, but only after ten pages of things that are obviously true, and only to frame the obviously true observations in the absurd formalism of the mainstream.

To begin with, *“conventional economic theory is built largely on the assumption of diminishing returns on the margin (local negative feedbacks); and so it may seem that positive feedback, increasing-returns-on-the-margin mechanisms ought to be rare.”* Standard neoclassical theory assumes competition pushes all into equilibrium, from which a deviation is punished by the negative feedback of reduced profits. So far, so good.

“Self-reinforcement goes under different labels in these different parts of economics: increasing returns; cumulative causation; deviation-amplifying mutual causal processes; virtuous and vicious circles; threshold effects; and non-convexity. The sources vary. But usually self-reinforcing mechanisms are variants of or derive from four generic sources: large set-up or fixed costs (which give the advantage of falling unit costs to increased output); learning effects (which act to improve products or lower their cost as their prevalence increases); coordination effects (which confer

advantages to ‘going along’ with other economics agents taking similar action); and adaptive expectations (where increased prevalence on the market enhances beliefs of further prevalence).”

Now we are getting into the meat of it. An example or two wouldn't hurt before applying this to entrepreneurship and markets.

Arthur likes Betamax versus VHS – which is a particularly good example in hindsight because we know that VHS won despite being mildly technologically inferior. Point number 1: If a manufacturer of VHS tapes spends an enormous amount on the biggest VHS (or Betamax) factory in the world, then the marginal costs of producing VHS will be lower from that point on. Even if the factory as a whole is loss making, the costs are sunk, and so the incentive is to pump out VHS by the gallon. The fact that this can be done so cheaply makes consumers more likely to choose VHS over Betamax, which will in turn justify the initial expense and contribute positive feedback (via profit).

Point number 2: doing so may give the owner of the factory the experience to learn how to do so even more efficiently in the future. By the same eventual mechanism as above, this contributes positive feedback via lower prices. (interested readers are encouraged to look into ‘Wright’s Law’, in particular a [recent paper](#) by Béla Nagy, Doyne Farmer, Quan Bui, and Jessika Trancik, which basically says that Moore’s Law happens for everything, just slower; or, we learn by doing)

Points number 3 and 4: if more people seem to be buying VHS tapes than Betamax, then producers of Betamax *players* are incentivised to shift production towards VHS players instead. Cheaper VHS players incentivise consumers to buy more VHS *tapes*. The *appearance* of VHS winning this battle causes economic agents to adapt their behaviour in such a way that makes VHS more likely to *actually* win. In glancing over an early draft of this essay, Nic kindly pointed out to me that this represents the dominant philosophy behind growth VC from 2015 until WeWorkGate, as if a bunch of zealous, born-again Arthurians were playing a [game of non-iterated prisoner’s dilemma with other people’s money](#). Anyway ...

Arthur writes, *“if Betamax and its rival VHS compete, a small lead in market share gained by one of the technologies may enhance its competitive position and help it further increase its lead. There is positive feedback. If both systems start out at the same time, market shares may fluctuate at the outset, as external circumstances and ‘luck’ change, and as backers manoeuvre for advantage. And if the self-reinforcing mechanism is strong enough, eventually one of the two technologies may accumulate enough advantage to take 100% of the market. Notice however we cannot say in advance which one this will be.”*

While Arthur mostly considers realistic examples in economics which have discrete end-states that are then ‘locked into’, such as settling on VHS over Betamax, or Silicon Valley over Massachusetts Route 128, my contention would be that every one of these features describes a part of the process of entrepreneurial competition. The fact of staking capital at all towards an uncertain end represents a fixed cost which must be matched by competitors, and after which unit costs fall. As we have mentioned several times, entrepreneurs learn from the result of their experiments and improve their own processes. There is a clear coordination effect for customers in the default assumption of doing whatever other customers are doing. And adaptive expectations are likewise fairly straightforwardly applied: we tend to assume that businesses will continue to exist and that we can continue to act as their customers. Businesses tend to assume the same of their customers within reasonable bounds of caution. The specific positive feedback as a result of each individual effect is that of ‘profit’ – it is positive in the sense that it can be reinvested in the enterprise and allow it to grow.

Of course, it is possible that these effects would diminish and the marginal feedback become negative. But what we are more tangibly proposing here is that any once-existing competitive advantage has been completely eroded away. This only happens when the product itself becomes either obsolete in light of a superior competitor, or completely commoditised. The former is simply more of the same at the macro level, but the latter we can in turn explain by uncertainty becoming so minimal

that we can more or less safely assume it is merely risk. Such circumstances are few and far between. Uncertainty is prevalent in all aspects of economic life, as we have discussed. My argument here is that, so, therefore, are increasing returns and positive feedback loops.

To bring in Arthur one last time:

“if self-reinforcement is not offset by countervailing forces, local positive feedbacks are present. In turn, these imply that deviations from certain states are amplified. These states are therefore unstable. If the vector-field associated with the system is smooth and if its critical points – its ‘equilibria’ – lie in the interior of some manifold, standard Poincaré-index topological arguments imply the existence of other critical points or cycles that are stable, or attractors. In this case multiple equilibria must occur. Of course, there is no reason that the number of these should be small. Schelling gives the practical example of people seating themselves in an auditorium, each with the desire to sit beside others. Here the number of steady-states or ‘equilibria’ would be combinatorial.”

Recall there is no way to know from the starting point which steady-state will be settled into. And of course, Arthur is only talking about specific economic circumstances, not the aggregate of all economic behaviour. The aggregate will likely have shades of evolution in a competitive environment (another concept we will soon encounter in more detail): many, many such interdependent sub-systems, always moving towards their own steady state, but almost all never getting there. And so, in summary, there is a solid mathematical basis to saying that economic behaviour in aggregate is wildly uncertain.

Before moving on, I just want to mention that Arthur should almost certainly be better known and respected in Bitcoin circles. Readers uninterested in the connection I am proposing between Bitcoin and complex systems (or unimpressed by my amateur passion for both) can skip ahead without missing anything. Arthur’s 2013 paper, [Complexity Economics](#), is an excellent place to start. Likewise, a good argument can be made that complex systems researchers should be a lot more interested in Bitcoin. Readers may well have picked up on the essence of Arthur’s analysis consisting of ‘network effects’. I avoided using the term because Arthur himself doesn’t use it. But he is considered the pioneer of their analysis in economics, and when you think about it, the concept of ‘increasing returns’ makes perfect sense in the context of a network. What greater competitive advantage can you have than everybody needing to use your product simply because enough people already use it? And what product do people need to use solely because others are using it more than ‘money’?

Although I have eschewed the idea of ‘lock-in’ as helpful for the analysis above, Bitcoin surely has amongst the strongest interdependent network effects of any economic phenomenon in history? Is it not a naturally interdisciplinary complex adaptive system *par excellence*? Is it not a form of artificial life, coevolved with economising humans in the ecology of the Internet? I mean, for goodness’ sake, Andreas Antonopoulos claims to have put ants on the cover of [Mastering Bitcoin](#) because,

“the highly intelligent and sophisticated behaviour exhibited by a multimillion-member ant colony is an emergent property from the interaction of the individuals in a social network. Nature demonstrates that decentralised systems can be resilient and can produce emergent complexity and incredible sophistication without the need for a central authority, hierarchy, or complex parts.”

Back in the SFI workshop, Arthur writes,

“When a nonlinear physical system finds itself occupying a local minimum of a potential function, ‘exit’ to a neighbouring minimum requires sufficient influx of energy to overcome the ‘potential barrier’ that separates the minima. There are parallels to such phase-locking, and to the difficulties

of exit, in self-reinforcing economic systems. Self-reinforcement, almost by definition, means that a particular equilibrium is locked in to a degree measurable by the minimum cost to effect changeover to an alternative equilibrium.”

I’m not sure anybody can sensibly describe what such a ‘minimum cost’ would be. Particularly because Bitcoin is set up in such a way that any move away from lock-in by one metric causes a disproportionate pull back to lock-in by another. It’s Schelling points all the way down.

Markets Aggregate Prices, Not Information

The most frustrating thing about the EMH for me is that even the framing is nonsensical. You don’t really need to get into subjective value, uncertainty, complex systems, and so on, to realise that in reading the proposition, *prices reflect all available information*, you have already been hoodwinked (*hoodwunk?*). What does ‘reflect’ mean?

Nic dramatically improved upon this by saying that markets *aggregate information*. I noticed this is typical of many more enlightened critiques of EMH, and it serves as a far better starting point, in at least *suggesting* a mechanism by which the mysterious link between information and price might be instantiated. Unfortunately, I think the mechanism suggested is simply invalid. It is not realistic at all and it implicitly encourages a dramatic misunderstanding of what prices really are and where they come from.

In making sense of this we have to assume some kind of ‘function’ from the space of information to price. I think it’s acceptable to mean this metaphorically, without implying the quasi-metaphysical existence of some such force. We might really mean something like, *the market behaves as if operating according to such and such a function*. Adam Smith’s ‘invisible hand’ is an instructive comparison. For the time being, I will talk as if some such function ‘exists’.

We can maybe imagine information as existing as a vector in an incredibly high-dimensional space, at least as compared to price, which is clearly one-dimensional. We could even account for the multitudes of uncertainty we have already learned to accept by suggesting that each individual’s subjective understanding of all the relevant factors and/or ignorance of many of them constitutes a unique mapping of this space to itself, such that the ‘true information vector’ is transformed into something more personal for each market participant. Perhaps individuals then bring this personal information vector to the market, and what the market does is *aggregate* all the vectors by finding the average. Finally, the market *projects* this n-dimensional average vector onto the single dimension of price. If you accept the metaphorical nature of all these functions, I can admit this model has some intuitive appeal, in the vein of James Surowiecki’s [The Wisdom of Crowds](#).

The problem is that this is clearly not how anybody actually interacts with markets. You don’t submit your n-dimensional information/intention-vector; you submit your one-dimensional price. That’s it. The market aggregates these one-dimensional price submissions in real time by matching the flow of marginal bids and asks.

This understanding [gets two birds stoned at once](#). First, it captures the mechanics of how we know price discovery in markets *actually works*. There is no mysterious, market-wide canonical projection function – no inexplicable ‘prices reflect information’ – there are just prices, volumes, and the continuous move towards clearing.

Second, it implies a perfectly satisfactory and not at all mysterious source of the projection of information into price: individuals who make judgments and act. Any supposedly relevant ‘information’ is subject both to opportunity cost and uncertainty. Individuals alone know the importance of their opportunity costs, and individuals alone engage with uncertainty with heuristics, judgment, and staking. If individuals are wrong, they learn. If they are very wrong, they are wiped out. Effective heuristics live to fight another day.

I am genuinely surprised that this confusion continues to exist in the realm of the EMH, given that, as far as I am concerned, Hayek cleared it up in its entirety in [The Use of Knowledge in Society](#). A superficial reading of Hayek's ingenious essay might lead one to believe something like *prices reflect information*. But, to anachronistically borrow our function metaphor once more, Hayek points out that the projection from the n-dimensions of information to the one dimension of price *destroys an enormous amount of information*. Which is the whole point! Individuals are incapable of understanding *the entirety of information in the world*. Even the entirety of individuals is incapable of this. Thanks to the existence of markets, nobody has to. They need only know about prices. 'Perfect information' is once again shown to be an absurdity. Of the 'man on the spot', whom we might hope would make a sensible decision about resource allocation,

"There is hardly anything that happens anywhere in the world that might not have an effect on the decision he ought to make. But he need not know of these events as such, nor of all their effects. It does not matter for him why at the particular moment more screws of one size than of another are wanted, why paper bags are more readily available than canvas bags, or why skilled labor, or particular machine tools, have for the moment become more difficult to obtain. All that is significant for him is how much more or less difficult to procure they have become compared with other things with which he is also concerned, or how much more or less urgently wanted are the alternative things he produces or uses. It is always a question of the relative importance of the particular things with which he is concerned, and the causes which alter their relative importance are of no interest to him beyond the effect on those concrete things of his own environment."

Hayek proposes this be resolved by the price mechanism:

"Fundamentally, in a system in which the knowledge of the relevant facts is dispersed among many people, prices can act to coordinate the separate actions of different people in the same way as subjective values help the individual to coordinate the parts of his plan."

Perhaps ironically, this points to the only sensible way in which markets *can* be called 'efficient'. They are efficient with respect to the information they manipulate and convey: as a one-dimensional price, it is the absolute minimum required for participants to interpret and sensibly respond. Markets have excellent [social scalability](#); they are the original *distributed systems*, around long before anybody thought to coin that expression.

Interestingly, this meshes very nicely with the complex systems approach to economics associated with Arthur at SFI, and perhaps more specifically with John Holland. His paper at the aforementioned inaugural economics workshop, [The Global Economy as an Adaptive Process](#), at seven pages and zero equations, is well worth a read. Holland recounts many, now familiar, difficulties in mathematical analysis of economics that assume linearity, exclusively negative feedback loops, equilibria, and so on, before proposing that 'the economy' is best thought of as what he calls an 'adaptive nonlinear network'. Its features are worth exploring, even if they require some translation:

"Each rule in a classifier system is assigned a strength that reflects its usefulness in the context of other active rules. When a rule's conditions are satisfied, it competes with other satisfied rules for activation. The stronger the rule, the more likely it is to be activated. This procedure assures that a rule's influence is affected by both its relevance (the satisfied condition) and its confirmation (the strength). Usually many, but not all, of the rules satisfied will be activated. It is in this sense that a rule serves as a hypothesis competing with alternative hypotheses. Because of the competition there are no internal consistency requirements on the system; the system can tolerate a multitude of models with contradictory implications."

We could easily translate ‘rule’ as ‘entrepreneurial plan’ or something similar. Entrepreneurial plans can contradict one another, clearly – if they are bidding on the same resources for a novel combination – and can and do compete with one another. Clearly, such plans are hypotheses about the result of an experiment that hasn’t been run yet. Holland then says,

“A rule’s strength is supposed to reflect the extent to which it has been confirmed as a hypothesis. This, of course, it’s a matter of experience, and subject to revision. In classifier systems, this revision of strength is carried out by the bucket-brigade credit assignment algorithm. Under the bucket-brigade algorithm, a rule actually bids a part of its strength in competing for activation. If the rule wins the competition, it must pay this bid to the rules sending the messages that satisfied its condition (its suppliers). It thus pays for the right to post its message. The rule will regain what it has paid only if there are other rules that in turn bid and pay for its message (its consumers). In effect, each rule is a middleman in a complex economy, and it will only increase its strength if it turns a profit.”

Much of this does not need translating at all: we see Menger’s higher orders of capital goods, and value of intermediate goods resting ultimately with the subjective value of consumers, who pass information up the chain of production. We see agents that learn from their experience. We see skin-in-the-game of staked capital in ‘bidding part of its strength’ and we see uncertain gain or reward ultimately realised by profit or loss. But most importantly – most *Hayekily* – we see agents who have no such fiction as ‘perfect information’, but rather responding solely to prices in their immediate environment, and whose reactions affect prices that are passed to other environments. In *Complexity*, Waldrop quotes Holland’s frustration with the neoclassical obsession with well-defined mathematical problems:

“‘Evolution doesn’t care whether problems are well-defined or not.’ Adaptive agents are just responding to a reward, he pointed out. They don’t have to make assumptions about where the reward is coming from. In fact, that was the whole point of his classifier systems. Algorithmically speaking, these systems were defined with all the rigor you could ask for. And yet they could operate in an environment that was not well defined at all. Since the classifier rules were only hypotheses about the world, not ‘facts’ they could be mutually contradictory. Moreover, because the system was always testing those hypotheses to find out which ones were useful and led to rewards, it could continue to learn even in the face of crummy, incomplete information – and even while the environment was changing in unexpected ways.

‘But its behaviour isn’t optimal!’ the economists complain, having convinced themselves that a rational agent is one who optimises his ‘utility function’.

‘Optimal relative to what?’ Holland replied. Talk about your ill-defined criterion: in any real environment, the space of possibilities is so huge that there is no way an agent can find the optimum – or even recognise it. and that’s before you take into account the fact that the environment might be changing in unforeseen ways.”

Hayek gives us the intuition of prices conveying only what market participants deem to be the most important information and actually destroying the rest, and Holland shows how this can be represented with the formalism of complex systems. But note that the EMH forces us to imagine that the information is somehow *in the market itself*. It is honestly unclear to me whether the EMH even allows for honest or ‘rational’ disagreement given it implies that the price is ‘correct’, and all other trading is allegedly ‘noise’. By my account (and Hayek’s) people can clearly disagree. That’s why they trade in the first place; they value the same thing differently. This is not at all mysterious if we realise that engaging with markets requires individuals to ‘project’ the n-dimensions of their information,

heuristics, judgments, and stakes onto the single dimension of price, and that markets do not project the aggregates; they aggregate the projections.

Markets Tend to Leverage Efficiency

So we know that entrepreneurial efforts will tend towards positive feedback loops if successful, which is a fancy way of saying, they will ‘grow’. And we know that the diversity of compounding uncertainty in markets for securities linked to these efforts will likely generate substantial volatility. But can we say anything more? Can we expect anything more precise?

It turns out that we can, and here we finally get to Alex Adamou, Ole Peters, and the ergodicity economics research program. It’s about time! The goal of the program is to trace the repercussions of a conceptual and algebraic error regarding the proper treatment of ‘time’ in calculations of ‘expectation value’ that pervaded mainstream economics over the course of the twentieth century. Interested readers are encouraged to visit the [program’s website](#), check out [this recent primer in Nature Physics](#), or just follow Ole and Alex on Twitter, which is where most of the action seems to happen anyway!

First, a down to earth [example](#). Imagine you want a pair of shoes. You can either go to the same shoe store every day for a month, or you can do to every shoe store in town all in one trip. If it turns out there is no difference between these approaches, this system is ‘ergodic’. If, as seems more likely, there is a difference, the system is ‘non-ergodic’.

Now with more technical detail, the *conceptual and algebraic error* is as follows: imagine some variable that changes over time, subject to some well-defined randomness. Now imagine a system of many such variables, whose ‘value’ is just the sum of all the values of the variables. Now imagine you want to find the ‘average’ value of a variable in this system in some pure, undefined sense.

How do you make sense of an ‘average’ of a system that will be different every time you run it? Well, you could fix the period of time the system runs for, and take the limit of where individual variables get to attained by running the system over and over and over to infinity. Or, you could fix the number of systems (preferably at ‘one’ for minimal confusion) and take the limit of where individual variables get to attained by running the system further and further into the future, to infinity.

These are called, respectively, the ‘ensemble average’ and the ‘time average’, and are easily remembered as the average achieved by taking x to infinity. ‘Ensemble average’ is commonly known as ‘the expectation’ but Peters and Adamou resist this terminology because it has nothing whatsoever to do with the English word ‘expectation’. You shouldn’t necessarily *expect* the expectation.

Now these values might be the same. This means you can measure one of these even if what you really want is the other. If so, your system is called ‘ergodic’. The concept first developed within nineteenth century physics when Ludwig Boltzmann wanted to justify using ensemble averages to model macroscopic quantities such as pressure and temperature in fluids, which are strictly speaking better understood as time averages over bajillions of classically mechanical collisions. If any regular readers of mine exist, they will remember me going through much of this in my essay, [Cargo Cult Math](#).

My point that time around was to go on to say that a great deal of financial modelling uses techniques – most notably expectation values – which would only be appropriate if the corresponding observables were ergodic. But they are not. Almost none of them are, to a degree that is both obvious and scary once you grasp it in its totality: clearly the numbers in finance are causally dependent on one another and take place in a world in which time has a direction.

My point this time around is more cheerful. I want to direct the reader’s attention to another of Peters’ and Adamou’s papers on the topic: *Leverage Efficiency* ([arXiv link here](#)). This subsection is a whistle-stop tour of what that paper says. The usual disclaimer about not doing it justice absolutely applies. The reader is heartily encouraged to read the paper too.

Imagine a toy model of the price of a stock that obeys geometric Brownian motion with constant drift and with volatility that varies by random draws from a normal distribution. It turns out that the growth rate of the ensemble average price – i.e. the price averaged over all possible parallel systems – is not the same as the time average growth rate of the price – i.e. the growth rate in a single system taken in the long time limit. Clearly what we care about is the time average, as we don't tend to hold stocks across multiple alternate universes, but rather across time in the actual universe. In particular, it turns out that the ensemble average growth rate is equal to the drift, while the time average growth rate is equal to the drift minus a correction term: the variance over 2.

This becomes very important when we introduce leverage via a riskless asset an investor can hold short. Let's call the model drift of the stock minus the stipulated drift of the non-volatile riskless asset 'the excess growth rate'. Then we can say that the ensemble average growth rate in situations with variable leverage is the growth rate of the riskless asset, plus the leverage multiplied by the excess growth rate. However, the time average has a linked correction, as above. As it is difficult at this point to continue the exposition in English, compare the formulae below:

$$g_e(l) = \mu_{riskless} + l * \mu_{excess}$$

$$g_t(l) = \mu_{riskless} + l * \mu_{excess} - \frac{(l * \sigma)^2}{2}$$

The relevance of the difference is that the latter formula is not monotonic in l . In other words, you don't increase your growth rate unboundedly by leveraging up more and more. This might seem intuitively obvious, and, in fact, the intuition likely strikes in exactly the right spot: in reality there is volatility. The more and more levered you are, the more susceptible you are to total wipeout for smaller and smaller swings. In fact, we can go further and observe that we can therefore maximise the growth rate as a function of leverage, implying an objectively optimal leverage for this toy stock:

$$l_{opt} = \frac{\mu_{excess}}{\sigma^2}$$

What might this optimal leverage be in practice? Well, Peters and Adamou propose the tantalising alternative to the EMH: the *stochastic markets hypothesis*. As opposed to the EMH's *price efficiency*, they propose *leverage efficiency*: it is impossible for a market participant without privileged information to beat the market by applying leverage. In other words, real markets self-organise such that the optimal leverage of 1 is an attractive point for their stochastic properties.

The paper continues in two directions: firstly, Peters and Adamou propose a theoretical argument for feedback systems that ought to be triggered over long enough periods of the theoretical value of optimal leverage in fact deviating from 1, which all ought to pull it back to 1. I will skip this as it is tangential to the point I am building towards, although obviously very interesting in its own right. Secondly, Peters and Adamou gather data from real markets to establish what the optimal leverage would, in fact, have been. I include some screenshots that strongly suggest this approach is quite fruitful:





Figure 1: Return-leverage curves for the S&P500. Total return from a constant-leverage investment in the S&P500, starting 4th August 1955 and ending 10th March 2017, as a function of leverage. Data analyses 1 (red), 2 (yellow), 3 (green), and 4 (blue). For descriptions of the computations, see text, section 4.2.



This chart probably deserves some explanation, but is very satisfying once grasped: as opposed to just the S&P500, above, both the German equity market (DAX) and Bitcoin show pretty much identical behaviours to the S&P500, which Peters and Adamou term “satisfying leverage efficiency”. That the Madoff curve is so different, and seems to have no clear maximum, indicates it is likely too good to be true. This is a nice result given that we know Madoff’s returns to be fraudulent!

However, what I really want to get to in all of this is a specific interpretation of the SMH; that markets self-organise such that optimal leverage tends to 1 in the long run. If we assume that the excess return of the stock price is generated by real economic activity (ultimately, the consistency of the stock’s return on equity) in the long enough run, this would seem to suggest that a certain amount of volatility is actually natural. Were a stock to consistently generate an excess return above that of the riskless asset, investors would lever up to purchase it. This mass act would (reflexively!) cause its volatility to shoot up as the price shoots up, and in the inevitable case of a margin call on these levered investors, volatility would increase further as the stock price comes back down.

This is a somewhat naïve explanation, but the gist is that the lack of volatility in the short run will tend to generate excess volatility in the medium run, such that a natural level is tended to in the long run. Or, markets are stochastically efficient. Readers familiar with the unsuspecting role that ‘portfolio insurance’ turned out to have played in 1987’s Black Monday - the single biggest daily stock

market drop in modern history that seemed to follow no negative news whatsoever - will find all this eerily familiar. Taleb calls Black Monday a prototypical *Black Swan* that shaped his formative years as a Wall Street trader. Mandelbrot cites it as sure-fire evidence of power laws and wild randomness in financial markets. It ties together many themes of this essay because, evidently, the information was not in any of the prices. Not in the slightest. I leave it to the reader to mull over what all this implies if interventions in financial markets are targeted solely at reducing volatility as a worthwhile end in itself, the rationale of which makes no mention of growth or leverage. Once again, search for “Raghuram Rajan Jackson Hole” or read about the so-called [Great Moderation](#) if unsure where to start. Volatility signals stability, in financial markets and [likely well beyond ...](#)

All this has a final interesting implication that I teased earlier: the resolution of the so-called ‘equity premium puzzle’; that, according to such-and-such behavioural models from the psychological literature, the excess return of equities ‘should be’ much lower than it really is. Cue the behavioural economists claims of irrational risk aversion, blah blah blah. Peters and Adamou provide an alternative with no reference to human behaviour at all. The difference between the growth rates of the risky ($l=1$) and riskless ($l=0$) assets is the excess return minus the volatility correction. If markets are attracted to the point at which leverage efficiency equals 1, then it follows by substituting the definition of the equity risk premium in terms of risky and riskless assets into the equation defining optimal leverage, that the equity premium ought to be attracted to the excess return over 2. Peters and Adamou delightfully write, “*our analysis reveals this to be a very accurate prediction ... we regard the consistency of the observed equity premium with the leverage efficiency hypothesis to be a resolution of the equity premium puzzle.*” QED.

I’ll note before wrapping up this sub-section that any readers triggered by such terms as *geometric Brownian motion* and *normal distributions* needn’t be. Peters and Adamou acknowledge that GBM is not realistically either necessary or sufficient as a mechanism for stock price movements. But their argument really only depends on the characteristics of an upward drift and random volatility, both of which *are* reasonable to expect. They choose GBM because it is simple to handle, well understood, and prevalent in the literature they criticise, but they also write that,

“for any time-window that includes both positive and negative daily excess returns, regardless of their distribution, a well-defined optimal constant leverage exists in our computations ...

Stability arguments, which do not depend on the specific distribution of returns and go beyond the model of geometric Brownian motion, led us to the quantitative prediction that on sufficiently long time scales real optimal leverage is attracted to $0 \leq l_{opt} \leq 1$ (or, in the strong form of our hypothesis, to $l_{opt} = 1$).”

We knew from previous sections that volatility is *likely*. It will exist to some extent due to the teased-out implications of subjective values and omnipresent uncertainty. But now we know that it is *necessary*. It is not [noise](#), irrationality, panic, etc, around a correct price. It is, at least in part, inevitable reflexive rebalancing of leverage around whatever the price happens to be.

Incompleteness

You can’t write ten thousand words on mathematical formalisms outlining the limits of human knowledge without mentioning Gödel’s Incompleteness Theorems.

Adaptation and Fractals

As mentioned in the introduction, of all the dissenting work on the EMH, I most recommend, by far, Andrew Lo's *Adaptive Markets Hypothesis* – the [original paper](#) and the [follow-up book](#) – and the various thoughts of Benoit Mandelbrot on fractals in financial markets – strewn across numerous academic papers, but lucidly conveyed in the popular book, *The (Mis)behaviour of Markets*. I assume familiarity with these works to avoid explaining everything from scratch, so if the reader is unfamiliar, I encourage jumping ahead to the next section.

My main critique of Lo is that he doesn't take uncertainty seriously enough. In covering the academic history surrounding the EMH, he only gives Simon a page or so, and Gigerenzer a paragraph. The key point of failure, in my view, is his treatment of the Ellsberg paradox. Or rather, the fact that he stops his rigorous discussion of uncertainty at this point.

The problem here is that the uncertainty in the Ellsberg paradox is confined to the odds, whereas we know from the previous discussion that the uncertainty in economics exists in the outcomes. This means that the odds aren't just uncertain, they are non-existent. By stopping here, Lo passes off the results of running the experiment that gives rise to the so-called paradox as simply indicating ambiguity aversion, which he presents as a kind of irrational bias – then a segue to behavioural economics. This prevents Lo from exploring the implications of Knightian uncertainty on entrepreneurship and competition, and ultimately gives him little ammunition to take on the EMH directly. In fact, he acknowledges that he never really does – he just proposes something he thinks is better.

That said, I agree that his model is better. Far better! Adaptation is a fascinating concept to employ here. As noted several times, it comes through very naturally in the complex systems approach. I won't comment on it too much as its roots in evolutionary biology are outside my academic pedigree. But the basic intuition of changing circumstances and responding agents I find rather compelling. As do, it would seem, several thinkers I have already cited. Consider this passage from Kirzner,

“it is necessary to introduce the insight that men learn from their experiences in the market. It is necessary to postulate that out of the mistakes which led market participants to choose less-than-optimal courses of action yesterday, there can be expected to develop systematic changes in expectations concerning ends and means that can generate corresponding alterations in plans.”

Also, there is a tradition of referring to heuristics as *ecologically rational*, and the biological analogy is no coincidence. This passage from Waldrop's *Complexity* on John Holland's conversion to complex systems thinking in his study of genetics is striking in the almost simple obviousness of the comparison drawn to economics (again, not at all a coincidence):

“it bothered Holland that [R.A.] Fisher kept talking about evolution achieving a stable equilibrium – that state in which a given species has attained its optimum size, its optimum sharpness of tooth, its optimum fitness to survive and reproduce. Fisher's argument was essentially the same one that economist use to define economic equilibrium: once a species' fitness is at a maximum, he said, any mutation will lower the fitness ... but that did not sound like evolution to me.

... to Holland, evolution and learning seemed much more like – well, a game. In both cases, trying to win enough of what it needed to keep going. In evolution that payoff is literally survival, and a chance for the agent to pass its genes on to the next generation. In learning, the payoff is a reward of some kind, such as food, a pleasant sensation, or emotional fulfilment, but either way, the payoff (or lack of it) gives agents the feedback they need to improve their performance: if they're going to be 'adaptive' at all, they somehow have to keep the strategies that pay off well, and let the others die out.”

One thing I especially like about Lo's approach is his idea of 'evolution at the speed of thought', often rhetorical as much as anything else. I think this provides a useful conceptual tool to

deal with what I deemed to be the only consistent deficiency in the material I covered on complex systems: Arthur, Holland, et al, seem to me so focused on the comparison to biological evolution, and on shifting the comparative conceptual framework from physics to biology as a whole, that they forget the role of purposeful human beings in all of this. Economic ‘mutation’ is not random, it is creative, intuitive, judgmental. It happens at the speed of thought because humans think *on purpose*. They do not cycle through the space of every thought that can possibly be had until they hit on one that happens to be a business plan.

To put this in a wider context and loop back to previously cited thinkers, I think Arthur is best read alongside Kirzner, and indeed Kirzner is best read alongside Arthur. Particularly in [The Nature of Technology](#), which is otherwise an excellent book, Arthur perfectly grasps *how* change happens, but not *why*. In *Competition and Entrepreneurship*, Kirzner perfectly grasps *why* change happens, but not *how*. Both the why and the how rely, in part, on understanding economic evolution as an essentially human phenomenon, because genes mutate, but humans *think*.

Mandelbrot’s ideas on fractals in finance are iconoclastic, to say the least. Unlike Lo, I see nothing to disagree with, and much that probably went over my head. But given Mandelbrot sets himself the task of demolishing the EMH out of left field, and seemingly succeeds, it is definitely worth grappling with.

The mildly boring part of *The (Mis)behaviour of Markets* is Mandelbrot showing that, empirically, financial data does not seem to fit the Brownian motion of the random walk hypothesis, and hence the EMH. The really juicy part is his explanation of *why*. To avoid getting into any really tricky mathematics and essentially rewriting his book, I will summarise his argument, again not doing it justice, as, *this isn’t random enough*. More suggestively, *it is too predictably random*.

Mandelbrot thinks that prices in financial markets are, up to a certain granularity, fractals. If true, this has many fascinating implications, but the most relevant here is that the self-similarity this implies means that any randomness in their fluctuation must be irregular. It should not be possible to ascertain any regularity just by changing the timespan because they look the same on *every* timespan (look! there’s ‘time’ again!) The randomness must itself be pretty random. And that randomness must be random, and so on. There are no genuinely normal distributions in finance, Mandelbrot believes, but rather they all tend towards Cauchy. We could be less hand-wavy about all this and point out instead that while a statistical test on some financial data might suggest the tail of a lognormal distribution, we are really looking at a power law. If parameterized to induce fat enough tails, such a distribution may not have a variance, and if fat enough, not even a mean. (And no, it doesn’t have an ‘infinite variance’ or ‘infinite mean’ because that is meaningless, but nice try). As Taleb and Mandelbrot both [wrote in Fortune](#),

“In bell-curve finance, the chance of big drops is vanishingly small and is thus ignored. The 1987 stock market crash was, according to such models, something that could happen only once in several billion billion years.”

Black swans, amiright?

What does this have to do with ‘complex’ markets? Mandelbrot doesn’t explore this idea, and I may be going out on a limb here, but I think this is almost exactly what you would expect if you thought markets were maximally uncertain, so to speak. If risk were predictable, then it could be hedged against. If it were unpredictable in and of itself, but were distributed predictably, then *that* could be hedged against. And so on and so forth. This all lends itself to a hand-wavy inductive proof by contradiction. We know that nothing can be perfectly hedged because it derives from uncertainty, and uncertainty on uncertainty, and uncertainty on *that* uncertainty, and so on. Financial markets can

shift uncertainty around, and selectively parcel it into more and less risky instruments, but uncertainty itself cannot be removed.

Bitcoin

Oh goodness, I guess I have to say something about Bitcoin now, lest I be accused of rickrolling an angry twitter mob into a sermon on armchair economics. Is the halving priced in?

[I have no idea](#). Which is sort of the moral of all of this. You can't predict the uncertain future, but you can bet on it. I'm not sure how you would bet on this exact hypothesis: perhaps a combination of options that pay off if and only if the price goes up (or doesn't go up) by whatever the stock-to-flow model predicts, within some bounds, when it predicts it, within some bounds? Obviously, you could just be long Bitcoin, but then you aren't isolating the essence of this claim, and you can benefit for all sorts of other reasons. If you do either, you'll move the price towards the outcome you are hoping for. But only by having put skin in the game. Also, you could believe something very specific about the halving, but have no way of testing it as you don't believe in stock-to-flow models, or any other valuation model, for that matter. That's the essence of my noncommittal answer above: I shouldn't tell you what I think, I should show you my portfolio, right? Well, I have no 'halving bet' in my portfolio, so I guess I think nothing. Which is what I said 😊

Even so, we can still make a few interesting observations that draw on the above discussion. Clearly, the question relies on reflexivity, which is interesting in and of itself. It's only derivatively a question about the fundamentals, and more about the extent to which the market is a well-oiled beauty contest. I don't think I know enough about the actual workings of the Bitcoin market to comment on this. It strikes me that, relative to global equities markets, at least, the range of heuristics that market participants in Bitcoin are using vary wildly from one another. If it makes any sense to say so, they likely have pretty dramatic variance. At the same time, the market itself is probably highly illiquid, relative to what we might be used to. This might suggest the halving *isn't* priced in, in the sense that the change in marginal bids and asks at which the market clears that we know is going to happen dwarfs the capital that is already deployed, including towards solely this essentially reflexive bet. But then again, maybe it doesn't.

Honestly, I just don't know. And if somebody does claim to know, tell them to show you their portfolio.

Conclusion

Value is subjective, which means uncertainty governs all economic phenomena. This creates a complexity that resists equilibria and is constantly changing besides. Within such a system, prices convey the minimal possible information necessary for economic agents to purposefully react. They do so with judgment and heuristics, not 'perfect information', which is nonsensical, as is 'perfect competition' and 'rational expectations'. Prices may pass statistical tests for randomness, but they are not themselves random (although it is plausible that their randomness is random, and that randomness is random, and so on) but rather are unpredictable on the basis of market data alone. They are, however, predictable to the extent that the predictor accurately assesses the future subjective valuations both of economic agents and fellow market participants, and backs up this prediction with staked capital. This act of staking changes the uncertainties at play, rendering any attempt at genuinely scientific analysis futile. You *can* beat the market, it's just really hard, and it depends on understanding *people*, not data. And it's meaningless if you do it in theory but not practice.

Markets have many characteristics. I suggest they are subjective, uncertain, complex, stochastic, adaptive, fractal, reflexive ... - really any clever sounding adjective you like – just not efficient.

Thanks to Nic Carter, Robert Natzler, Alex Adamou, and Sacha Meyers for edits and contributions.

Version 1.2

I may upload revisions to correct typos, broken links, nonsensical arguments, etc. – please get in touch if you find any.

*Feel free to contribute to the Allen ideas fund:
bc1q8utvneuvn3hf2lm5nvt3dregead6hh5sda2sa*