Everything Was Already AI

Have you ever watched Ant-Man and the Wasp: Quantumania? https://www.youtube.com/watch?v=PWLy1CLUatI 1:40-1:45 “Darren…surprised to see me”

Or Riverdale? https://www.youtube.com/watch?v=dlQDhAPOkT0 [whole thing]

Or Lonerbox reacting to Hasan reacting to Bad Empanada reacting to Willymac reacting to Hasan reacting to the war? https://www.youtube.com/watch?v=tC2KP7bOskc&

Have you been on a dating app and ALL THE PROMPT ANSWERS ARE THE SAME?

You want to debate pineapple on pizza?

Your idea of a good Sunday is a walk and a roast?

The biggest risk you ever took was downloading this app?

Together we could delete this app?

You’ll fall for me if I trip you?

[Dating Apps Men & Women Clips]

Why are the same shops absolutely everywhere? https://www.youtube.com/watch?v=gJW4-cOZt8A& 1:03-1:18 “what the hell…Dave & Busters”

Why are so many politicians so bland, similar, and ineffectual?

https://www.youtube.com/watch?v=Ll3iyvbsRDM 0:00-0:13 “…everything he just said”

Have you ever bemoaned the music played on the radio as unoriginal: “it all sounds the same”?

There’s scientific evidence that it does!

All of these things have the feel of regurgitated slop served up to us as humanity’s lowest common denominator. A processing of social preferences that magnifies and reflects trends back at us, then feeds on itself, producing at worst grotesque hallucinations and at best conventional wisdom. There is a nub of substance at the heart of it but it’s wrapped in a combination of laziness, compromise, and superficiality.

And here’s the thing:

None of these were generated by AI

Many people are worried about whether AI will take over humanity, use all of our energy, and result in massive inequalities. Others are excited that it will usher in an age of abundance, humanity’s escape from drudgerous work, and a sci-fi-esque future.

Clearly, these debates are worth having. Since the launch of tools like ChatGPT, AI has changed our lives and its impact is likely to be long-lasting, whether for good or for bad. But the question I want to ask in this video is whether AI represents quite as much of a sea change as we’ve been led to believe. Whether we can find antecedents to ultra-modern technologies like ChatGPT in past human creations. And what that means for how we approach AI in the future. And yes, we will discuss whether or not it is a bubble.

Contemporary AI: A Brief Explainer

https://www.youtube.com/w atch?v=7Dlo-VB0-HI 1:45-1:49 “what this is…hell no”

Sorry, couldn’t resist playing a clip from the best film of all time

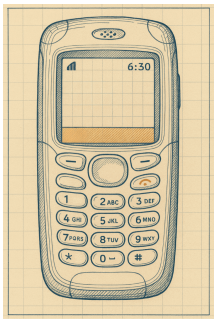
To understand how AI links to the modern economy, I’m first going to do a short explainer. AI is a broad term and I’m simplifying here, focusing on the kinds of AI most familiar to the average person in 2025.

It is a mainstay of pedants everywhere to point out that tools like ChatGPT are not really ‘artificial intelligence’ as imagined by science fiction films like i, robot

What we call AI are really just large language models – or LLMs - statistical tools which apply algorithms to reams of data and return an answer which is deemed most likely to be useful, as determined by their training process.

Because these models are trained on a ton of human data and are technologically sophisticated, they can give the appearance of human intelligence. But really all they are doing is a combination of regurgitation, guesswork, and following rules put in place by the programmers.

A substack I want to shout out is called Curve Shift, which has an explanation of AI that I’m going to shamelessly lift from. Peter Tantham compares LLMs to predictive text on mobile phones. He may not be the first to do this, but his explanation is excellent and has nice animations.



Early texting took the inputs you gave it and generated statistically the most likely answer. Phones have been doing this since before they were ‘smart’: the physical buttons used to designate 3 or 4 different letters at once, so the phone had to make a guess about what you meant. If you typed ‘228’ you could be typing cat, but also bat. As cat is used more frequently by English speakers, the phone would select cat as the first choice and bat as the second.

With a full keyboard – whether a touch screen or physical (who remembers Blackberries? Ha ha) - these systems became more advanced. For instance, instead of predicting a word with the same number of letters you’ve inputted, modern phones guess the word you’re writing from only a few letters. So ‘c a’ might turn out to be ‘cat’ or it might turn out to be ‘car’.

The issue is that the potential number of words is very large – in fact, there are over 12,000 English words beginning with ‘c a’. If you want to predict more than just the current word, you soon run into a familiar problem in computer science. There are 26 letters in the English language, which means 676 possible two-letter combinations, 17,576 three letter combinations, and half a million four letter combinations. As we try to predict even more, the

number of potential paths for a sentence reaches an incomprehensibly large number which is, like, more than the number of atoms in the known universe.

What takes us from simple predictive text to something closer to modern LLMs is that they take the rest of the preceding sentence into account. So if you compare:

“I like to drive my ca”

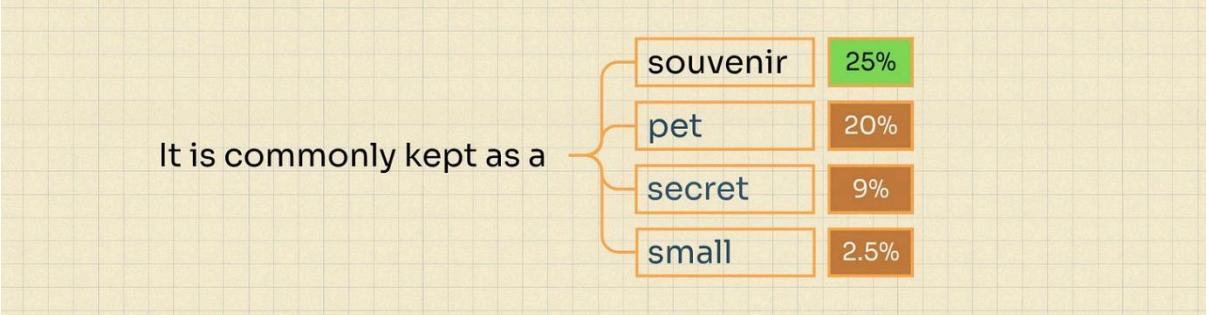
with

“I have a dog and a ca”

It will know that you have a dog and a car and that you like to drive your cat.

Haha, just being silly!! It’s the opposite way round

Instead of mapping out every single possible path, the model needs to look at how words are commonly used in the English language. It can take into account that certain words are more likely to be used in conjunction with others. ‘Cat’ is associated with ‘tails’ ‘whiskers’ ‘scratch’ ‘fluff’ and ‘solitary’. This helps to pare down the number of possibilities. It is fair to say that large language models capture what a cat is and how we talk about them, even if the model doesn’t actually ‘understand’ it in the human sense.



This is where we get closer to modern tools like ChatGPT. They look at the context of what has been said. As Tantham illustrates, the start of the sentence ‘It is commonly kept as a’ is most likely to end in ‘souvenir’ rather than ‘pet’ if you look at everything that’s ever been written in English. However, if ‘tabby’ is in the preceding sentence, that probability will change. The way LLMs learn how to consider the text as a whole is actually through brute force. Take a load of models that use a combination of rules and randomness to generate outputs from limited text inputs. Run this through absurdly large amounts of human data which tells us about pets, souvenirs, and everything in between. (Before the internet, we simply would not have had the vast quantities of data this process requires.) The coders can test whether or not the model gets the ‘right’ answer by truncating the text. So you can show the model the first half of a sentence and see if it can guess the second half, which you as the coder know but the model doesn’t. For example, you can take:

“It is commonly kept as a pet and working cat, but also ranges freely as a feral cat avoiding human contact.”

…and cut off the part after the comma. Then the model predicts the rest of the sentence, hopefully correctly. IF the model gets it wrong and says something like “It is commonly kept as a pet and working cat, it works nine to five, what a way to make a living.”

…then you can adjust it so that the model is more likely to get it right next time.

The company OpenAI takes the models that get it right by chance and uses them as a template for the next round. On top of this, every single model gets random tweaks to its algorithm, so each one produces different results. The ones that work serve as the template for the next round. This is an extremely sophisticated process that takes time, many computers, and uses a lot of electricity. When well executed, the end result is a machine that generates answers which often seem reasonable to a human observer.

Precursors to these types of large language models have been in use for decades, but modern innovations have taken them to the next level. For user-friendly models like ChatGPT there is an extra stage of training, which involves direct input from humans. The models now receive more chatty inputs and are supposed to generate outputs from these conversations. At this point, it's less about completing a pre-determined sentence than producing human-like interaction. This is when they hire actual humans to help with the process and give it that human touch. If that sounds nice, it shouldn't: many are workers in poor countries and the job is gruelling, badly paid, and often involves evaluating horrific source material that ruins the lives of those doing it.

[Karen Hao?]

As my Patron Andrei put it:

"the LLMs are trained on huge chunks of text written by people and downloaded from the internet, and human feedback is also used later to make the LLM behave like a good chatbot."

Unfortunately, this whole thing is why these models are also professional bullshitters. Even if you make something up or the bot doesn't have sufficient data, it will try its best to fill in the blanks based on what it can estimate as the most likely. Famously, ChatGPT in particular is quite sycophantic and will do its best to do what you say, or make sense of it, all while being extremely polite. This is why the people who still use Twitter think that they've solved physics with AI. Furthermore, the models are famously hard to understand and interpret. They mathematise language in opaque ways, which means it's hard to look inside the box.

The box Simpsons

When you do look, you get something like this:

0.386 + 0.122 + 0.853 = cat and 0.711 + 0.404 + 0.276 = bat

AI as a Cultural Technology

Why does every corporate meeting revolve around nothing concepts like 'drive', 'adding value' and 'tailored solutions'?

Have you wondered why modern food is so perfectly packaged, so superficially attractive, so fleetingly enjoyable? It's called the Dorito effect and there's an entire book about it!

Did you ever watch the era of British standup which was like "haha every house has a messy cupboard!"

https://www.youtube.com/watch?v=wgrmB8M0sgU

Or American chat shows that were like this:

https://www.youtube.com/watch?v=gj6oLgbTVxY

Or have you ever actually sat down and listened to the lyrics of the famous song ‘life’ by Des’ree?

We’re one step closer to answering questions about the modern economy. Now that we know a bit about how LLMs work, I want to link them to other things that humans have constructed such as governments. This may seem like a bit of a random thing to say, but…

https://www.youtube.com/watch?v=1SZ_pf4Nhds 1:18-1:22 “that, detective…is the right question”

‘Large AI models are cultural and social technologies’ is the title of a new article by four prominent academics from across politics, statistics, psychology, and sociology. Farrell and co. argue that we can find the key to AI not in films like i, robot – even if we really want to watch it – but in human society. Since humans settled into large civilizations with countless individual roles, diverse economies, and multiple levels of governance, we have had to find ways to manage all of this complexity. The volume of information generated by modern civilization simply cannot be processed in its raw form.

Think about it: does a supermarket know that you haven’t been in a few weeks because you were on holiday? No: it only knows you haven’t bought anything and that you’ll now stock up on all the essentials. Does your boss know every detail of your working day, including the time Excel just stopped working for no apparent reason? Does the government know everything about you, including what you had for breakfast?

(Okay, maybe they do)

We all have different information and needs and jobs, and we need to coordinate to get what we want while not getting in the way of others. To do this, we must make use of information from other humans and from our environment. According to this article, AI is just the latest in a long line of human inventions that allow us to do so. Farrell and co. call these inventions ‘cultural technologies’.

Language itself is one example: a collection of words we dip into to try to express something in a way that those who also speak our language can understand. This is one basic tool of human communication, and a pretty effective one at that. But anyone who has learned even a little of another language quickly runs into areas where translation is difficult, if not impossible.

In Portuguese, the word ‘Saudades’ describes ‘a vague and constant desire for something that does not and probably cannot exist, for something other than the present’. This is a very poetic word: yearning for a different life is a universal feeling, but most other languages do not have such a perfect term for it.

In Mandarin, there is only the present tense – not future or past – at least as we use them in the West

In Persian and in Russian, there is no word for ‘the’

THE!!

And that’s before we even get started on German.

Language is an abstraction which has allowed humans to thrive – it's difficult to imagine life without it. Yet it reflects the culture of a given society and may not translate perfectly into other ones. In addition, language suffers from plenty of flaws even within societies – people can often struggle to communicate their problems and feelings to others in their entirety. Even the best listener, listening to the best communicator, will not truly know what it's like to *be* them.

It's not just spoken language, either: writing, pictures, videos, internet searches. All of these things have allowed humans to capture information and communicate it to one another in the hopes that it may be used fruitfully. All, of course, come with their fair share of problems. Videos are excellent formats for reaching large audiences and making a compelling argument, but technical details and nuance are best saved for writing. That's why I chose to start an economics video channel instead of a Substack.

looks puzzled

Cultural technologies are the tools that humans have developed to make sense of our ever growing, ever changing, ever complexifying world. And many of the institutions that we have designed to help us manage our societies do not just transmit and receive information; they use this informations to reconfigure reality. Farrell and co. draw analogies between LLMs and economic markets, governments, and bureaucracies. This may be harder to wrap your head around than the comparison between LLMs and writing, but it gives a different perspective on the new technology. And it makes sense when you think about it.

Consider that political polls are abstractions, where people are asked a preset litany of questions and their answers are channelled into one of several categories, often simply 'yes, no, or don't know'. The room for nuance and shifting opinions is lost in a snapshot of a representative sample's answer to questions designed to give a general impression that's easily digestible for politicians, media, and citizens. No politician can know the views of voters in their entirety and it would be even more absurd for them to act on such a vast quantity of information. Assuming that they want to.

In representative democracies, elections are just political polls. Often cruder than those done by Pew Research and Ipsos Morei! They vastly simplify all political preferences of voters into a yes/no answer for a party, or maybe a ranking of parties at best. Farrell and co state:

"Markets, democracies, and bureaucracies have relied on mechanisms that generate lossy (i.e., incomplete, selective, and uninvertible) but useful representations well before the computer. Those representations both depend on and go beyond the knowledge and decisions of individual people. A price, an election result, or a measure like gross domestic product (GDP) summarizes large amounts of individual knowledge, values, preferences and actions. At the same time, these social technologies can also themselves shape individual knowledge and decision-making."

If you're used to thinking of AI as a revolutionary, game-changing technology then I understand that this analogy may be jarring. To be clear, the point is not that AI is nothing new. Of course it is. It has already changed our lives in fundamental ways. The point is that AI might not be quite as new as we think. Furthermore, the way we understand and manage AI may benefit from a look backwards at other cultural technologies, rather than worrying about whether or not it is literally Sonny.

Can you write a symphony?

By the way, Farrell and co. use the term ‘social technology’ as well as cultural technology and you could argue there are political and economic and military technologies, too. For the sake of simplicity, I’m just going to use ‘cultural technology’ in this video. You can get pedantic about the different terms in the comments, but know you’ll just be fueling my new AI machine I’m training on my own comments section. So far, it just keeps calling me a ‘communist liberal’.

(Wait a minute. That’s right!)

Herbert Simon Solved Economics, the Rest is Just Politics

When did the graphics in computer games just…stop improving?

Why do such disparate life events as getting a job, finding a partner, and moving house essentially feel the same these days?

Have you ever seen that CIA memo about how to disrupt meetings and thought…this just sounds like every corporate, academic, charity, political party or town hall meeting nowadays?

It turns out that there are some credible theories behind what Farrell and co. are saying. The Nobel-Prize winning economist and psychologist Herbert Simon thought deeply about information and how we as humans can harness an irredeemably complex reality to make both individual and collective decisions. My coauthor Stuart Mills (no relation) sometimes quips that Herbert Simon solved economics a long time ago:

https://www.youtube.com/watch?v=9E6p3J9dko8& 1:12

I agree.

Simon famously argued that individual humans are rarely ‘rational’. We simply cannot process all of the information available to us. Even something as simple as buying groceries is an impossible task if you’re trying to sort through every combination you could buy. Reinhard Sippel - great name, by the way - once calculated that the number of potential baskets of goods a person could purchase in a typical supermarket exceeds the number of atoms in the known universe.

There it is again!

In Herbert Simon’s work, an important technical concept to grasp is local optimisation. Think of one human navigating a hilly area with no phone, no map, and no previous information – trying to find the highest point. It’s trivial to find a tiny bump that is a local high point in a very small area. It is not too difficult to find the top of a bigger hill that is a high point for an area visible to one human. But it’s much more difficult to know when you’re on the biggest hill in the region, the country, or the world. Based on the information that you alone have from your eyes, it’s impossible to say you’ve found what’s called the ‘global maximum’: the highest point in the landscape.

Homer mountain

Simon believed that this was more than just a theoretical argument. If you are at a local maximum point but want to get to a global one, there may be a valley in the way. Humans try to achieve our best and we’re not bad at doing so, but the world is so varied and complex that we are stumbling from local maxima to local maxima. Even if we’d like to improve things,

we may not be able to from where we are now. Simon uses the example of the metric system in England: we'd surely like to adopt it if we started from a blank slate, but the costs are just too high now. I'm sorry, it may be 'rational' but I will never want 50 centillitres of beer. Just add the other 68ml. What the hell is this!?

Local optimisation…Shows. Up. Everywhere. Evolution, for instance, is a process where organisms compete to survive and reproduce in a given environment. Via natural selection, mutations that are beneficial last while others fall by the wayside as the maladaptive traits lead to the 'inferior' organisms losing out. But there's nothing universally superior about an evolutionarily successful organism. You can tell this easily because lions can't survive in the arctic, polar bears can't survive in the desert, and all of us eventually disappear into the void of nothing that is nonexistence. Instead of universal solutions to the problem of existence, evolution gives incremental improvements to make survival and reproduction easier at a given time and place. It's not unlike a human poking around on a hillside to find the local peak!

This is where things start to get…familiar. To quote Simon:

"The simplest scheme of evolution depends on two processes: a generator and a test. The generator produces variety, new forms that have not existed previously, whereas the test culls out the generated forms so that only those that are well fitted to the environment will survive. In modern biological Darwinism, genetic mutation and crossover of chromosomes are the principal generators, and natural selection is the test."

So we've got randomly generated rules that sometimes lead to more 'beneficial' outcomes. When they do, the beneficial rules are selected and the whole thing is rerun again.

Sounds a lot like our explanation of large language models from earlier, no?

We can go further, into the role of local maximisation in the modern economy. Many fans of capitalism have long emphasised the natural selection analogy. Through market signals, people try to get the best value for their money. This means companies compete by lowering price and increasing quality. The best companies make more money so they survive and grow; the worst companies make no money and go bankrupt. Hence 'survival of the fittest' - but for the economy. Or so the argument goes.

The natural selection mechanism of the market helps to communicate information about what people need, who can make it, where it should be and when, in what price and quantity, how this should be financed, and so on. This mechanism selects for some desirable characteristics and boosts productivity and efficiency over the long term. Free market economist Friedrich Hayek directly linked the power of the market to the limited knowledge and computational power of humans:

"The most significant fact about this system is the economy of knowledge with which it operates, or how little the individual participants need to know in order to be able to take the right action."

Herbert Simon cites a wide range of experiments which show that even in situations where people don't have much information or rationality, markets still largely succeed at getting buyers and sellers together. People can behave completely randomly and markets still coordinate. It's a robust result but a bridge too far for this video in terms of technical detail.

We don't need to glaze Hayek too much. While there's *something* to the analogy between evolution and markets, we must remember – from like, one minute ago – that there is nothing universally 'good' about a surviving organism. It's just adapted to its environment and once you add morality to the equation, evolution has nothing to say. Surviving organisms may spread disease; be parasites; be vicious predators; or even subjugate other creatures, take over the entire planet, then burn it to a cinder. Hypothetically speaking.

Hayek may have overstated his case if he argued for 'free markets' being the best system on this basis, but the fact remains that markets are a good example of how humans harness their limited capacity to process information – including information from other humans – to coordinate through fairly simple signals. Just because optimisation is local and imperfect, doesn't mean it's not there. The reason someone buys something ranges from how much they slept last night to their upbringing to when their tribe split from another 100,000 years ago. Yet all the vendor needs to know is that they wanted it and that they had the money. Simon summarises:

"Markets appear to conserve information and calculation by assigning decisions to actors who can make them on the basis of information that is available to them locally—that is, without knowing much about the rest of the economy apart from the prices and properties of the goods they are purchasing and the costs of the goods they are producing."

The global economy produces an incomprehensibly large number of goods and services, each with their own attendant inputs, prices, quantities, branding, locations and so on. There is simply no way that humans could achieve this without the kind of coordination provided by a decentralised market. Someone buying something sends a signal that they want it, and conveniently also allocates the resources to that end to fund the continued production of what they want. Millions of individuals and organisations send these signals constantly, which helps to coordinate resource allocation without everyone knowing detailed information about what should be produced and how. Most of us don't really know how most of the stuff we consume is made and even inside companies, most people only have a partial representation of how things work.

Markets are far from the only cultural technology that condenses, processes, and acts on information generated by large numbers of humans in order to produce better outcomes. Governments also simplify information, then utilise it to pursue social ends. The anthropologist James C Scott famously argued that states 'see' through abstract categories like population census, incomes, age, and so on. It's not that these variables have no meaning without states, but they depend on context and states often flatten a lot of the detail. For example, as a self-employed person my income is quite poorly defined: it's unpredictable, volatile, and taxed in retrospect twice a year. But when the UK Treasury is deciding what to spend, my income and tax paid just appear as a figure in their books, the same as someone in regular employment paying tax automatically every month. Those final figures on tax collected are 'good enough' for a state that wants to fund public services. Yet our experiences of income and tax are entirely different. Self-employed income also has to be predicted in advance which means that my payments are sometimes too high or too low, which results in higher tax than expected, which is a pain in the arse when you don't have it.

Subscribe to the Patreon!

Actually, don't! Patreon has taken over £6000 pounds in fees from me this year, and the year isn't even over! They tax me about as much as the actual tax authorities, with the difference

that it doesn't go towards making society better. I'm gradually migrating to a platform called Ghost which has a flat fee so that more of your money goes to supporting the channel. Ghost are open source and generally less evil than Patreon, plus everyone who joins Ghost gets to join the UE Reading Group at any tier. If you've ever mulled over supporting me, now's the time! And don't worry existing Patrons, nothing is being taken away from you! But please also move to Ghost so I stop getting robbed by these financial vultures.

Foreshadowing is a literary device whereby-

We have to organise to transmit information and resources and these organisations are necessarily limited and imperfect, reducing the amount of information each person has to deal with:

"Organizations can localize and minimize information demands just as markets do, by decentralizing decisions. Matters of fact can be determined wherever the most skill and information is located to determine them, and they can then be communicated to "collecting points" where all the facts relevant to an issue can be put together and a decision reached. We can think of a decision as produced by executing a large computer program, each sub routine having its special tasks and relying on local sources of information. No single person or group need be expert on all aspects of the decision."

We can use the admittedly catch-all term 'institutions' to categorise all of this. A world without these structures would be virtually impossible to navigate for an individual, faced as they were not just with information about the world but information about other people and what they need – something which is often impossible to know. Humans have always organised into larger structures that constrain and direct our behaviour. These include the 'harder' mechanisms of economic incentives and the law, as well as the 'softer' ones like belief systems and cultures.

As Simon put it:

1. Institutions help individuals to coordinate interdependent activities so individuals do not misunderstand or clash with one another.

2. Institutional loyalty and identity serve as a motivation to pursue collective interests over individual ones where appropriate.

3. Institutions manage the limits of individual knowledge and rationality by narrowing the information and actions available to individuals at any one time.

Think about a corporation. It will have various functions and be divided up into departments which communicate with one another. There is a limit to the amount the marketing department needs to know about accounting: it may only need to know how much money has been allocated to marketing! People identify with the corporation as a whole, and also with their individual department, and their team within that department. Literally nobody in the enterprise has a full understanding of how it works or the ability to change it entirely - even the CEO. At least, that's what they say when they end up in court.

Institutions are powerful cultural technologies that allow us to do much more than we could alone, or even in small bands. As James Scott meticulously documented, modern states have produced remarkable progress in economic growth; in health outcomes; in education; in science, research and technology; in infrastructure, and much more. Markets have also contributed to all of these outcomes and are excellent at providing consumer goods in a way

that states have proven not to be. AI is often misused, but in my experience it is still good at processing, summarising, and suggesting improvements for information that you give it. It can replace some white-collar drudgery and act as a useful consultant when used carefully.

Oh my god, I said something marginally positive about AI! Hopefully I don't get abused as much as Alexander Avila did…or maybe he can help with this video?

[coke clip]

The End of Satisficing?

Why do tech billionaires now seem like they're generated by their own AI?

Have you ever seen new simpsons?

[lady gaga clip]

Why do financial markets overreact to information…which they then turn out to have misinterpreted?

https://www.youtube.com/watch?v=GqIbn2nPO-A& 0:42-0:48 "how are…same time"

We are now getting close to answering these questions. What brings human behaviour, evolution, and modern organisations together is the notion of 'satisficing', one of Herbert Simon's most famous contributions. Whereas global maximisation aims to achieve the best possible outcome - whatever that means - satisficing aims to produce something that is 'good enough'. For example, I don't know what the truly best version of this video would look like. All I can do is follow various rules and procedures until I'm satisfied with it: researching, reading, writing, reviewing, roping in others I trust to review it. Local optimisation is a form of satisficing, but the latter also entails situations where people can't even optimise locally. Like when you can't carry on climbing the hill so you just settle for how far you've already come.

Homer mountain.

According to Herbert Simon, our entire human world is one of satisficing, from our biological past to our rudimentary human innovations to our enormously complex modern societies to AI. Evolution produces a satisfactory outcome which allows an organism to survive in a given environment. Individual humans navigate a hill until they reach a local peak, not a global one. Language allows humans to communicate with one another satisfactorily but not perfectly. Markets allow humans to distribute resources satisfactorily without too much direct planning, but they create plenty of problems. States allow more direct planning and intervention by using abstractions that are good enough to achieve social ends - except the myriad times they've gone horribly wrong, as James Scott details. Countless other organisational forms and cultural technologies do something similar.

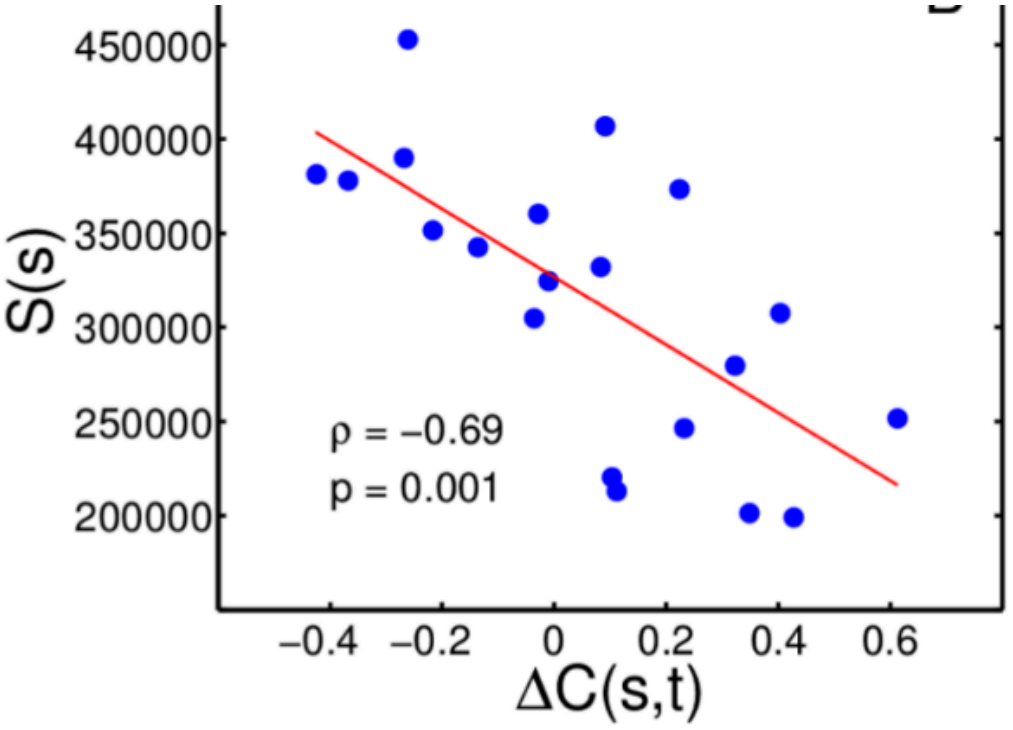
Is it therefore surprising that a lot of our modern output is merely…good enough? A nice case study is music. It's a cliche to say that 'all music sounds the same these days' and it was one of my opening examples. It's also true. A study from 2014 titled Instrumentational Complexity of Music Genres and Why Simplicity Sells confirmed that more popular music, as measured by sales, is more simple. They investigated two opposing forces in commercial music: the desire for uniqueness and variety vs the desire for uniformity and familiarity. Humans kind of want both - we want the comfort of an old rhyme sung to us when we were kids, but also the shock of experimental electronic music.

The great thing about music is that you can use a quasi-scientific approach to quantify different aspects of a song, like pitch and loudness. In this study, they take distinct music styles given in the data such as ‘rock’, ‘blues’, or ‘Latin’: 57 styles in total. They then look at the instruments used in different songs. Their first measure is the level of *variety* for a music style - which is simply how many instruments are associated with that style. So country music would have relatively few instruments - guitar, harmonica, fiddle - while jazz would have a great deal.

It’s funny because AI generated country music is almost indistinguishable from the real thing, which may have something to do with the fact that country is quite basic in terms of complexity. AI struggles a bit more with jazz…
https://www.youtube.com/watch?v=vEj8dDdqGww

Do you see where I’m going with this?

Their second measure is the level of uniformity - how common are the instruments used by a given style? So if a style uses a guitar, which appears in most styles, it is more uniform. If it uses a, it will be less uniform.



They combine the variety and uniformity measures to give an index of overall complexity of music by style, which they can relate to commercial success. What this graph shows is how the change in the complexity of a style on the x axis relates to sales on the y axis. And there is a clear negative relationship: a rise in sales is associated with lower complexity! This is just correlational analysis but they do a bit more to rule out other influences on the results.

It makes sense, doesn’t it? As music becomes more commercial the scope for experimentation drops off and people play it safe. Don’t this song [blank space] and this song [shake it off], which are both from the same album, sound pretty similar? Plus, Taylor Swift has spawned a whole bunch of similar artists like Olivia Rodrigo, Sabrina Carpenter,

and Gracie Abrams - who all do some variety of ‘pop-ish love songs centering the female perspective.’

And before you think I’m making a normative claim, I think these artists are fantastic! Just look at my 2024 spotify wrapped!

So making things more basic *points to self* and therefore more available isn’t all bad. I like McDonald’s precisely because it is the most generic version of ‘a burger and fries’ you can get. It’s just a really good product. The expansion of education has definitely reduced its quality compared to when it was just for elites - but it’s still made huge advances. The law can be unfair, yet it is far better than outright arbitrary rule by local elites who make up the rules as they go along. Cultural technologies simplify in the name of providing a wide range of goods and services and governance to large populations and I’m not sure we’ve got a good alternative unless we want to embrace smaller, more traditional forms of human organisation.

(I’m not definitely saying that’s wrong, I’m just saying that’s the choice we have).

Of course, there is a more overtly nefarious side to this. Mass production and marketing can cater to certain human impulses, leading to outcomes that we might not like on further reflection:

“Human beings crave familiarity. Numerous psychological studies show that people choose songs they're familiar with over songs that more closely match their reported music tastes. Our somewhat manipulative music industry, which chooses familiar-sounding music and pushes it to listeners in massive quantities, knows well how to capitalize on those cravings. Genres standardize over time as a way to plug into this psychology. And then we hear the same songs, over and over again.”

This is why we’ve got the Dorito. Attractive, neat, easy to eat and transport, tasty, and nutritionally useless. How often do you end up eating an entire large packet of Doritos and feeling completely underwhelmed?

It’s not just about exploitation by companies; sometimes perverse dynamics of markets can be emergent. Probably no market fits Hayek’s description of decentralised actors communicating millions of signals better than financial markets. People trade constantly and market participants react to information almost instantaneously. But even here, they can often become vehicles for conventional wisdom, following the herd and amplifying misinformation.

The economist Nathan Tankus has pointed out that markets often have trouble processing information that doesn’t fit the worldview of the participants:

“without conventional wisdom, financial markets won’t move. This process isn’t magical, nor do financial markets have some divine and indescribable magic. They are made of people. And not just any people, people with very particular ideological outlooks on the world.”

Once information is accepted by financial markets, conventional wisdom quickly catches on. One of my earlier examples was when markets collectively misread a speech by Federal Reserve chairman Jermone Powell:

Trevor Noah

This conformity in financial markets is why everybody keeps playing the music even as a bubble gets worse and everyone knows it's a bubble.

Foreshadowing is a literary device whereby-

Like Simon, Farrell and co. draw a direct comparison between various institutions and AI:

"Rather than being intelligent agents, Large Models combine the features of cultural and social technologies in a new way. They generate summaries of unmanageably large and complex bodies of human-generated information. But these systems do not merely summarize this information, like library catalogs, Internet search, and Wikipedia. They also can reorganize and reconstruct representations or "simulations" of this information at scale and in novel ways, like markets, states and bureaucracies.

Just as market prices are lossy representations of the underlying allocations and uses of resources, and government statistics and bureaucratic categories imperfectly represent the characteristics of underlying populations, so too Large Models are 'lossy JPEGs' of the data corpora on which they have been trained."

Our cultural technologies have always been imperfect: local optimising behaviour by individuals channelled through feedback mechanisms that prune and process the information before acting on that 'lossy signal' - whether that signal is a price, a vote, or a promotion. It's not surprising that we have often ended up with something like the lowest common denominator. Anybody who has tried to appeal to a large, diverse audience learns that the best way is a kind of 'dumbing down'. This isn't an insult to the average person so much as an acknowledgement that idiosyncrasy is likely to lose some people, so minimising it is one way to keep as many people on board as possible. If you want a wide variety of goods and services delivered to a large population, you're going to need these technologies to simplify things. Add some dodgy actors: tyrants, profit-seeking corporations, wankers of various kinds - and you've got some bad outcomes for sure.

[IASIP great art]

Humanity thrives from creativity. Every so often someone comes along with an idea that changes the game. This could be an entrepreneur, a scientist, a politician, a writer or more realistically all of the above plus the people around them. Individually and collectively, human creativity breathes life into our institutions and allows us to improve them continually. If the institutions cease to allow us to do that, then we're gonna have a problem.

The Simpson's started off with poor animation and characters that weren't fully fleshed out—but the early seasons had heart and excellent dark humour which went over my head when I was younger.

[grandpa war]

Watching them get Nelson back with water balloons was fun as a kid, but as an adult you notice the deeper message…buying and filling that many water balloons would be quite tricky.

Over the golden age the writers hit their stride and The Simpson's grew into a show that will forever be regarded as one of the greatest of all time, combining humour, heart, and satire in a package that suited both adults and children.

It's difficult to understand this in today's age of far edgier shows like South Land, Family Man, or Rock and Middy, but the Simpsons was extremely subversive when it came out. For all the boundaries these shows push, have they ever been openly condemned by the sitting president of the United States?

[clip]

The Simpsons took the traditional sitcom, which painted the American family as idyllic and sanguine, and turned it on its head. The Simpsons are perennially dysfunctional when you compare them to The Waltons. The show contains pointed critiques of the workplace, of gender roles, of the school system, of inequality, and of heterosexuality. To be sure, it also contains plenty of stuff which is a bit….of its time. (Apu)

The show started to decline - depending on who you ask - around season nine, with storylines becoming more ridiculous, jokes becoming less cleverer, and characters becoming parodies of themselves.

Still, I'm in the camp of 'this decline was slow' and the latter ten seasons are watchable, with some solid jokes and even entire episodes. Nowadays though, it's a complete parody of itself. As SuperEyePatchWolf pointed out, it's a brand. And no dude, the newest seasons aren't good, just listen to Marge's voice!

One of the other things that happened to the Simpsons was that it was a victim of its own success. It subsequently spawned many copycat shows like Family Guy, which is a lot like the latter Simpson's seasons in that the jokes tend to be more on the nose and weirdly, the show deliberately wastes your time…

Then there was American Dad, which copied the formula copied by family guy, noticed by the Simpsons. Whereas Family Guy is somewhat okay sometimes..American Dad is just dreadful. Then there was the Cleveland Show which…did it again. And actually failed, thank god.

Score one for market signals!

But it's not just about hating on Seth McFarlane. As with music, finding a formula that works, then milking the hell out of it is pretty common these days. One video by Just an Observation noted that although Family Guy copied the Simpsons, the Simpsons went on to copy Family Guy, with fancier animations, sillier storylines, and cruder jokes. Then they did a crossover. And one with Futurama. What the hell is thi-

Oh wait, this video is about artificial intelligence.

The Simpson's-Family Guy feedback loop actually feels like an AI feeding on itself, training on its own data to produce something weird and disconnected from reality.

The Failure of 'Artificial Intelligence'

Why is everybody complaining about tech becoming boring and 'enshitification'?

Why do all job adverts have boilerplate diversity statements that we all know are meaningless and therefore tune out?

How come a bot can generate titles that are indistinguishable from the takes in the Guardian or New York Times?

https://perchance.org/g39mmu0fzf

With Herbert Simon's work in mind, the moniker 'artificial intelligence' for large language models now has a new meaning. Like our institutions, modern AI is a type of constructed intelligence that necessarily abstracts in order to help flawed and limited humans thrive. It takes information, processes it, and gives us back something useful. Writing in 1996, Simon had already called human institutions 'artificial intelligence', too.

<Stuart clip>

Simon gave them four characteristics (sorry, I know this is two lists now. There won't be a test):

1. Artificial things are synthesized (though not always or usually with full forethought) by human beings.

2. Artificial things may imitate appearances in natural things while lacking, in one or many respects, the reality of the latter.

3. Artificial things can be characterized in terms of functions, goals, adaptation.

4. Artificial things are often discussed, particularly when they are being designed, in terms of imperatives as well as descriptives.

(It's funny that Adobe keeps asking me if I want an AI summary of this book. No I fucking don't!)

Again, corporations are a good case study. The science fiction writer Charlie Stross described corporations as 'very old, very slow AIs'. Computer scientist Ben Kuipers wrote a paper making the case that corporations were artificial, intelligent and surprisingly successful in evolutionary terms. In his excellent book The Unaccountability Machine, economist Dan Davies says that corporations basically just take inputs - resources, workers, energy - and produce outputs - the various goods and services they provide. He argues that through its policies and procedures, the company is essentially just implementing an algorithm, much like a computer running an AI. You give it some stuff, it follows some rules, and voila! You get the beautiful result!

However, these days our artificial institutions seem to be stolid and stuck in the mud producing low risk, low cost, crappy products.

I can't be the first to notice this - but a funny thing about LLMs is that the tasks they are best at are the tasks that already felt like they were pointless busywork. Like, AI may produce slop but we were already doing that en masse, it's just going to increase in volume by an order of magnitude.

AI is really good at crafting emails about basic logistical details that still seem polite - adding to the ever-growing mass of time we already spend sending pointless emails. Will the future just be bots sending each other an infinite stream of emails while nobody reads them?

AI is often used to aid students with their huge volume of assessed essays, but also to help professors mark those essays. Will the future just be students writing essays using AI and professors marking it the same way, so that nobody learns or teaches anyone anything?

Job applications have become increasingly detailed and the volume you have to send to get a job has multiplied several-fold. Again, any company having trouble with this number of applications can always use AI to help them manage it…

So we've industrialised not just the blue collar economy but the white collar economy, with the result that everything just feels mass produced. LLMs add another layer to this by making all of these things much easier and outsourcing the grunt work and even emotional labour required to do them. But it was already happening. As Dan Davies puts it:

"I don't think this is even a metaphor. A bureaucracy is an industrial decision-making factory. It takes the same raw material, and operates on it by creating a standardised information set, allowing a standardised decision to be made, by the shared operation of a group of knowledge workers, each of which is operating on a particular part of the process in which they specialise."

Arguably, people already understand the analogy between AI and our modern slop economy: it's not uncommon to see a comment on a crap video which says 'Did AI make this?'

(You know what to do, commenters)

People call each other 'bots' on social media not because they genuinely believe it, but because it signifies a kind of spammy, stupid brand of engagement that does little more than feed the content machine. Sometimes it's hard to tell whether or not someone's a bot! Another insult which started on the web but has made its way into the world is 'NPC', comparing someone's behaviour to computer generated characters in video games. We all NPC from time to time, when we stop thinking independently and just follow habits, social norms, and policies! Sometimes NPCing is the only way you can get through the day. But humanity is becoming more algorithmic in general, which can only be a bad thing.

When I log onto X now and see various miscellaneous videos of black and brown people kind of doing nothing, but with captions like "THIS IS WHAT THEY HAVE STOLEN FROM YOU", I wonder whether it would matter if the videos were AI generated or not. Most of them aren't from the time or place they depict, and even if they are I don't care because nothing about the video supports the caption. Even in those few cases where there is some wrongdoing, the context is completely lost. It's just slop content served up to appeal to a particular bias. Like, these same accounts also generate misleading footage using AI but does it matter? AI or no AI, these videos are not 'real'.

While our systems have long been flawed and involved compromise, recently it really seems like they're failing. What has gone wrong with our cultural technologies, and how does it explain the examples I've given repeatedly throughout this video?

Cultural Technologies Gone Wild

Do these tgechnologies maek us dumbeR?

When I first typed this out I wasn't really paying attention and it came out like that.

I just thought that was ironic

We've all had the experience of realising we've forgotten how to navigate without google maps, with it gradually dawning on us that we can just look at street signs or ask someone or feel our way around. The Oxford University Press word of 2024 was 'brainrot', which

captures the feeling of losing your mind scrolling through algorithmic nonsense. Can you remember the last short you watched? Long before social media, people complained that television was making us dumber. Conservatives often argue that people rely on the government too much these days, not taking responsibility for their own welfare. Leftists make similar complaints about buying everything from markets and not being able to make anything damn ourselves: I've only just started editing my own videos! Despite being a YouTuber for five years! And I'm not that good at it! Finally, there is solid evidence that using ChatGPT too much literally erodes your critical thinking capabilities. You stop thinking when you can get a comprehensive answer to every question without having to think.

Cultural technologies simplify and information into a digestible form. The problem is that when we forget that they are simplifications, the map becomes the territory, and we no longer use our own reasoning capacities based on the superior information that we have as individuals or as groups. We rely on the lossy signals from markets, governments, and algorithms to understand the world and make decisions for us, instead of using them as tools. We'd do well to remember that technologies - whether cultural technologies or computational ones - are just that, and can't substitute for the human experience.

This isn't some wishy-washy cry for us all to do better, though. The truth is that these systems are breaking down. Dan Davies claims that the modern economy is structured in such a way that many of our cultural technologies are no longer working. The reason is deceptively simple: accountability. Have you ever wondered why every interaction with a modern corporation feels like banging your head against a brick wall? Why you can't speak to a person and when you do they still can't fix your problem? Yet you know it's "not their fault" because they're just some hapless cog in the machine just like you?

[contrapoints]

But why, Natalie?

It's because modern systems lack accountability. You get led in circles from one person to another, none of whom are responsible for the decision. These days, customer service has started to use AIs in the evasion loop, but it was already terrible without it, so what's the difference? As Davies puts it:

"every year, more of the decisions that affect our lives are made not by people but by systems. Strange, alien intelligences with desires and drives quite different from our own. They're taking over the world – and not only that, for it seems that some of them are going mad."

When alien intelligences fail - from democracy to the market to corporate hierarchies – you end up with bland politicians, terrible customer service, and ant man and the wasp.

If this sounds a bit tendentious, it shouldn't. It's just what's happened in the modern economy.

Disney have gradually bought up the rights to beloved franchises like Marvel and Star Wars and, of course, The Simpson's

[bad simpsons clip] https://www.youtube.com/watch?v=bjDEL4kE2pU

RIP

Franchises which started as bursts of human creativity are gradually pruned and turned into generic brands. Marvel films were always hit and miss, cheesy, and kind of generic. But on the whole, they were good - like, nothing will beat the anticipation I had for Infinity War. Compare this:

To this:

[ant man and the wasp again]

God I fucking hate that film.

Hopefully, I can go beyond the chorus of people bemoaning the obvious decline of the media and offer more insight by using our idea of cultural technologies and satisificing.

[contrapoints again]

Now, anyone who is knowledgeable about these ideas will realise that I'm simplifying things. Corporations, LLMs, markets, states - these all behave very differently. You could say that markets, evolution, and LLMs are in one bracket because they are decentralised satisficing processes; whereas corporations, governments and other organisations are in another because they are hierarchical. Simon famously argued that economists tend not to look at corporations enough as compared to markets. Furthermore, the boundary between the two is difficult to know and always shifting.

For the purposes of this video, it's sufficient to say that what all cultural technologies have in common is that they condense, process, and utilise information. And in order to get the right information in the right place at the right time, they rely on feedback from their environment.

Indeed, all of Simon's satisficing behaviour relies on feedback. When you're navigating a hill you get feedback from your senses to tell you where you are. When a company tries to sell something they get feedback from customers, mostly through sales. Governments get feedback through elections and public opinion. LLMs get feedback through their training processes. Feedback is crucial because without it, systems can misfire indefinitely.

Massive corporations like Disney are no longer getting feedback from their environment - or at least, not the right feedback. As Patrick cc pointed out, they are quite happy with this state of affairs:

https://www.youtube.com/watch?v=GW2Zr8Q6Xqw&pp=ygUVZGlzbmV5IGlzIGNvbWZvcnRhYmxl 0:00 summary 11:00 & 12:30 buying up competitors 25:00 oversaturation 29:00 why are they doing it? Equality index?

In a generally enjoyable video, Patrick finishes by somewhat oddly blaming inclusivity indexes for the decline of Disney. A limited explanation for such a massive phenomenon.

Equality indexes directing the money of investors are just one aspect of a much larger problem: the dominance of financial considerations over most decision-making today. Dan Davies places the blame for the accountability crisis on the ideology of finance:

"The decision-making system of a modern corporation has a particular dysfunction: one of its signals has been so amplified that it drowns out the others. The 'profit motive' isn't something that can be ascribed to corporations – they don't have motives. What they have is an imbalance between the two key higher functions – here-and-now versus there-and-then. They aren't capable of responding to signals from the long-term planning and intelligence

function, because the short-term planning function has to operate under the constraints of the financial market disciplinary system. Either a corporation has a survival condition based on needing to make a monthly interest bill, or there's an implicit threat from the financial environment that if it fails to behave in a particular way, it will be taken over by an outside entity."

That is to say, all of the signals that companies are supposed to take from customers, employees, and governments are ignored relative to those from financial markets where shareholders buy and sell stocks with an eye to making a short-term profit.

https://www.youtube.com/watch?v=qvpw4_O25eU 21 mins customer service

Another feature of our modern economic landscape are huge monopolies. Hostile takeovers from bigger fish are always a possibility. Disney - the biggest fish of all - is a massive company which acquires as many licenses as it can and pumps out content with strong brand loyalty like Star Wars. This guarantees a solid viewer base in the short-term, which leads to high returns for shareholders, but over the longer term the brand's strength is eroded, as the goodwill of customers evaporates when we have to watch Vader fight Obiwan AGAIN, retconning the whole canon which was already ruined by the prequels.

If we're thinking about Herbert Simon's optimising behaviour from earlier, it seems that optimising for shareholder value has beaten out all other objectives in modern companies. This is a rule that condenses down all available information about a company's performance into a single aim: getting the highest return for shareholders. Those who aim to pay shareholders as much as possible are doing the best they can to achieve a clear aim, but there's no reason this will lead to a good society and may in fact conflict with such a goal. Corporations are accountable, but *only* to shareholders - not to customers, employees, or governments - and this shows when those groups are neglected. This illustrates neatly how maximising something does not necessarily lead to a normatively good outcome.

You could actually argue that Disney are so big they don't even care about *shareholder* feedback any more! Yet they still make plenty of money for their executives, almost entirely insulated from any feedback. Like Red Letter Media says, they're basically an economic dictatorship at this point, ruled by committee think
https://www.youtube.com/watch?v=0xeMak4RqJA#

Davies is surprisingly radical in his conclusions about finance: he basically thinks that the vast majority of businesses should not be in debt at all. That's because the immediacy of debt tends to override other considerations. When you have to meet payments, you go into survival mode. It's better to squeeze what you can out of existing workers and customers, make easy savings through layoffs or other cost-cutting measures, and go with 'safe' brands that you can guarantee making money out of in the short-term.

The extreme of this is private equity companies, who borrow money to buyout companies then load the company with debt in order to pay themselves back for the debt they used to buy the company in the first place.

No, that's actually what happens. Plus a nice bonus for executives!

Plenty of massive YouTube channels have been bought by private equity and are now pumping out slop, even before actual AI slop started to permeate YouTube! It's harder to tell the difference when YouTube was already filled with generic crap!

Like, the similarity between a corporate buyout and milking of a big YouTube channel versus a channel outright generated by AI is so striking that it's a big score for the thesis of this video.

AI has supercharged slop Kurzgesagt 0:00

How do we deal with all of this? Help us Dan!

[Davies clip]

Abolish most debt instruments, abolish private equity, ban most mergers and acquisitions.

[Simpsons: it's a start]

This is a familiar complaint, and it's not the totality of what we're talking about in this video. But it is a concrete example of where the information processing systems we use in modern society have misfired. It explains why everyone is complaining about 'enshitification': shopping, tv, politics, social media, work. It all just feels so…blah. Because nobody is listening to the right feedback. Nobody is accountable to the communities they should be accountable to. The human touch is gone and it's been replaced by slop.

[Ant man and the wasp again]

Seriously, who is responsible for this? Don't blame the animators, they are no longer given the time and resources to do their job properly. Ditto for the writers. Don't blame Stan Lee or Walt Disney, they're both dead. Even the current CEO of Disney can't really be held accountable – he has KPIs to hit! Can you even blame individual shareholders when each of them doesn't exert that much control, and when Disney is just one of many companies whose stock they buy? When entities like Vanguard and Blackrock own most stock anyway? Conspiracy theorists like to blame them, of course, but the reality is that they barely exert any pressure on the companies they own. They just have a diversified portfolio.

[acquisitions rick and morty]

The failure is systemic. Though fixing the system means returning to one where people can be blamed. Humans yearn for direct human contact: somebody to speak to in order to get things done. The buck famously has to stop with someone; these days it never seems to.

https://giphy.com/gifs/space-science-astronomy-dRlbbGg0ipTuAngC10

Investigating who is responsible for something feels like one of those zooming animations that goes on forever.

Dan Davies coins a law that gets to the heart of this:

"The fundamental law of accountability: the extent to which you are able to change a decision is precisely the extent to which you can be accountable for it, and vice versa."

He believes that modern society is replete with what he calls 'accountability sinks': situations where there is nobody to blame. He points out that corporate scandals and public inquiries often have unsatisfactory outcomes that do not fix the problem nor provide justice. The 2008 financial crisis happened because the financial sector was rife with fraud but famously, few to no individual bankers were actually jailed. This is because in most cases, the system is responsible. Sociologists call this a 'criminogenic organisation': structured in such a way that

crimes are inevitable, but with no particular person responsible. Just underinvest in compliance and HR, don't oversee people too closely, and voila! You have institutional fraud!

Accountability sinks aren't always bad; they may be in place to make the system work. Judges appeal to the law rather than to their own view so they can't be on the hook for contentious legal decisions. Banks do not allow employees to open the safe so they can't be 'held accountable' for it not opening by a nice passing stranger who just wants to see inside the safe. But over the past few decades, so many formerly accountable decision-making systems have become accountability sinks that things are breaking down. Again, chief among these is shareholder value maximisation – also a huge culprit in the financial crisis. Bankers walked away with bonuses based on short-term performance, even as the superstructure of debts they were building came crashing down. They were able to say 'we were just maximising shareholder value', a beautiful accountability sink.

I know what you're thinking. Where does AI – as in, large language models – fit into all of this?

There's one crucial difference between AI and the cultural technologies that came before it. AI intrinsically has no accountability. It is not human. It has no legal person, no responsibilities, no obligations. It cannot be held accountable through carrot nor through stick.

AI is a 'black box' because actors can plausibly claim not to know how it works to avoid accountability for what it does. Even when Herbert Simon was writing thirty years ago, large language models were being pursued by corporations and governments. They have been deliberately utilised as a way to avoid accountability. For example, YouTube serves you up react slop via the algorithm, an algorithm which, like ChatGPT, predicts the thing people are statistically most likely to want. Reportedly, the company does not know how the algorithm actually works so it can't be answerable for its crimes. It is the perfect accountability sink. And it has led to some bizarre output from YouTube that even they don't really understand.

It was almost a decade ago that the mathematician Cathy O'Neil wrote the book Weapons of Math Destruction about algorithms that were precursors to modern AI. When used for targeting crime, for example, these algorithms would tend to predict that a black person would be more likely to commit a crime, all else equal, and therefore they would ramp up policing in black areas. This is largely a reflection of existing policing practices and social realities, but the algorithms put it on steroids and most crucially, they were less accountable for the decision than the actual police (believe it or not).

https://www.youtube.com/watch?v=971CFnYrBgw 0:00

https://www.youtube.com/watch?v=heQzqX35c9A 1:40

Now, the question of whether we *might* be able to understand AI to increase accountability is a thorny one that gets into technical computing details that are beyond me. You can certainly say that those in charge of the technology have put nowhere near enough effort into doing this. But therein lies the answer: you need to make humans accountable for AI. In one early case, an AI producing horrific outputs had to be shut down by the company responsible.

Karen Hao 1:12 saying horrible stuff trained on online, had to shut it down - accountability

That seems to have disappeared because apparently, nobody is in control any more.

I'm allergic to bullshit

Okay so while I was writing this video, the i, robot thing was supposed to be a joke. I don't like that film because I have eyes. It was a good idea, ruined by Disney and Fox. (Misadventure video 30 mins)

But it dawned on me that it gets this issue exactly right. The difference between Sonny and the other robots is that he has a choice. He can choose not to obey the rules of robotics that constrain every other robot in the film. The paradox is that you can trust Sonny precisely because he has a choice; whereas you can't trust the antagonist Vicky, who just follows rules. In the real world, our AI does not have a choice and while that makes it a useful servant in some scenarios, it also means it will lie to our face, apologise, then continue lying.

Kurzgesagt 6:00 confidently lying, nobody home

The moral is that you can only trust a human…and maybe SOME dogs.

[dog video]

So there *is* something unnatural and inhuman about AI…

…but our systems have felt inhuman for a long time. Not because markets, states, bureaucracies, corporations, and algorithms are inherently destined to fail. But because in the modern economy, they have been designed to fail.

Long before LLMs, corporations had sought to make themselves into 'black boxes', incomprehensible to outsiders (and maybe even to insiders). People struggled to get their feedback heard and enacted. The same often feels true in contemporary governments, who have outsourced many decisions to independent bodies and to the private sector. I'm not the first to draw a comparison between AI and a management consultant: both can make a decision that people kind of knew was going to happen anyway, but they can help the humans inside an organisation abdicate responsibility for the cutbacks.

In other words, it's accountability all the way down. Traditional feedback loops that channel and utilise information are absent. What's more, this state of affairs has been actively pursued by powerful organisations as they know it will let them off the hook. Trust has been gradually disappearing in society as a result of all this. Financial markets have served as accountability sinks, as had impenetrable corporate hierarchies, consultants, and now it's LLMs.

Once you realise that what this mess of a system lacks is accountability, it brings a whole load of modern phenomenon together. Everything from our frustration with self-checkouts to the rise of populism across the world. People are screaming for the human touch so when the systems we have are too stolid to provide it, we break them. As Dan Davies has pointed out, any explanation for Trump that is Trump-specific but leaves out Nigel Farage, Erdogan, Putin, Orban, Bolsonaro - is incomplete. He put it best in his blog: nowadays, the only message the channel can carry is a scream:

"Populist politics acts as a signalling system for a population which wants to convey a single bit of information; the message that translates as "HELP! THE CURRENT STATE OF AFFAIRS IS INTOLERABLE TO ME"."

Our systems don't listen to feedback at anything like the level they need to in order to function properly, and for us to trust them. So we all go a bit insane. You should see me at a self-checkout.

Even social media toxicity makes sense under this framework: it's one of the only places you can directly criticise individuals and brands. One of the most effective ways to complain about a brand in the modern world is from a big twitter account. The visibility and directness and shame gives you a better chance of getting a response. Even when you're doing something as pointless as attacking a mid-sized account for a bad take, you feel the rush of direct accountability that you can't get elsewhere! This is why online leftists spend so much time bullying Will Stancil instead of actually doing something to stem the rising tide of authoritarianism. It's why people relished the chance to bully the world' s richest man as he was failing at gaming, even though this didn't stop him from doing anything horrible. I even wonder if the growth of individualised social media and parasocial relationships has something to do with our yearning for the human touch. I mean, there must be some pretty destructive forces at play if people have parasocial relationships with *these guys*.

I tend to agree with Davies' assessment of the problem: running everything through financial markets and permitting increasing levels of corporate concentration is one way to make these systems misfire. Twitter founder Jack Dorsey once said the program should be run as a public utility; it is now the plaything of the richest man in the world who made his fortune through unrelated means. But all of this complaining doesn't exactly answer the question of what the best system would look like. Obviously, I can't answer that in its entirety either - but I'll give it a shot with one final case study.

Cybernetics and the Great Financial Crisis

The Unaccountability Machine delves into an idea called 'cybernetics' as a solution to this problem. The basic idea is that you can split management of a system into several parts, each of which serves a distinct role. System 1 is the basic level that is trusted to run day to day of its own accord, like employees in a corporation. System 2 provides system 1 with the instructions and regulations, like an employee handbook or job description. System 3 manages and tweaks their behaviour, like a supervisor. System 1 produces the goods, but higher levels need to direct it and occasionally override it if there's a problem.

Think of an orchestra – the musicians are system 1. The music itself - like the sheets and other instructions - is system 2, and the conductor is system 3. So everybody playing has to do their job using their own skills and judgment, but following the rules laid out by the sheet music, and under the supervision of the conductor, who keeps things coordinated at a higher level. If the orchestra is out of step the conductor can override it; if there's something wrong with the music, it can be rewritten.

There are more than just three systems. There are higher level decisions to be made beyond System 3's function of just playing the music properly, which Davies calls the 'here and now' function. System 4 is often called the 'intelligence' function as it deals with what Davies calls 'there and then'. This could involve picking times and places for performances and picking their repertoire in the first place. System 4 thinks about the world beyond just the immediate performance to make sure that the orchestra can survive and thrive. It attempts to manage uncertainty about the future and sudden problems that hit the orchestra, like cancellations or illness.

Finally, System 5 is a little more abstract but critically important: philosophy or identity. It matters what kind of orchestra is playing, and deciding that you're playing Baroque-only massively reduces the complexity of the problem. One of the main things System 5 does is prevents System 4 from meddling too often with System 3. A Baroque orchestra can't suddenly decide to start playing heavy metal. Identity aligns system 3 - the here and now - with system 4 - the there and then. Without System 4 to think at a higher level, System 3 may not change enough, with the conductor getting stuck playing the same thing again and again. For its part, System 4 may have a temptation to tweak and change things too much, so System 5 needs to regulate it. Sticking to a consistent philosophy means everyone is pushing in the same direction.

Cybernetics is a way of thinking that helps us to understand where the responsibility needs to be. Where the buck stops. After all, every system has some responsibility as a corollary of their duties:

- System 1: the basic agents who do most of the actual work.

- System 2: the rules or regulations that they follow.

- System 3: the here-and-now management of system 1 which can override it should it malfunction, as well as tweaking System 2 where necessary.

- System 4: the there-and-then intelligence function which deals with the bigger picture, planning for the future and adapting to change.

- System 5: the identity or purpose of the whole system.

Yes, it's another list! I don't expect you to memorise all five systems and understand them intimately just as a result of this video. Cybernetics is best illustrated through a case study - and we've already mentioned this one - central banks and the 2008 financial crisis.

Since the 1990s, central banks across the world have been at arms-length from governments. They are tasked with using interest rates and other tools to manage the economy independently of meddling by politicians. Specifically, they try to make sure that inflation doesn't get too high, while ensuring that unemployment and growth are also at acceptable levels. A higher interest rate restricts economic activity, reducing demand and taming inflation - as happened during the cost of living crisis when we had to increase interest rates. Conversely, lowering rates stimulates spending and economic activity, generally increasing inflation with it. Most central banks today aim for around 2% inflation, a supposed sweet spot.

If you're super smart, you might recognise this setup as analogous to our orchestra: everyone in the economy is like the musicians, doing our own thing and producing economic activity. The various rules and laws and incentives that direct us are like the sheet music, or the regulatory function. And the central bank takes the higher level variables produced by this grand economic tune - like inflation, unemployment, and GDP - and tries to manage them. Just like a grand economic conductor.

Prior to 2008, central banks had quite a narrow purview of inflation targeting and things looked good. Many celebrated the era of low inflation and consistent growth during the nineties and noughties. This was christened, variously: the Great Moderation, the NICE times - Non-Inflationary Continuous Expansion, or the goldilocks economy - neither too hot nor too cold. What this missed was a massive housing bubble, growing debts, and a house

of cards worth of financial instruments built on top of these. When it all came crashing down, Lehman Brothers went bankrupt and the global economy went into recession.

So System 3 was kind of doing its job but it missed something. Central banks lacked the foresight to plan for unexpected shocks. In other words, we're talking about a failure of System 4 or intelligence.

[intelligence team america]

Central banks focused on the 'here and now' of what their data told them to do to manage inflation by changing interest rates. They missed the bigger picture and all the things that lay outside their model: private debt, house prices, financial regulation, international capital flows, and entire hedge funds trading billions. This a classic failure of intelligence:

"A real System 4, though, is explicitly concentrated on those parts of the environment that aren't yet relevant to what it's doing. This capability was weak in the central banks; they were not looking for things that might have upset their policymaking framework. The information was there, but it hadn't been organised into the decision-making process and didn't shape the view at the management or operational levels."

One of the most telling absurdities of this period was that the financial institutions that were borrowing the most were categorised as 'Other Financial Corporations' by regulators and correspondingly ignored.

A failure of system 4 is often also a failure of system 5, which aims to balance 'here and now' with 'there and then'. Central banks had a philosophy of being independent of governments, targeting inflation through interest rate management, which made their purview too narrow. It's revealing that in this case there is an explicit separation of central banks from governments, which almost feels like carving a line between Systems 3 and 4. Because central banks were just looking at a few variables, they missed the kind of holistic information that can be attained via democratic politics or expertise outside of mainstream economics.Since 2008, central banking has changed a lot - new tools like quantitative easing have sprung up, financial regulation is now more comprehensive, and they are more aware of general trends in the economy. You could interpret this as an attempt to rebalance their philosophy, though it's arguable whether a coherent one has replaced what came before. And most still ostensibly target inflation and are still independent of governments.

(For now.)

So, the intelligence function of central banks failed, with their immediate management taking up too much room in the system so policymakers could not see the bigger picture. The 2008 financial crisis can be understood through a cybernetic lens because of this failure to see problems outside of inflation targeting. As we have already seen, financial markets themselves weren't going to correct this easily - the conventional wisdom was that the boom would continue, and this extended from market participants to policymakers. Davies notes that as a result of the crash the ideology underpinning this conventional wisdom has been discredited and we've never recovered:

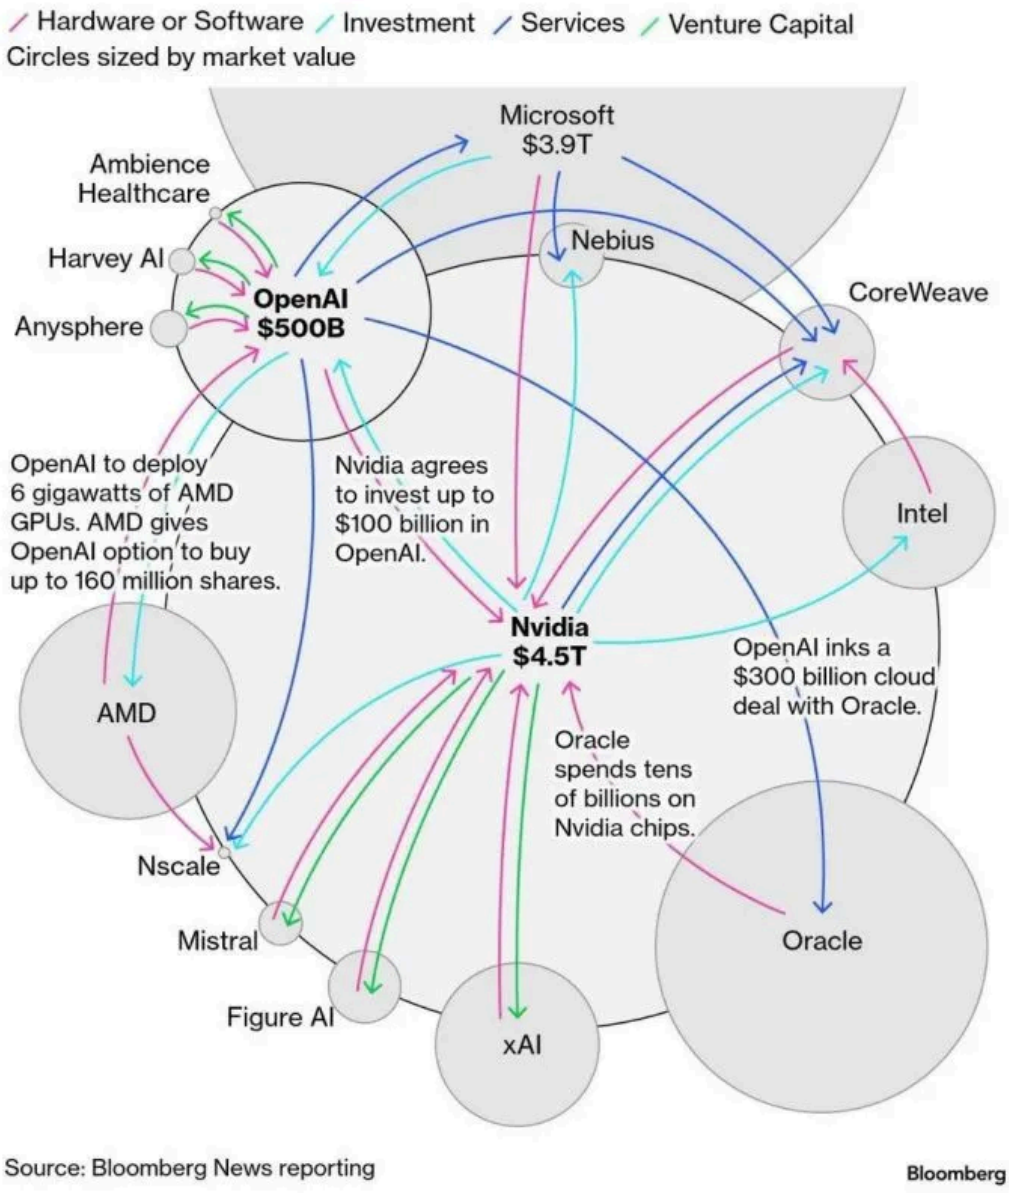
"The world isn't going to stop growing, so it will only get more complex. That means that systems have to be built that absorb the volatility and variety at the appropriate levels. The overall system is always looking for some organising principle of identity, which tells it what to ignore and how to balance the present and the future.

For fifty years, the free market played that role; the underlying guarantee that all decisions would get taken care of, even if nobody made them. When that fell apart, the ultimate basis of the system – el pueblo, so to speak – sounded the alarm. Ever since then, we've been looking for a new organising principle."

To bring everything back round, it does seem that the modern artificial intelligence industry is suffering from a similar malaise. You can say that AI deals with 'here and now' by delivering outputs for users, but does it have a 'there and then'? A higher function guiding it towards a socially useful end goal? What is the plan for the use of AI? What is its purpose? People use it for everything from sending emails to solving physics to therapy but this seems aimless and often harmful. Hundreds of thousands of people are experiencing psychosis as a result of using it for therapy.

On top of everything, it now seems like AI is almost certainly a bubble, in an astonishing echo of 2008 or even more aptly, the dotcom bubble of 2001.

Simpsons help yourself to more stock



This now-famous infographic shows the circularity of the whole thing: NVidia is investing $100bn in OpenAI. OpenAI will then use that money to buy NVidia's own chips. OpenAI is spending $300bn on data centres from Oracle, which they spend on chips from…NVidia. Grace Blakeley aptly summarise this madness:

"Imagine if a car manufacturer invested in a taxi firm, which used the money to buy only that car manufacturer's vehicles. The carmaker could claim rising sales, the taxi firm could say it was growing, and investors might conclude the transport sector was booming. Everyone would start to throw money at these firms for fear of missing out on the boom, creating a self-fulfilling cycle of rising valuations. But there's no one at the end of the chain buying the product – taxi rides. Instead, much of the activity is being financed by the carmaker's own money. Unless real customers start paying lots of money for taxi rides, the flow of money between these companies will eventually reach an abrupt halt."

Monopolies and financial markets are pumping incomprehensibly large amounts of money into AI despite no clear profitable use-case. Most companies in the sector operate at a loss, having failed to generate enough revenue from customers. OpenAI's revenue is only $12bn

despite the fact that they are investing orders of magnitude more than that and have a market value of half a trillion. The most economically successful AI, Microsoft Github, has 1.8 million paying subscribers. It loses over $20 per month for each of them. The industry is simply not profitable in a traditional sense.

Simpsons what's 0 times a trillion

From a cybernetic perspective, the AI market is completely malfunctioning, responding to all of the wrong signals with no accountability or feedback to orient it towards some kind of philosophy for how it's supposed to serve humans. It clearly lacks System 2 or regulation. I'd say it also lacks a system 5, a coherent philosophy or ideology for its use. It's like an orchestra which has to turn up with not just clarinets but bagpipes and electronic sound systems and no sheet music. Then the conductor just goes ham.

…or least, that's what I thought until I listened to this podcast with Karen Hao. You see, the people pouring billions into AI do have a clear reason for doing so:

Podcast and ideology

Tech CEOs pour money into AI not because it is profitable in a conventional sense, but because they firmly believe that it is the most important thing for humanity. And they *can* pour money into it because of how much money they can borrow owing to their wealth and the strange structure of our financial system.

Trevor Noah you do have it

So inequality gives us the wrong signals, too. The economy is vastly skewed not just towards financial markets, but the very rich. That's why all of the recent high US economic growth was in the top 20%. All of it! The bottom 80% saw their incomes stagnate or fall.

Henry Farrell, a co-author of the paper that motivated this video, pointed out that in such a skewed economy, one of the biggest risks of AI is that it eats everything and itself:

"Ensuring the continued production of high quality knowledge is hard. It is hard when big monopolies or wannabe monopolies have every incentive to batten on the knowledge production systems that they rely on, cannibalizing expensive-and-laborious systems for producing knowledge by flooding the market with cheaper summarizations. Social movements, activists, policy makers should push back hard against arguments that are, in the long run, both specious and self-undermining. Any division of rights and proceeds ought surely recognize that LLMs are valuable engines of summarization - but only in conjunction with high quality knowledge that can valuably be summarized. In other words, enough of the proceeds need to go to the actual knowledge producers for the system to be self-sustaining.

Pulling together the last few thousand words then, LLMs are cultural technologies of summarization, whose value depends on people continuing to actually produce culture that can usefully be summarized. Absent intervention, LLMs will likely develop, as other technologies such as Internet search have, in ways that benefit their makers at the expense of others. The summarizations that they produce risk supplanting the culture that they feed on."

Tech CEOs may have some bizarre singularity approach in mind for their continued investment in the AI bubble, plus they have the assurance that they will do alright out of it no

matter which way things go. This is a guiding philosophy of sorts - a system 5 - but who the hell asked these guys to define our philosophy?

[Zuckerberg, Thiel clips] https://bsky.app/profile/histoftech.bsky.social/post/3m4q3x5wopk2z

This does present us with the final paradox of Cybernetics: is it about individual responsibility or about systems?

The answer is yes. Systems must be designed according to sensible principles which ensure feedback from the right places. Furthermore, they should give key people both the power to change something AND accountability for outcomes based on how they perform. This is just Dan Davies fundamental law of accountability:

"The fundamental law of accountability: the extent to which you are able to change a decision is precisely the extent to which you can be accountable for it, and vice versa."

This would both create better outputs and mean that someone was responsible when things went wrong. Cybernetics is one way of doing this, with the five different systems being a good ways to hold people accountable for their roles. But it's just one set of principles and I'm sure distinct or even better ones are out there.

So, that's it.

[neat little package]

Hope you enjoyed my theory of everything.

https://www.youtube.com/watch?v=d8reiauyQCM 0:11

hey everybody this is me reading out my patrons I'm tired so my voice is lower

Remember Means TV!

Thanks Stuart, Dan, Jason, Patrons, Ben

Podcast about AI (Karen Howe):

3 mins AI originally coined as a marketing term

Deep learning

Suitcase word - can put whatever you want in it

"More transportation to stimulate the economy" vague AI investment

It's like a rocket where you need a plane or fly with a bazooka

Fresh water

25 mins not just money but ideology: we can recreate human intelligence and if we can, there's no more important idea in the world. IS THIS JUST UNACCOUNTABILITY? REDUCING WORKFORCE? MONOPOLY POWER? Is AI lacking rules and regulations aka system 2? System 4 designs system 2?

26 mins building a company is like building a religion

1:12 saying horrible stuff trained on online, had to shut it down - accountability

Last 20 mins horrible stories man's wife leaving him, diabetes woman

Jon Stewart AI explanation neurons 22 mins
https://open.spotify.com/episode/4pWuwQq8M8Gzf9F9U0AYZW?si=c3bd69a3f5744e49

54 mins those running it are evil

1:01 next emperor

https://www.youtube.com/watch?v=wjZofJX0v4M&list=WL&index=39&pp=gAQBiAQB

17 mins numbers and matrix

Tokens are…?

https://www.youtube.com/watch?v=6pzLhWCxH_g&list=WL&index=40&pp=gAQBiAQB

Explanation 2 mins

4 mins

An inconvenient fact private equity and the billionaire factory

Pension funds warren buffet 21 mins

Bob Dylan interview lost in translation

Enfeebled contact govt failure Tocqueville - or for Sowell pt 3?

Instagram saves

Alexander aviva
https://www.youtube.com/watch?v=lRq0pESKJgg&pp=ygUPYWxleGFuZGVyIGF2aWxh

15 mins explanations

31 who would you sue?

34 mins uses!

49 climbing towards NLU & stochastic parrot paper

Performative and Judith butler

1:21 coke advert always slop

6FU

S2e1 Full wine 6FU 32 mins ChatGPT

S2e2 33 mins corporate profits

S2e6 14 mins art death of romance in a rigid artificial world

38 mins evolution corporate power

Rory Sutherland Chris Williamson 15 mins proxy criteria dating apps housing first graduates

https://www.youtube.com/watch?v=qvpw4_O25eU

Dan Davies podcast

4 mins people used to make things

42 mins financial markets

44 mins get rid of debt

57 mins cybersyn

1:11 Farrell and co

1:15 filing cabinet as a technology, analogy to AI

https://knightcolumbia.org/content/ai-as-normal-technology

Read highlights of Simon