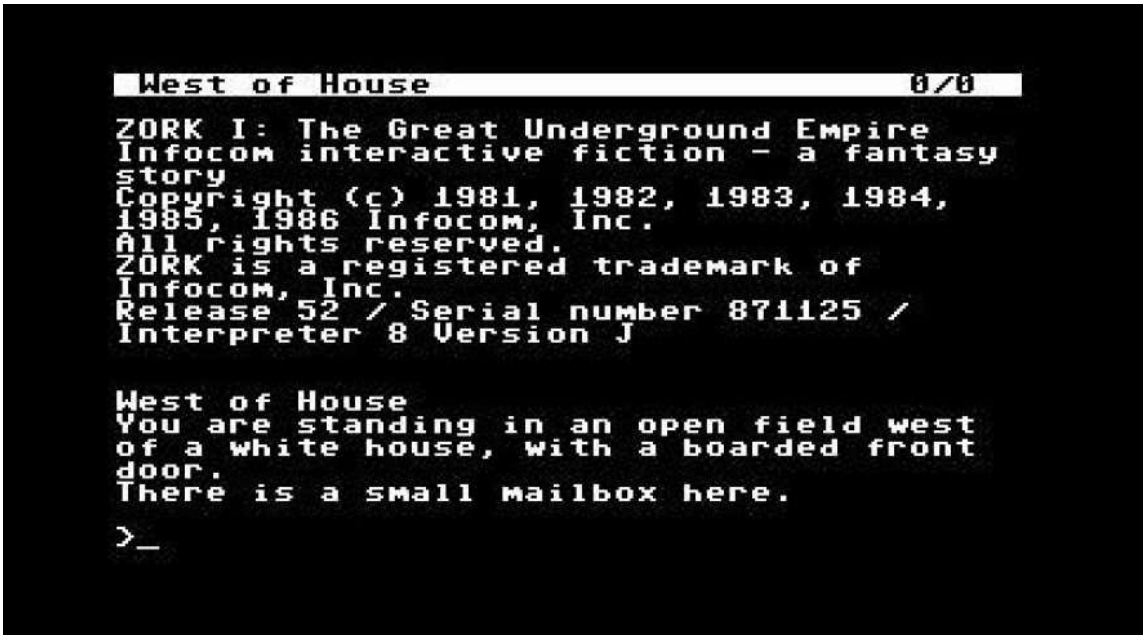


Building the Narrative Machine

December 31, 2018 | [Neville Crawley](#) | [No Comments](#) | [In Brief](#)

[Twitter](#) [LinkedIn](#) [Email](#) [Print](#)

< 7+



This week, inspired by Ben’s note [We’re Doing It Wrong](#), I thought that I would write a follow up note on ‘*Eight learnings on incorporating Narrative observations into a system*’.

This note is not written to be deeply technical, nor is it written from an envelope-pushing theoretical perspective. Rather it is from the perspective of a practitioner and therefore from a practical – and hopefully digestible – point of view on how to use visualizations of Narrative-space in real-world decisions and analysis.

By way of background, I’ve stared at and tried to make actionable sense of many, many thousands of machine-generated (or semi-machine generated) narrative analysis over the

1

POST REMAINING

x campaigns to military intelligence to corporate disaster

SUBSCRIBE | LOGIN

And yet I feel like I've only touched the very surface of the power of this approach and what can be done. There is so much work to do in this area, so much potential, and while the Narrative is as probably old as language itself (probably older), it feels more relevant that ever right now to be able to observe, understand and incorporate the Narrative into our systems for making sense of the world.

Couple of framing comments before we get going with *'the learnings'*:

- I'm assuming that most Epsilon Theory readers are not deeply technically engaged in Machine Learning so I've used terms such as 'vector' and 'dimensionality reduction' in somewhat liberal, expressionist ways that I think are intuitive and most usefully convey core concepts, rather than in precise, formal ways (and stayed away from terms like 'centrality' and 'vertices' and other less intuitive but more precise graph theory terminology).
- If you are interested in exploring some of the concepts mentioned below in more depth and with more formality, I would suggest starting by working your way through the book [Python Machine Learning](#) by Sebastian Raschka. It is a good, foundational blend of theory and practical application of Machine Learning and requires no prior knowledge of, or experience in, coding (of course, there are many other good resources on this topic too, I just happen to be familiar with and like this one).
- For full transparency, I was previously the CEO of [Quid](#), a pioneer in Natural Language Processing (NLP) and Graph Theory that generates many of the Narrative maps seen on Epsilon Theory. This note is not Quid-specific, but rather takes on the overall topic of machine-assisted Narrative observation. Some of the examples below can be very well executed on [Quid](#), some can be performed using free open source software (with some configuration work), some require custom extensions of both / either.

Learning #1: It's not unstructured data, it's really complex structured data

Firstly, and foundationally, to state the obvious, the Narrative is most often expressed in language (sometimes in pictures, dance and other forms, too, but for today let us stick to language)

1

POST REMAINING

x

SUBSCRIBE | LOGIN

Oftentimes language is referred to as ‘unstructured data’, which is database data model terminology.

However, ‘unstructured data’ has gradually crept into general parlance and led to this kind of implicit, half-formed notion that the Narrative is therefore intractable except in a thin, reductionist way.

This is a mistake.

Clearly we know that language itself is not unstructured, otherwise you could not read this; it just has a high degree of complexity in its structure.

This notion of language and the Narrative being intractable except in a thin, reductionist way is the first mental block to rid ourselves of.

We will come back to this point, but this is why using high dimensional graphs is a useful technique for analyzing the Narrative and why squishing the dimensionality way down into a ‘factor’ to put alongside categorical and numerical variables in a regression-type approach is typically pretty unrewarding.

Learning #2: Narrative \neq Sentiment

The only word that is more irritatingly misused than ‘factor’ when talking about Narrative is the ‘S word’: Sentiment.

Sentiment is to NLP as sushi is to Japanese food. It’s fine, it’s in the set, but is very far from the whole cuisine.

And most sentiment analyzers are like the \$10 takeout lunch special with green-dyed horseradish instead of wasabi – it’s not capturing the right flavor. At all. For example, imagine you are trying to evaluate the sentiment of equity analyst reports, where the word ‘overweight’ is a positive thing, but you have trained your sentiment analyzer on general data (such *New York Times* articles), where ‘overweight’ is part of the same negative vector as ‘morbid obesity’. Your sentiment analysis isn’t going to be just in error, it’s going to be perversely in error.

1

POST REMAINING
SUBSCRIBE | LOGIN


× niyaki, yakitori, tempura, onigiri, sukiyaki and hundreds of
are ignoring.

Beyond Sentiment we can classify and score language by:

- Affect (level of emotion)
- Assurance (level of confidence)
- Technicality (level of subject-area specificity)
- Partisanship (level of social organization embeddedness)
- Fiat-ness (level of opinion-leading effort, as Rusty has been doing [here](#))
- ... and thousands of other vectors we can conceive of

Looking exclusively or even primarily at Sentiment as your vector of meaning in a narrative map is almost always a recipe for confusion. For example, below is a short document marked up with ‘sentiment’ in green and red, but also with “Growth” vocabulary and “Value” vocabulary highlighted in fuschia ellipses.

Question: Is the Barclays report below ‘positive’?



BARCLAYS

Tesla Motors Inc.

Positive Sales Momentum but Valuation Still Rich

3Q13 sales likely above guide – but we still see \$141 as fair value. TSLA's 3Q13 highlights the dilemma we have around TSLA stock. On the one hand, we still believe our \$141 price target is a fair value for TSLA stock, as we fundamentally balance TSLA's near-term success of the Model S with the longer-term potential of building a significant mass-market position. On the other hand, we think 3Q is likely to show a more rapid sales progress than TSLA has guided too, which, along with improvements in gross margin, should continue to propel TSLA's stock price above its fair value.

Gross margin likely to continue positive expansion, but trimming EPS to 17c on higher operating expenses. We assume gross margin excluding regulatory credit is a better lens; we see this metric improving from 14.8% in 2Q13 to 19.4% in 3Q on the basis of a stronger pace of deliveries and likely manufacturing process improvements. While our non-GAAP revenue estimate has increased from \$597mn to \$598mn, we have trimmed our 3Q13 EPS estimate to 17c from 27c (vs consensus 12c) as we anticipate higher R&D and SG&A spending than originally modeled, and we have cut the ASP to better reflect the roll-off of regulatory credits.

Answer: No, if you like Value-oriented constructs like EPS and lower Operating Expenses.

Answer: Yes, if you like Growth-oriented constructs like Sales, Gross Margins and Non-GAAP revenue.

Taken alone, Sentiment tells you almost nothing. Combined with other vectors of narrative meaning, Sentiment can be one (of many) useful dimensions of narrative analysis.

So, please don't be (or let your friends be) that person who thinks Japanese food is sushi – there's a whole world of cuisine and language analysis out there to discover!

1

POST REMAINING
SUBSCRIBE | LOGIN

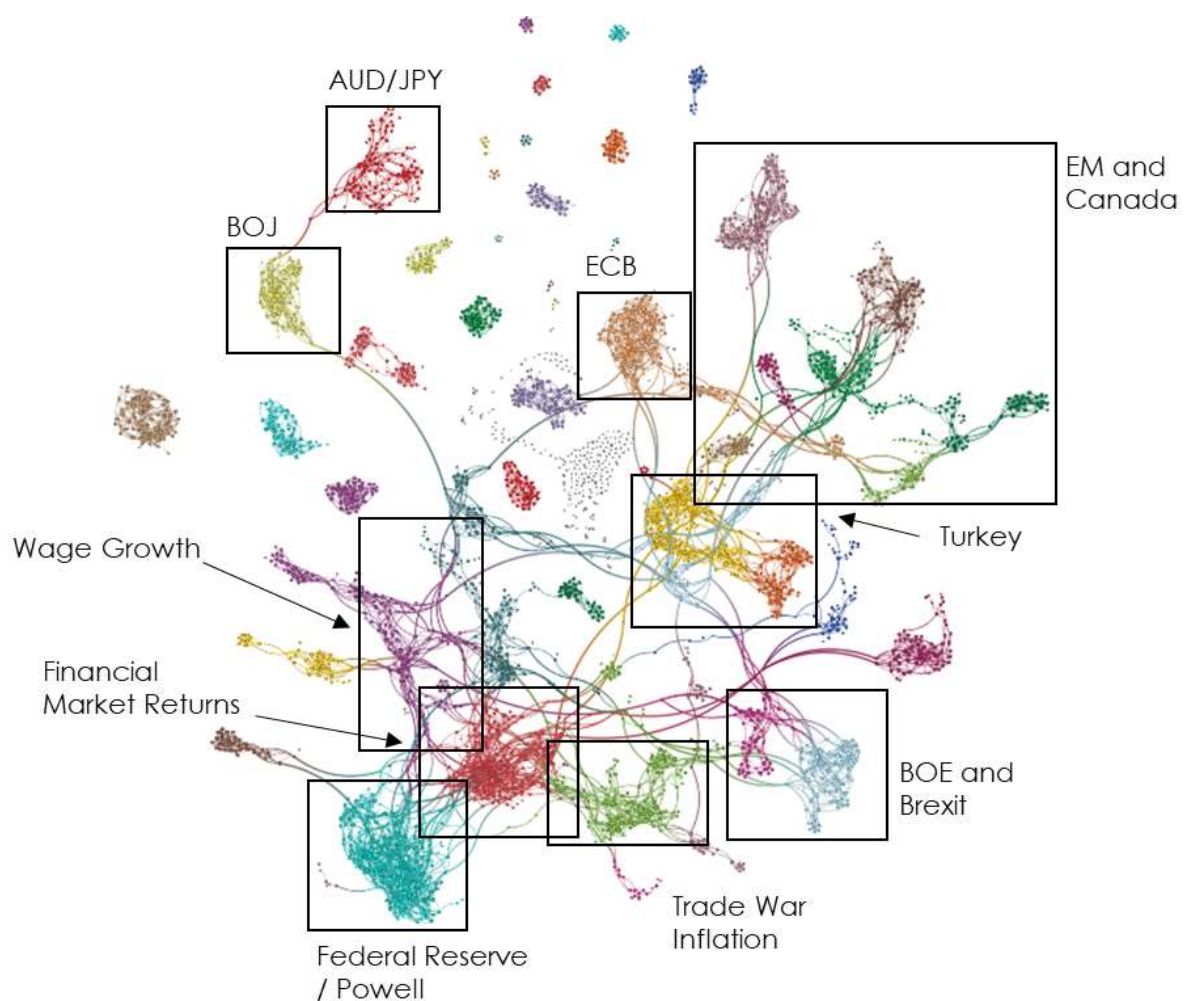
x

Learning #3: Graphs are our friend

To do the analysis noted above of 'Sentiment' or 'Growthiness' for a large corpus of documents, an auto-clustering graph can be very helpful to help prune outliers, boost the score of 'canonical' documents etc., but it is not necessary.

For other type of analysis graphs are super, super useful.

With a graph we can group (cluster) and measure the similarity of documents. For example, here is a graph that Rusty recently generated on the Inflation Narrative using Quid where he has clustered documents based on their linguistic similarity in order to see the key themes:



This is very, very helpful in order to be able to understand the Narrative.

1

POST REMAINING
SUBSCRIBE | LOGIN

x

Learning #4: The art of the graph


Once you really get into NLP clustering graphs you realize that there is a real art – a very human art – to conceiving of and extracting insight by observation from these graphs.

This is because (with flexible software) the permutations are, for all practical purposes, infinite. As a result, being hypothesis driven on the question you are asking, having domain expertise to craft a query, and then having a certain type of intuitive analytical ability to iterate the analysis creates a strong edge.

To try to give an example of this, we might have a question about how the future looks like for Facebook, and so we spin up a graph about this.

Starting with a graph clustered by topic we might then have a sense that (using some of the examples from Learning #2) it would be insightful to score and then observe the documents in the graph on the following dimensions:

- Technicality
- Affect
- Confidence (see below for an old CIA mapping of language to probability)



Having done this, we can then re-cluster and, for example, distinguish a *'highly technical, low emotion, high confidence tightly clustered cluster'* from a *'non-technical, high emotion, low confidence loosely clustered cluster'*.

This is clearly valuable and very, very hard to do without a graph that is visualized.

Learning #5: Enter the Missionary

The worst sin of all is looking at Narrative without considering the Missionary.

The number of times someone has proudly told me *"we analyze a gazillion tweets in real time with our super bad-ass data mungers"*... but in a way that does not distinguish whether there is,

1 POST REMAINING x a specific topic, and at a specific time ...
 SUBSCRIBE | LOGIN

Ben has written so extensively on this I will not re-hash here, but if you don't know who the Missionaries are – ***on the specific topic, at a specific time*** – then it is very unlikely your Narrative analysis will be very fruitful (except in racking up AWS fees for munging massive, fruitless data sets).

Learning #6: Be the Centaur

We used to talk about Centaur chess – a combined human and machine intelligence as the most formidable chess player possible – until DeepMind ruined it for everyone by constantly winning with a machine intelligence alone.

But Narrative analysis is still way, way, way more complex than chess and so a Centaur approach is the right one for Narrative. But in a specific way.

Clearly, computers are useful for computing graphs based on language similarity.

To be clear we don't absolutely *need* a machine to calculate a linguistic similarity graph: it is conceivable to take a stack of 500 documents (say, analyst reports), mark up each document with a highlighter by its unique n-grams, copy the n-grams onto a post-it note per document, score each n-gram on each post it note by a global tf-idf score, and then use our judgement to group post-it notes together by high score post-it notes.

But, man, would that be tedious, and probably not very accurate.

So, computers are useful.

But so are humans. Especially ones with domain expertise, quantitative abilities and non-linear creative minds (especially, non-linear creative minds).

Bottom line is that in all but the simplest systems, we are currently at the stage where the machine is best at generating a base map of current Narrative reality, but a human is better at then asking insightful questions of the map and / or predicting how the map will evolve and / or its interaction with 'other bodies' (see below), which then leads to actionable insight.

Learning #7: Narrative analysis in markets is a 'three-body problem', but with each body having a continuously changing, long-term unpredictable mass

Generally, simplistically, for the purposes of developing a system, we can think of Narrative in equity price movements as one body in a real life 'three body problem':

1. Fundamental information
2. Technical information
3. Narrative information

If we are going to get serious about building a system that performs in the real world, we must understand that Narrative information does not replace the first two or simply reflect the first two, but is a separate body.

This is really, really, really important.

(note: I add the '*continuously changing, long-term unpredictable mass*' clause, as it does not seem conceivable to me that this 'full system' can be brute force 'solved forwards' like the classical physics example Ben described in his note [The Three-Body Problem](#), whereas it does seem conceivable that the Narrative alone can be a Three Body Problem as Ben describes – more on this and quantum computing at the end of this note)

Now, please take a minute to accept or reject the notion of Narrative as one body in a '*three-body – but with each body having a continuously changing, long-term unpredictable mass – problem*', as I feel fully resolved on this point and so am going to carry on this section with the '3rd separate body' taken as fact (while considering that you should please feel free to add more bodies or sub-divide the first two to conform to your mental model of markets and make it an n -body problem, that's not the important point, the important point is that Narrative is its own body in a three or greater body problem) ...

... so, accepting as fact Narrative as a third body and that the first two bodies are super well understood conceptually and many smart people are hyper-optimizing and hyper-incorporating the available information, it logically leads us to certain practical approaches, such as looking for periods when the other Fundamental and Technical informational inputs are 'weak form' and the Narrative information can become dominant and strong form.

1 POST REMAINING ^x money then this is the really key point.
 SUBSCRIBE | LOGIN

It's not that the other forces will ever be zero, or that you can really predict how long they will stay weak (per the '*continuously changing, long-term unpredictable mass*' clause), but when the Narrative is strong it will set direction (or cause volatility), and even if other forces such as new fundamental information re-emerge in an unexpected way, the Narrative will help create an asymmetric elasticity bound around price movement driven by the new information.

Inversely, as Ben pointed out, if you are betting on Narrative and the Narrative is very weak form (i.e., no one is looking) you won't get paid.

Learning #8: Need more resolution!!!

Does the Narrative Machine work? – Yes, it works.

I've seen it so many times in so many different contexts with systems I would consider to be really quite naive delivering really surprisingly strong results. Oftentimes to the extent that I'm suspicious of the results and spend days and days digging into them. But invariably the results are the results, we can know *why* as it's not a black box, and by working hard the systems improve.

So, if it works and is relatively un-picked over, then why isn't everyone doing it?

I think there are three main reasons:

1. Mental model

Ad agency folks, political campaign folks, Department of Defense folks and most creatives seem to get it very well, so it is not a general mental model problem, but I have bumped into a distinct mental model problem around Narrative for hedge fund / markets folks:

Candidly, the mental model problem is so strong I have pretty much avoided discussing this 'computational approach to Narrative' stuff with hedge fund folks for the past year or so as it almost always goes the same:

- The fundamental-type people tell me about their 'process' and require the Narrative to fit into their existing process as a subservient input into their mental model (including POST REMAINING x ... sobering) as, for example, a way of improving the timing

1

SUBSCRIBE | LOGIN

dy doing.

- Meanwhile, the ‘quantitative’ folks just seem bemused that I’m so excited about graphs and say ‘*we already do NLP and have a proprietary model*’. Considering the amount of time and money it takes to get this to work as an even half-decent system (at least in the way that I’m describing it), this seems improbable given the resourcing and backgrounds of technical people at most funds I’ve met, at least as of 12 – 18 months ago.
- The ‘quantamental’ folks are truly the worst, as they are obsessed with reducing Narrative to a ‘factor’ in a regression alongside thousands of other structured data sets.

To be clear, my experience is that all categories mentioned above are full of very smart people, but somehow the mental model I believe is true about the Narrative just doesn’t seem to fit with the mental model of folks who build systems for trading / investing in markets for a living. Whereas folks who market shampoo and yoga pants for a living seem to really quickly and intuitively get it. Odd, perhaps, until you consider that FMCG ad agency people, DoD intelligence analysts, and political campaigners all truly *live* the Narrative and are relatively unencumbered by the Physics Doctorate and MBA predictive analytics orthodoxy.

Anyway, as I started with, it was actually this point as noted by Ben in [We’re Doing It Wrong](#) that prompted me to write this note.

2. Quite a bit of time and money investment

As noted above, building a scalable, reasonably accurate, extensible system that will allow you conceive of fairly arbitrary Narrative questions, quickly spin them up, and have the system learn from them is, in my view, at least a ‘15 strong engineers for 12 – 18 months’ problem to get to V1.0. So ~\$3M – \$5M cost to get to a basic system that can then be built upon, and then at least a \$5M run rate cost (ideally more) to keep extending, with really interesting results taking a couple of years to get to. So, a reasonable sized dollar and time commitment for a new approach.

I think this is why you really don’t see many shops building their own systems and rather see people using sell-side text analytics tools that are relatively thin applications (i.e. they are ‘hard coded’ to extract certain features from certain types of reports which are then output as a ‘buy/sell’-type signal) or 3rd party point solution providers (like DataMinr, Bottlenose etc.) which are focused on single point solutions sold to many clients (e.g., Dataminr for getting early warnings of event by processing Tweets) but are not true Narrative solutions as we are talking about it here.

1

POST REMAINING

x


SUBSCRIBE | LOGIN

The first two reasons given here are primarily bias barriers, but there is a 3rd barrier – a technical barrier – which is very real: Resolution.

Per above, for about \$5M a year you can build and run a decent, fit-for-purpose, reasonably flexible system for this type of analysis.

However, make no mistake about it, the resolution with which you can see the Narrative through this system is relatively low, and your ability to take a constant stream of images (graphs) and in particular to compare images (graphs) is very, very limited.

My sense is that we are today at the animation equivalent of ~6 dpi, ~1 frame a minute. It's blurry, you can kind of follow the story, but you gotta interpret (guess) quite a bit and it is pretty painful after a while.



Unfortunately, it is just a basic physical fact that comparing large, complex graphs using classical computing is very, very compute intensive (at least with currently available algorithms). So this puts a limit on the computational approach and is why human interpretation is still critical in anything other than quite simple systems.

To be clear, these low-res Narrative observation and calculation systems work really surprisingly well and you can make money from them, they are just nowhere near what they can, should and will be.

So, how will we get to the full Narrative Machine?

To my mind this ability to compare complex objects at low cost and at high frequency will be

1

POST REMAINING


SUBSCRIBE | LOGIN

x and will make this stuff really work.

I remember speaking on the same bill as quantum computing companies D-Wave and 1Qbit two or three years back (an age ago in quantum computing land!) and even then this hit me as absolutely true. This 2016 paper by 1Qbit sets out the case well.

We are not there yet with QC, but my bet is that we will get there within the coming years and that we will then truly, finally be able to achieve the Narrative Machine.

To calibrate what an acceleration of compute power looks like, I leave you with this image of the evolution of Lara Croft rendering resolution as a function of GPU processor improvement.



My point here is a simple one: once a technology starts on a path of increasing resolution, it ALWAYS follows a Moore's Law-esque trajectory of improvement, with uses and implementations at higher levels of resolution that were never even considered in early days. It took 17 years for Lara Croft to become a digital character indistinguishable from imagery of a human actor. It won't take nearly that long for a similar resolution intensity of Narrative-space.

Is it early days with the Narrative Machine?

Yes. But not as early as you might think.

1

POST REMAINING
SUBSCRIBE | LOGIN

x 1 you suspect.



Twitter

LinkedIn

Email

Print

Leave a Reply

Please [Login as a Paid Member](#) as a Paid Member to comment

Subscribe

The Latest From Epsilon Theory



When the Story Bends

By Rusty Guinn | March 24, 2020



One for the Road

By Rusty Guinn | April 18, 2020



Exigent Circums

By Peter C 2020

DISCLOSURES

This commentary is being provided to you as general information only and should not be taken as investment advice. The opinions expressed in these materials represent the personal views of the author(s). It is not investment research or a research recommendation, as it does not constitute substantive research or analysis. Any action that you take as a result of information contained in this document is ultimately your responsibility. Epsilon Theory will not accept liability for any loss or damage, including without limitation to any loss of profit,
 1 POST REMAINING x / from use of or reliance on such information. Consult
 SUBSCRIBE | LOGIN ng any investment decisions. It must be noted, that no

one can accurately predict the future of the market with certainty or guarantee future investment performance. Past performance is not a guarantee of future results.

Statements in this communication are forward-looking statements. The forward-looking statements and other views expressed herein are as of the date of this publication. Actual future results or occurrences may differ significantly from those anticipated in any forward-looking statements, and there is no guarantee that any predictions will come to pass. The views expressed herein are subject to change at any time, due to numerous market and other factors. Epsilon Theory disclaims any obligation to update publicly or revise any forward-looking statements or views expressed herein. This information is neither an offer to sell nor a solicitation of any offer to buy any securities. This commentary has been prepared without regard to the individual financial circumstances and objectives of persons who receive it. Epsilon Theory recommends that investors independently evaluate particular investments and strategies, and encourages investors to seek the advice of a financial advisor. The appropriateness of a particular investment or strategy will depend on an investor's individual circumstances and objectives.

[Terms of Use](#) [Disclosures](#) [Privacy Policy](#) [Glossary](#)

© 2020 Epsilon Theory

1

POST REMAINING

X

[SUBSCRIBE](#) | [LOGIN](#)