

## Through the Looking Glass, or ... This is the Red Pill

There are unwritten rules for almost all social phenomena, from investing to writing epic poetry. Not that these weekly notes aspire to Homeric levels (although they do get pretty lengthy), but Epsilon Theory does follow one unwritten rule taken from the *Iliad*, the *Aeneid*, etc. in that I began this story *in media res* – in the middle of things. [Last week's note](#) looked at two very current issues – the jobs report on Friday June 7<sup>th</sup> and the Fed's market communication policy – through the lens of game theory to develop what I think are some non-intuitive results ... namely that the market's Narrative around US labor conditions is fundamentally at odds with the Fed's communications, creating a major source of instability for global markets. And if you look at the [archived notes I've posted](#), they are almost all focused on some specific event or issue.


But at some point it's important to step back and show in a more general way how game theory works to shape markets. This may strike some readers as too academic, but *c'est la vie*. We will return to our regularly scheduled entertainment next week. This week I want to lift the hood a bit on Epsilon Theory and show you how the engine works, or at least that there is an engine, so that you will trust me enough to get in the car and let me drive you to this or that destination. Epsilon Theory is not a collection of musings, and it's not a blog. There's significant intellectual property here, and you deserve to be convinced of that before you invest more of your time reading what I've got to say.

Along the way, though, there will be payoffs for the patient reader. I'll show you the difference between volatility and instability, why it's the latter we are suffering through today, and why the difference is critically important for your portfolio. Also, I can give you an answer to the lead story in the JP Morgan Market Intelligence note from this Wednesday, June 12<sup>th</sup>:

**Market Update** – “*why are the futures up?*” – as has been the case on most mornings for the last few weeks, **the futures are making a large move for no apparent reason**. On the whole it was a quiet night as far as incremental news is concerned and as a result people had a lot of time to contemplate some of the big recent themes driving trading (bond weakness, EM rout, Fed tapering, whether stocks are whistling past the graveyard, etc).

A game theoretic perspective reveals the all too real dynamics behind the “no apparent reason” for these market swings, and there's nothing academic about that.

An unwritten rule is also called a Convention, and both are just alternative names for Common Knowledge. The best example of a Convention is ... language. There's no inherent reason why we should call a rabbit by the name "rabbit" instead of some other word; it's just a behavioral Convention that English-speaking people have developed over centuries in order to improve their mutual lot in life. There was no Saxon chieftain that commanded people to start calling the long-eared rodent that jumps around a lot a "rabbit" instead of a "gavagai", to use Quine's famous example. Instead, over time it somehow became clear to this group of people that everyone knows that everyone knows that a long-eared rodent that jumps around a lot is called a "rabbit". A behavioral equilibrium to call a rabbit a



Van Quine, photograph by Steve Pyke (1990)

"rabbit" developed without coercion, and as a result hunting long-eared rodents that jump around a lot got a whole lot easier for the group of people who shared this Convention. If you lived in this group but didn't share the Convention – if you insisted on calling a rabbit a "gavagai" and had your own words for lots of other things – well, you probably didn't last very long in this group. Similarly, if you're an investor and you don't share the Conventions of the market ("don't fight the Fed" and its like) – well, you're probably not going to last very long, either.

**The best game theoretic work on Common Knowledge comes from linguists like Brian Skyrms and evolutionary biologists like Edward O. Wilson, not economists.** There's an enormous intellectual depth to these fields that I can't do justice to here, but for our purposes of applying Common Knowledge game theory to markets I want to highlight a few ideas that underpin modern linguistic and biological studies of Convention.

Conventions *evolve* over time – whether we are talking about Conventions governing language or market behavior or any other social behavior – which is why it's so important to know something about history in order to understand behavior. There is nothing eternal or written in stone about any behavioral Convention, and even the most socially entrenched Convention can change with amazing speed. For example, dueling and slavery as state-sponsored Conventions were considered part of the "natural order of things" for thousands of years; within a span of about 80 years in the 19<sup>th</sup> century they

were wiped out globally. That said, it's very hard to see Conventions changing when you're living in one. John Lennon wrote that "it's easy if you try" to imagine a future with alternative sets of Conventions, but for most of us it is nearly impossible. Of course, that doesn't mean that change isn't coming. It always does.

Conventions are formed by communication. Language is an obvious form of communication, but so is buying 100 shares of Apple. Any behavior, if made publicly, is a communication of sorts, whether communication was intended or not. It is a signal.


As social animals our brains are hard-wired to look constantly for communication signals and respond to them. As social animals we train each other from birth to look constantly for communication signals and respond to them. **We can no more ignore a speech delivered by Bernanke than ants can ignore a pheromone emitted by their queen.** At first blush this might seem like a weakness, as something to be avoided or at least mitigated. But it is precisely this heightened sensitivity to signals that makes us, like ants, such a successful species! Human behavior in response to signals – what is more commonly called decision-making – is not chaotic or illogical or counterproductive. On the contrary, it's the finely honed product of millions of years of biological evolution and hundreds of thousands of years of social reinforcement. It's why there are 8 billion of us on the planet today.

The insight of evolutionary studies of linguistic Convention is that because we have been socially organized as a certain type of social animal for millennia and because the wiring of our brains for social success hasn't changed in a lot longer than that, there is an identifiable pattern to our behavior around signals. There is an underlying *behavioral logic* at work in humans. The parameters of that behavior – the Convention itself – may be socially constructed and constantly changing (i.e., there's no natural reason to call a rabbit a "rabbit", or to value gold more highly than peacock feathers), but the logic and pattern of strategic human decision-making are constant over time.

**If you can measure the signals that investors are biologically and socially wired to respond to, and if you can map out the likely behavioral pattern of those responses ... then you can predict how markets will respond to new signals.**

That's what I'm trying to do with Epsilon Theory.

There is a methodology for measuring and analyzing signals. It's called Information Theory. To conceptualize how signals and patterns of strategic decision-making work together to create predictable market outcomes, I have developed what I believe is a novel way of depicting the informational structure of markets. This is the intellectual heart of Epsilon Theory. And not to get all Matrix-y, but once you start to see the market in terms of its informational structure, that in fact, **the market IS an informational structure**, nothing more and nothing less, then you will have a very difficult time going back to seeing it as you once did.




Defining the strength of a signal as the degree to which it changes assessments of future states of the world dates back to Claude Shannon's seminal work in 1948, and in a fundamental way back to the work of Thomas Bayes in the 1700's. Here's the central insight of this work: **information is measured by how much it changes your mind**. In fact, if a signal doesn't make you see the world differently, then it has zero information. As a corollary, the more confident you are in a certain view of the world, the more new information is required to make you have the opposite view of the world and the less information is required to confirm your initial view. There's no inherent "truth" to any signal, no need to make a distinction between (or even think of) this signal as having true information and that signal as having false information. Information is neither true nor false. It is only more or less useful in our decision-making, and that's a function of how much it makes us see the world differently. As a result, the informational strength of any signal is relative. The same signal may make a big difference in my

assessment of the future but a tiny difference in yours. In that case, we are hearing the same message, but it has a lot of information to me and very little to you.

Let's say that you are thinking about Apple stock but you are totally up in the air about whether the stock is going up or down over whatever your investment horizon might be, say 1 year. Your initial estimation of the future price of Apple stock is a coin toss ... 50% likelihood to be higher a year from now, 50% likelihood to be lower a year from now. So you do nothing. But you start reading analyst reports about Apple or you build a cash-flow model ... whatever it is that you typically do to gather information about a potential investment decision.


The graph below shows how Information Theory would represent the amount of signal information (generically represented as bits) required to change your initial assessment of a 50% likelihood of Apple stock going up over the next year to a post-signaling assessment of some new percentage likelihood. These are logarithmic curves, so even relatively small amounts of information (a small fraction of a generic bit) will change your mind about Apple pretty significantly, but more and more information is required to move your assessment closer and closer to certainty (either a 0% or a 100% perceived likelihood of the stock going up).



Of course, your assessment of Apple is not a single event and does not take place at a single point in time. As an investor you are constantly updating your opinion about every potential investment decision, and you are constantly taking in new signals. Each new update becomes the starting point for

the next, *ad infinitum*, and as a result all of your prior assessments become part of the current assessment and influence the informational impact of any new signal.

Let's say that your initial signals regarding Apple were mildly positive, enough to give you a new view that the likelihood of Apple stock going up in the next year is 60%. The graph below shows how Information Theory represents the amount of information required to change your mind from here. The curves are still logarithmic, but because your starting point is different it now only requires 80% of the information as before to get you to 100% certainty that Apple stock will go up in the next year (0.8 generic bits versus 1.0 generic bits with a 50% starting estimation). Conversely, it requires almost 140% of the same negative information as before to move you to certainty that Apple stock is going down.



What these graphs are showing is the *information surface* of your non-strategic (i.e., without consideration of others) decision-making regarding Apple stock at any given point in time. Your current assessment is the lowest point on the curve, the bottom of the informational “trough”, and the height of each trough “wall” is proportional to the information required to move you to a new assessment of the future probabilities. The higher the wall, the more information required in any given signal to get you to change your mind in a big way about Apple.


Now let's marry Information Theory with Game Theory. What does an information surface look like for strategic decision-making, where your estimations of the future state of the world are contingent on the decisions you think others will make, and where everyone knows that everyone is being strategic?

I'm assuming we're all familiar with the basic play of the Prisoner's Dilemma, and if you're not just watch any episode of *Law and Order*. Two criminals are placed in separate rooms for questioning by the police, and while they are both better off if they both keep silent, each is individually much better off if he rats his partner out while the partner remains silent. Unfortunately, in this scenario the silent partner takes the fall all by himself, resulting in what is called the "sucker pay-off". Because both players know that this pay-off structure exists (and are always told that it exists by the police), the logical behavior for each player is to rat out his buddy for fear of being the sucker.

Below on the left is a classic two-player Prisoner's Dilemma game with cardinal expected utility pay-offs as per a customary 2x2 matrix representation. Both the Row player and the Column player have only two decision choices – Rat and Silence – with the joint pay-off structures shown as (Row , Column) and the equilibrium outcome (Rat , Rat) shaded in light blue.

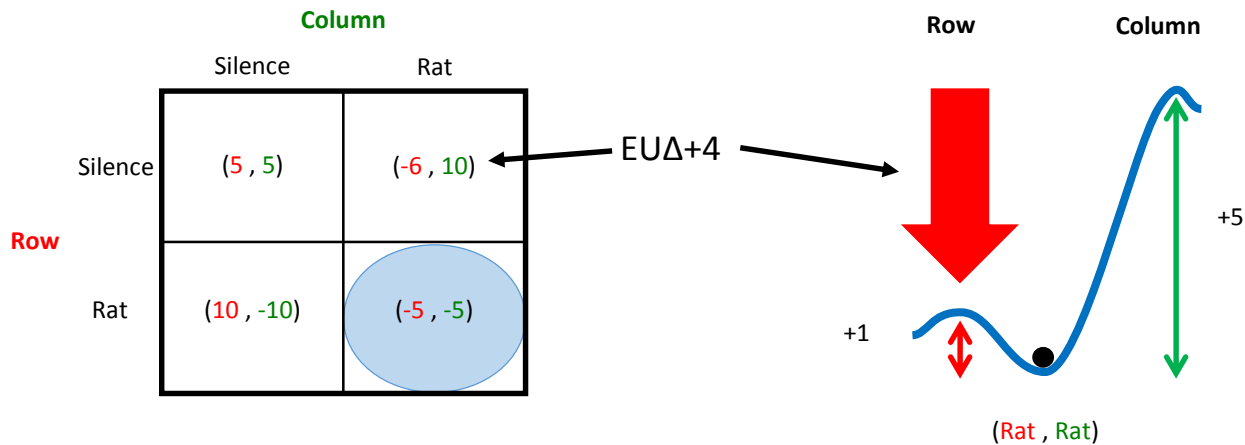
The same equilibrium outcome is shown below on the right as an informational surface, where both the Row and the Column player face an expected utility hurdle of 5 units to move from a decision of Rat to a decision of Silence. For a move to occur, new information must change the current Rat pay-off and/or the potential Silence pay-off for either the Row or the Column player in order to eliminate or overcome the hurdle. The shape of the informational surface indicates the relative stability of the equilibrium as the depth of the equilibrium trough, or conversely the height of the informational walls that comprise the trough, is a direct representation of the informational content required to change the conditional pay-offs of the game and allow the ball (the initial decision point) to "roll" to a new equilibrium position. In this case we have a deep informational trough, reflecting the stability of the (Rat , Rat) equilibrium in a Prisoner's Dilemma game.

		Column	
		Silence	Rat
Row	Silence	(5 , 5)	(-10 , 10)
	Rat	(10 , -10)	(-5 , -5)






Now let's imagine that new information is presented to the Row player such that it improves the expected utility pay-off of a *future* (Silence, Rat) position from -10 to -6. Maybe he hears that prison isn't all that bad so long as he's not a Rat. As a result the informational hurdle required by the Row player to change decisions from Rat to Silence is reduced from +5 to +1.



The (Rat, Rat) outcome is still an equilibrium outcome because neither player believes that there is a higher pay-off associated with changing his mind, but **this is a much less stable equilibrium** from the Row player's perspective (and thus for the overall game) than the original equilibrium.


With this less stable equilibrium framework, even relatively weak new information that changes the Row player's assessment of the current position utility may be enough to move the decision outcome to a new equilibrium. Below, new information of 2 units changes the perceived utility of the current Rat decision for the Row player from -5 to -7. Maybe he hears from his lawyer that the Mob intends to break his legs if he stays a Rat. This is the equivalent of "pushing" the decision outcome over the +1 informational hurdle on the Row player's side of the (Rat, Rat) trough, and it is reflected in both representations as a new equilibrium outcome of (Silence, Rat).






This new (Silence , Rat) outcome is an equilibrium because neither the Row player nor the Column player perceives a higher expected utility outcome by changing decisions. It is still a weak equilibrium because the informational hurdle to return to (Rat , Rat) is only 1 informational unit, but all the same it generates a new *behavior* by the Row player: instead of ratting out his partner, he now keeps his mouth shut.


The Column player never changed decisions, but moving from a (Rat , Rat) equilibrium to a (Silence , Rat) equilibrium in this two time-period example resulted in an increase of utility from -5 to +10 (and for the Row Player a decrease from -5 to -6). This change in utility pay-offs over time can be mapped as:



Replace the words “Column Utility” with “AAPL stock price” and you’ll see what I’m going for. **The Column player bought the police interrogation at -5 and sold it at +10.** By mapping horizontal movement on a game’s informational surface to utility outcomes over time we can link game theoretic market behavior to market price level changes.

Below are two generic examples of a symmetric informational structure for the S&P 500 and a new positive signal hitting the market. New signals will “push” any decision outcome in the direction of the new information. But only if the new signal is sufficiently large (whatever that means in the context of a specific game) will the decision outcome move to a new equilibrium and result in stable behavioral change.







In the first structure, there is enough informational strength to the signal to overcome the upside informational wall and push the market to a higher and stable price equilibrium. In the second structure, while the signal moves the market price higher briefly, there is not enough strength to the signal to change the minds of market participants to a degree that a new stable equilibrium behavior emerges.

**All market behaviors – from “Risk-On/Risk-Off” to “climbing a wall of worry” to “buying the effin’ dip” to “going up on bad news” – can be described with this informational structure methodology.**


For example, here’s how “going up on bad news” works. First, the market receives a negative Event signal – a poor Manufacturing ISM report, for example – that is bad enough to move the market down but not so terrible as to change everyone’s mind about what everyone knows that everyone knows about the health of the US economy and thus move the market index to a new, lower equilibrium level.



Following this negative event, however, the market then receives a set of public media signals – a Narrative – asserting that in response to this bad ISM number the Fed is more likely to launch additional easing measures. This Narrative signal is repeated widely enough and credibly enough that it changes Common Knowledge about future Fed policy and moves the market to a new, higher, and stable level.




So what is the current informational structure for the S&P500? Well, it looks something like this:



The market equilibrium today is like a marble sitting on a glass table. It is an extremely unstable equilibrium because the informational barriers that keep the marble from rolling a long way in either direction are as low as they have been in the past five years. Even a very weak signal is enough to push the marble a long way in one direction, only to have another weak signal push it right back. This is how you get big price movements “for no apparent reason”.

Why are the informational barriers to equilibrium shifts so low today? Because levels of Common Knowledge regarding future central bank policy decisions are so low today. The Narratives on both sides of the collective decision to buy or sell this market are extremely weak. What does everyone know that everyone knows about Abenomics? Very little. What does everyone know that everyone knows about Fed tapering? Very little. What does everyone know that everyone knows about the current state of global growth? Very little. I’m not saying that there’s a lack of communication on these subjects or that there’s a lack of opinion about these subjects or that there’s a lack of knowledge about these subjects. I’m saying that there’s a lack of Common Knowledge on these subjects, and that’s what determines the informational structure of a market.

The unstable market informational structure today is NOT a volatile informational structure, at least not as “volatility” is defined and measured by today’s market Conventions. Here’s what a volatile market structure looks like:



This is an asymmetric informational structure where the signal barriers for the market to go down are much lower than the signal barriers for the market to go up. This structure does not mean that the market will definitely go down; it simply means that the market *can* go down, and *will* go down, with “ordinary” bad news on either increased macroeconomic stress or reduced policy support. The market could still go up, but it would take extremely positive signals on either the macro or policy front to overcome the high informational barrier. Given anything close to a normally distributed set of new market signals, a market with this informational structure is much more likely to go down than to go up, which will be reflected by higher market volatility measures.

I’m a big believer in calling things by their proper names. Why? Because if you make the mistake of conflating instability with volatility, and then you try to hedge your portfolio today with volatility “protection” – VIX futures, one of the VIX ETF’s, S&P 500 puts, etc. – you are throwing your money away. **You are buying insurance for a danger that doesn’t exist right now, and you are leaving yourself unprotected against the danger that is staring you in the face.**

So if you’re reading Epsilon Theory for specific trade ideas, here’s one: *short VXX every time it pops its head up on a big down day as investors rush to buy “volatility protection”*. Hedge it out with some sort of long straddle on the S&P 500 or short positions on underlying stocks if you want to be cautious, but you really don’t need to. So long as most investors mistake instability for volatility (and unless Epsilon Theory gets a LOT more distribution I think it’s safe to say that will persist for a loooooong time) this is an archetypal behavioral trade. I’ll let you know if the informational structure shifts so that volatility really does raise its ugly head.

On that note ... I need to 'fess up about something. The informational structures I'm showing in this note are a rudimentary version of what I'm using in my current research, in at least four ways.

- There's a fractal structure to game-playing, where the same patterns occur on different time frames and different market aggregation levels. Read anything by Benoit Mandelbrot if you don't know what I'm talking about.
- There's a meaningful distinction between backward-looking Event signals, like the release of macroeconomic data, and forward-looking Narrative signals, like the communications of central bankers and politicians.
- There's also a Market signal – what George Soros calls “reflexivity” – that plays an important role in market game-playing. If you've ever watched a stock drop violently without any news showing up on Bloomberg or your traders hearing anything on chat, and then you've sold the stock because “somebody must know something” ... that's reflexivity.
- There's a dimension of time to all this, so that an information surface is three-dimensional (much like a volatility surface in options trading), not two-dimensional as shown in this note. Actually, the information surface is four-dimensional when you add uncertainty and what game theorists call “the shadow of the future” into the algorithms.

And you'll notice I'm not saying anything about the methodology for actually measuring any of this.

I mention all this, not to be coy, but because I want to make clear that there is a depth to Epsilon Theory beyond some interesting but abstract perspective on markets. I mean for Epsilon Theory to have a strong *practical* application to active investment management, and to that end, I think that I'm pretty far along in developing the necessary tools and instruments.

Of course, I also mean to signal that there's a lot more to Epsilon Theory than what I am distributing publicly! I don't want Epsilon Theory to be a black box, not because I have anything against black boxes, but because I think the Convention of trusting a black box had been dying for a long time before Madoff put the final nail into that coffin. I'm not sure what the black box Convention has evolved into, but I'm trying to find out.

As always, thanks for your time and attention.