

## A Game of Sentiment

**We do not sow.**

– *House Greyjoy motto, “A Game of Thrones”*

**It would be ideal if we could have an uncontrolled flow of information.**

– *Poindexter, former Director of DARPA Information Awareness Office*

**It’s SCIENCE, Mr. Felix!**

– *Poindexter, cartoon boy genius*

Sentiment is affect. Sentiment is emotion. Sentiment is fickle and transitory. In the context of investing, Sentiment is how you *feel* about a stock. Sentiment is NOT how you feel about a company, and sometimes it can be difficult to separate our feelings about a company from our feelings about a stock. Difficult, but necessary.

Your feelings about a stock, as opposed to your feelings about a company, should be completely determined by your beliefs about *other investors* and *their* feelings about the stock. In the lingo ... your preference functions for stocks qua stocks are entirely exogenously derived and epiphenomenal. There is no rational internally-developed preference for one stock versus another stock, any more than I prefer a \$5 chip from Harrah’s to a \$5 chip from Caesar’s Palace. The only thing I care about is whether other investors, *for whatever reason*, will value the Harrah’s chip at \$6 tomorrow.

But each of these other investors is thinking about his or her feelings about stocks in the same way that you are. And you all know it. And you all know that you all know it. **The strategic interaction of all investors trying to figure out how all other investors feel about a stock, each of whom knows that everyone else is going through the same decision process, is Sentiment.**



**I know.**

**You know I know.**

**I know you know I know.**


**We know Henry knows, and Henry knows we know it.**

**We’re a knowledgeable family.**

– *Prince Geoffrey, “The Lion In Winter”*

Any strategic interaction is a *game*, in the formal sense of the word. Every game has rules, and every game has a smart way to play and less smart way to play.

### What are the rules for the Game of Sentiment? And how can we play the game smartly?



In 1935, John Maynard Keynes bemoaned the Sentiment game as dominating markets to such a degree that “genuine long-term expectation is so difficult today as to be scarcely practicable.” Sound familiar? Keynes described the Sentiment game as a “Newspaper Beauty Contest”, a media promotion that was familiar to his readers (this was, after all, the heyday of the Miss America contest and “bathing beauty” pageants everywhere), less so to us. Here’s how it works.

A newspaper would run a page of photographs of pretty girls, and readers were invited to mail in a ballot with their choice of the prettiest. If you picked the girl who got the most votes, you were entered into a drawing for some sort of prize. Voting for the girl you think is the prettiest is what Keynes would call the first degree of decision-making.

Now it doesn’t take a lot of thought before you realize that choosing the girl who you truly believe is the prettiest is probably not a winning strategy. *To win, you need to choose the girl who gets the most votes as the prettiest*, and your personal preferences aren’t nearly as useful in that task as figuring out who everyone else is going to vote for as the prettiest. Voting for the girl you anticipate more people will consider to be the prettiest is what Keynes would call the second degree of decision-making.



But there's a big problem with the second degree. It assumes that everyone else is making a first degree decision, that everyone else is making a choice "on the merits" of the photographs and you're the only one smart enough to think about the average preference of the group. As a result, you quickly realize that everyone will be thinking exactly like you are, so you need to make a third-degree decision – *who will get the most votes when all the voters are basing their votes on who they think will get the most votes?* This is the Sentiment game!

Note that this third-level decision probably has *nothing* to do with the relative or objective prettiness of the girls. If "everyone knows" that the brunette with the biggest smile tends to win, then that's where you should place your vote regardless of your personal preference or your knowledge of everyone else's personal preferences. It's the "everyone knows" component of the contest – *regardless of what the contest is fundamentally supposed to be about* – that determines voting behavior and contest winners. To get beyond the third degree of decision-making requires a superior identification of whatever it is that "everyone knows". As Keynes wrote, "We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practice the fourth, fifth and higher degrees."

**So ... how do we practice the fourth, fifth and higher degrees of decision-making? How do we identify the context and strength of what "everyone knows"? How do we *win* the Sentiment game?**

What Keynes called the Newspaper Beauty Contest is known to modern game theory as a variant of the Common Knowledge (CK) game. To illustrate how this game is played on a fundamental level, I'll use another classic example: The Island of the Green-Eyed Tribe.

On the Island of the Green-Eyed Tribe, blue eyes are taboo. If you have blue eyes you must get in your canoe and leave the island at dawn the morning after you learn the awful truth. But there are no mirrors or reflective surfaces on the island, so no one has actually ever seen the color of his or her own eyes. It is also taboo to talk or otherwise communicate with each other about eye color, so if anyone saw that you had blue eyes he wouldn't tell you. Similarly, when you see a fellow tribe member with blue eyes, you are not allowed to tell him the true state of his eye color. It's a small island, so every tribe member knows the eye color of everyone else, but not his or her own eye color.

A Missionary comes to the island one day and announces to everyone, "At least one of you has blue eyes." What happens?

Let's take the trivial case of only one tribe member having blue eyes. He has seen everyone else's eyes, and he knows that everyone else has green eyes. Immediately after the Missionary's statement this poor fellow realizes, "Oh, no! I must be the one with blue eyes." So the next morning he gets in his canoe and leaves the island.

But now let's take the case of two tribe members having blue eyes. The two blue-eyed tribe members have seen each other, so each thinks, "Whew! That guy has blue eyes, so he must be the one that the Missionary is talking about." But because neither blue-eyed tribesman believes that he has blue eyes himself, neither gets in his canoe the next morning and leaves the island. The next day, then, each is very surprised to see the other fellow still on the island, at which point each thinks, "Wait a second ... if he didn't leave the island, it must mean that he saw someone else with blue eyes. And since I know that everyone else has green eyes, that means ... oh, no! I must have blue eyes, too." So on the morning of the second day, both blue-eyed tribesmen get in their canoes and leave the island.

The generalized answer to the question of "what happens?" is that for any  $n$  tribe members with blue eyes, they all leave simultaneously on the  $n^{\text{th}}$  morning after the Missionary's statement. Note that no one forces the blue-eyed tribesmen to leave the island. They leave voluntarily once public knowledge is inserted into the informational structure of the tribal taboo system, which is the hallmark of an equilibrium shift in any game. Given the tribal taboo system (the rules of the game) and its pre-Missionary informational structure, new information from the Missionary causes the players to update their assessments of where they stand within the informational structure and *choose* to move to a new equilibrium outcome.

Note also the role of time in this example. It takes time for strategic observations to take place (the other blue-eyed tribesmen didn't leave right away) and time for strategic calculations to occur. The more ambiguity associated with the Missionary's message or the more tribesmen who have blue-eyes, the more time required for behavior to move the system to a new equilibrium position. Both of these concepts – message ambiguity and prevalence of non-conforming players – play a crucial role in the operationalization of CK game theory into practical investment applications.

But the most interesting aspect of the CK game played on the Island of the Green-Eyed Tribe is the role of the Missionary. **It is the public statement of information, not the prevalence of private information or beliefs, that forces movement in the CK game.** The public statement is what *creates* Common Knowledge,

even if all of that knowledge was already there privately. Everyone must see that everyone else sees the same thing in order to unlock that privately held information and drive individual decisions and behavior.


**The context of what “everyone knows” is whatever the Missionary says it is.**

**The strength of what “everyone knows” is a function of the Missionary’s credibility and the loudness of his voice.**

For better or worse, as investors we love Missionaries. Our reaction is almost never this:



But almost always this:



Who are our Missionaries? Well, there are the obvious choices:



But the more precise answer is that the role of the Missionary is served by any signal that is propagated widely enough and publicly enough so that everyone thinks that everyone has heard the signal. The important thing is NOT that lots of people actually hear the signal. The important thing is that lots of people *believe* that lots of people heard the signal. The power source of Common Knowledge is not the crowd seeing an announcement or a press conference. **The power source of Common Knowledge is the crowd seeing the crowd seeing an announcement or a press conference.** This is why sitcom laugh tracks exist. This is why *American Idol* is filmed in front of an audience. This is why the Chinese government still bans any media mention of the Tiananmen Square protests more than 20 years after they occurred. The power of a crowd seeing a crowd is one of the most awesome forces in human society. It topples governments. It launches Crusades. It builds cathedrals. And it darn sure moves markets.

How do we “see” a crowd in financial markets? Through the financial media outlets that are ubiquitous throughout every professional investment operation in the world – the *Wall Street Journal*, the *Financial*

*Times*, CNBC, and Bloomberg. That's it. These are the only four signal transmission and mediation channels that matter from a financial market CK game perspective because "everyone knows" that we all subscribe to these four channels. If a signal appears prominently in any one of these media outlets (and if it appears prominently in one, it becomes "news" and will appear in all), then every professional investor in the world automatically assumes that every other professional investor in the world heard the signal. So if Famous Investor X appears on CNBC and says that the latest FOMC announcement is a great and wonderful thing for equity markets, then the market will go up. It won't go up because investors agree with Famous Investor X's assessment of the merits of the FOMC announcement. The market will go up because every investor will believe that every other investor heard what Famous Investor X said, and every investor will be forced to update his or her estimation of what every other investor estimates the market will do. The greater the signal strength (from an Information Theory perspective) of Famous Investor X's words, the greater the change in each investor's estimation of every other investor's estimation, and the more the market goes up. That's how the Sentiment game works.



The key to winning the Sentiment game, then, is to listen to all of the signals emanating from all the Missionaries in the world, human and otherwise, as they are mediated through the Big Four channels, and calculate the impact of these signals on investor estimations of all other investor estimations before the actual decision-making process occurs. That may sound like an impossible task, but it's really not.

Just putting together the listening system would have been a logistical challenge even a few years ago, both in terms of data access and computer processing power. Today this part of the solution is almost trivially easy and cheap. And while I'd like to claim that the underlying estimation updating algorithms are ferociously complex, the truth is that we're talking about a pretty simple Bayesian approach.

Fortunately for me, what remains less than trivial is to operationalize Information Theory appropriately in order to gauge the relative strength of this signal versus that signal, this Missionary versus that Missionary, in order to process quickly the market's Bayesian update within the CK game structure. I've [written extensively](#) about what I think is the appropriate framework for just such an application of [Information](#)

---

[Theory](#), so I won't repeat that here. But you don't have to take my word for how to win the Sentiment game ... that a keen ear and a quick Bayesian update can carry the day ... because we all know people who are already using this strategy successfully today. We call them good traders.

What I'm proposing to do systematically is exactly what effective, alpha-generating traders do non-systematically every day. They have internalized an ability to gauge whether a piece of news hitting the market has legs or should be faded. Maybe it's an earnings announcement; maybe it's a jobs report; maybe it's an FOMC announcement. Maybe they trade e-mini's; maybe they trade currency futures; maybe they trade stocks. Regardless of the signal context or the security implementation, **what all effective traders have in common is the ability to, as Keynes wrote, "guess better than the crowd how the crowd will behave."** Keynes would call these traders effective players of the Newspaper Beauty Contest. I would call them effective players of the Common Knowledge game. We can all call them effective players of the Game of Sentiment, and they are one of the last sources of pure alpha in this world.

There's an obvious business opportunity for Epsilon Theory in the full-blown systematization of the alpha-capture process described above ... an externalization of the internalized abilities of effective traders across their various niche-y specialties. But one thing I've learned in 25 years of operationalizing content analysis on a massive scale is that an experienced human is still far and away the most accurate and quickest *interpreter of meaning* in these public signals that are at the heart of any CK game. Computerized systems, particularly learning or adaptive systems, are wonderful for solving various pieces of this puzzle, and they can provide extremely useful tools and guidelines for the entire puzzle. But an experienced human has a knack for contextual pattern recognition of content meaning and salience that cannot be duplicated by either a computer or a roomful of less-experienced humans.

All this suggests an alternative model for applying Epsilon Theory to alpha mining, one where the Epsilon Theory tools and algorithms *support* effective traders to make them *better* at what they have already intuited – that they are playing a complex game with thousands of other market participants, some human and some not, where the decision-making processes of those participants in response to some new signal are somewhat predictable. This is what traders have always done, well before there was even a concept of a stock market or securitization.



Traders are all about buying something, anything, for two copper coins and selling it for three. They are doing this in an environment where lots of other traders are trying to do the same thing, where everyone is playing the same game, and where your “edge” is your ability to assimilate informational signals of all sorts and translate that into an anticipation of everyone else’s behavior. *This is the game that has been played for thousands of years, and the basic rules are the same whether you’re sitting in Hong Kong in 2013, New York in 1935, or Tyre in 800 BC.* But traders today are like those Phoenician captains in 800 BC



roaming the Levantine coast in their little wooden ships looking for their next score ... they have a good ear for news, a good nose for the weather, and a good eye for a bargain, but they’re doing all this by the seat of their pants and with rudimentary tools based on naïve empiricism and personal experience. Can you imagine what a Phoenician trader could have done with a decent astrolabe, much less a weather forecasting service? Just thinking about trading activities within the

right theoretical framework – recognizing that what traders do in their everyday life is to play a game with a certain set of Information Theory rules and a certain set of Game Theory behaviors – will pay dividends in ways that it’s impossible to predict in advance. *This is why theory matters.* Once you put a smart person on the right theoretical path ... once you show him or her how to see the world in a new light, one that illuminates and makes clear behaviors and decisions that were formerly shrouded in myth or custom ... well, there’s just no end to what’s possible.

I’m not sure of the best way to implement this sort of hybrid system, where both human traders and Big Data systems work together under a common theory of how the Game of Sentiment is played. There are many variations on this theme that could be composed, from selling astrolabes as a tool to predicting the weather as a service to forming a network of a few skilled traders and arming them with tools that no one else has. But whatever the final composition, it’s a melody that rings true to my ears, and one that I hope you’ll follow with me at [Epsilon Theory](http://epsilontheory.com).