

# PENSÉE

*Monthly updates from Cognitive Science*

We would be lying if we said we weren't looking forward to the end of 2020. The year was wrought with changes and unanticipated challenges. At the same time, it was an ode to human resilience and science. Together, we stood up against injustice, took responsibility to ensure each others' well-being, and made the world a kinder place. We strived to do good science, and when we couldn't, we cut ourselves some slack.

We hope that in 2021, we will collectively make science kinder and inclusive. We will continue to bring you updates that matter and keep them concise as always. If you have feedback or suggestions, please don't hesitate to [write to us](#).

## IN THIS ISSUE

---

### WHAT'S NEW IN THE WORLD OF COGNITIVE SCIENCE RESEARCH

---


### MEET CGS: JEFFREY LEES ON COMBATING MISINFORMATION

---

### DIGITAL GEM: LEARNING SALON

---

### WE RECOMMEND: DECODE THE FEP



Cascade of rumor spread on Facebook  
Source: Rumor Cascades. ICWSM (2014)

# Rumors, Moods, Physics: Research Updates

---

- **How mood tunes prediction: a neurophenomenological account of mood and its disturbance in major depression**

Karl Friston's Free Energy Principle has been deployed to explain a variety of cognitive phenomena. In this article, the authors combine the first-person phenomenological experience of mood and a neuroscientific FEP-based account of it. Although the relevance of Free Energy Principle over other predictive processing theories can be debated, this neurophenomenological account is significant in its attempt to combine phenomenology with a rigorous framework.

- **Broken Physics: A Conjunction-Fallacy Effect in Intuitive Physical Reasoning**

Conjunction Fallacy is a phenomenon wherein people tend to think of the conjunction of two events as more likely than the individual constituents. The authors demonstrate the existence of conjunction fallacy in tasks involving physical predictions. They claim that people's flawed performance contradicts theories of reasoning that claim that humans simulate physical dynamics of the world while making likelihood judgements.

- **Rumor Cascades**

This paper co-authored, ironically, by three researchers from Facebook, examines the evolution and dynamics of rumors. They find that rumor spread is spontaneous and can be immune to contradictory comments (i.e. comments that question or bust the veracity of the rumor)

- **Using social and behavioural science to support COVID-19 pandemic response**

This paper from 2020 with concerted efforts from scientists across disciplines and universities presents a good template for how science can address emerging crises through an interdisciplinary lens. The paper marks the maturation of Cognitive revolution with its distinct impact.

---

## Events and Opportunities

### **Applications invited: MSc (Cognitive Science)**

Centre for Cognitive and Brain Sciences, IIT-Gandhinagar invites applications for the MSc program in Cognitive Science. More details [here](#).

### **7th Annual Conference of Cognitive Science**

The 7th edition of ACCS will take place online, between the 23rd and 25th January. More details [here](#).

# Handpicked from the World Wide Web

---

## Digital Gems

### **Ten simple rules to colorize biological data visualization**

As scientists, we tend to tell a lot of stories through plots and graphs. Color can have a vital influence on the understandability of data visualizations. In this detailed editorial piece, the authors lay down some thumb rules which will make your data aesthetically pleasing and easy to comprehend.

### **Learning Salon**

A weekly forum to discuss biological and artificial learning, Learning Salon features weekly discussions about core theories and assumptions in the study of learning and intelligence. Watch livestreams and recordings of past sessions on their [YouTube channel](#).

---

## We Recommend

### **Trying to decode the Free Energy Principle**

The Free Energy Principle is conjectured to be the Theory of Everything of Cognitive Science. But people rarely ever understand the specifics of it. [This](#) is Slate Star Codex's 4000+ word-long attempt to decode and critique it.

If you want to go further down the rabbit hole, listen to Karl Friston himself [here](#) and [here](#).

### **The Science of Social Problems**

Manual Galvan is a Psychology graduate student who studies "issues related to racial and economic inequality". Read [his blog](#), if you're interested in wealth gap, a scientific take on racism, and the intersection of the two.

# Misinformation Deep Dive : Jeffrey Lees

Jeffrey Lees is a Visiting Assistant Professor at Clemson University's Wilbur O. and Ann Powers College of Business. His research examines the psychological causes and consequences of inaccurate beliefs. He completed his Ph.D. in organizational behavior and social psychology at Harvard University in 2020.

---

**Tell us a little about your education background. How did you transition from a PhD in Organizational Behavior from Harvard Business School to the work you're doing now?**

I am the first academic in my family and didn't know what academic research looked like until I got to college. I went to a small State University for my undergraduate studies, while I was exposed to research broadly, I didn't find opportunities to engage in it directly. I wanted to get a PhD in Social Psychology eventually and with that goal in mind, I got a Masters at the University of Chicago and then worked as a lab manager at Columbia business school for 3 years. Columbia is where I got exposed to this domain of basic, straightforward social and cognitive psychological research being conducted in a business school context, which differs slightly from the kind of research that is done in a Psychology department's context.



Often, it is more interdisciplinary and more organizational focussed. However, it is not that business schools are taking ideas from Psychology and applying them directly.

The faculty I was working with in Columbia business school were primarily people with their PhD in Psychology, whose interests happened to be a little more related to business than other psychologists. I found that environment more interdisciplinary than traditional Psychology departments, and thus, I applied for my PhD primarily in schools that looked like that; ie business schools where people you work with and the degree you get effective made you a psychologist. While at HBS, I collaborated with faculty in the business school but also faculty in the Psychology department. So my research varied from completely traditional Psychology to more organizational and management work. I have just graduated after 5 years of my PhD and now work as a visiting Professor at Clemson University's business school.

**Within domains like decision making, morality, there are many ideas that are at play when viewed from an interdisciplinary lens and you want to study at the point when all of them intersect. For instance, misinformed beliefs, inherent biases, how does trust arise and the real world consequences of these. Given that there are so many factors at play, what kind of methods do you use and where does one begin to probe these questions?**

There is a global question and a specific question here. The global question is how do you approach interdisciplinary topics. For any topic that is scholarly, you have to understand the audience, you have to understand where the literature is now. You can't come in and evangelize to everybody that they're wrong, because they actually are not wrong. It's that they have one particular perspective, and all perspectives are right in a certain respect but also blind to some other things. It is about bridging the gap strategically so that people are willing to engage with what you are saying.

To give you an example, for psychology, it's about making really strong methodological arguments; that those methods are not designed to capture real world behavior, giving very detailed critique of existing methods. For management it is about updating existing frameworks to say this framework is based on an older theory which we now know isn't the most accurate, here is what we actually know about moral cognition and how it might change the way that you're making predictions. The methods I use, I look like a traditional social psychologist and primarily use behavioral experiments, in person and online, and survey methods.

**What are the factors that lead to some people believing in misinformation/disinformation? One of the factors for instance is that people tend to just be lazy and do not cross check validity. How can we fix this in the real world in a way that makes it less likely for people to believe in misinformation?**

You've prefaced what I was going to say first which is the lazy-take, by which I'm referring to the paper title which summarizes this way of thinking [Lazy, not biased]. The idea is that much of misinformation is not spread by people who are processing information in a biased fashion, but instead that people aren't processing information at all. That does not mean it implies people are sitting there saying this could be misinformation but I don't really care and are unmotivated. It's this old idea in psychology for which there's a lot of existing research which is known as system 1 and system 2, sometimes referred to cold and hot cognition. System 1 is very fast, habitual, automatic and sometimes unconscious cognitive processes. System 2 is very deliberate, effortful, focused and logical. The idea is that much of misinformation spreads passively because of system 1, because we share in the same way as when you wake up, you put your clothes on, it's habitual, you do it automatically because you've done it a million times.

Disinformation is also quite similar in the sense that people don't even get to the point where they're asking themselves if this is true or false. Effective interventions often tend to get people to do that and slow down. Even asking somebody "do you think this is true or false", regardless of their answer, is usually enough to kick in this kind of more deliberate cognition and people are then more likely to notice disinformation. My expertise is in why people believe things that are inaccurate and everything I just described is definitely at play, but there's a much larger question which is when is that actually the culprit. Where I think that holds true is with lay people in countering misinformation in the real world. However, often the sources or primary seeds of misinformation aren't lay people, it's instead people we might call extremists, or partisans for whom we need a different account.

When I think broadly about models of inaccurate beliefs and the behaviors that come out of them, I think of several things. There is the extent to which people hold inaccurate beliefs, very specifically what those inaccurate beliefs are, how did they get them, what are the incentives that create these beliefs. I think a lot of people often have the 'people are dumb or biased' as the end of explanation. People are not dumb, they're pretty good at determining things. Typically when they are not, there are larger environmental factors, social incentives that are pushing them to believe something or behave in a particular way. Then even if you get a pretty robust understanding of why people believe something inaccurate, the next step is to ask what is its relationship to behavior. I think one pitfall a lot of researchers fall into is adopting a very rationalist model, wherein information processing leads directly to behavior, therefore we can show that information processing is flawed, and assume a direct relationship to behavior, but that's not always the case.

There is a lot of evidence showing a pretty weak relationship between the perception of whether a particular piece of news is false and people's willingness to share it. To me that implies there is something other than the inherent truth value of that particular statement which is driving the behavior we care about, which is sharing. So we need to think about the consequences of inaccurate beliefs in a larger social context and the social incentives that drive behavior and interact with true or false beliefs.

**Then how then do we think about the relationship between how people process information while deciding on whether it's true or not and how that process relates to behavior?**

There is an experiment I am going to be running soon that takes a different view of the relationship between true/false beliefs and behavior. I want to introduce two things into that classical model. I want to introduce ambiguity in the judgement of whether something is true or false and I want to introduce reputational motives. Imagine that you are partisan and you are interacting with other members of your in-group and you see a piece of information that a friend posted. You don't know whether that piece of information is true and have some concerns that it is potentially false. Given that you're not sure, another potential motive comes into play, which is the reputational motive, which says that by showing my support for this information or by signaling to my other group members that I believe it, I will gain positive reputational benefit. Because this piece of potential misinformation is either hostile towards an out-group, or is evangelizing the in-group. I will choose to share it because I am unsure of its truthfulness but I think it will garner me positive reputational benefits. So in this context, what I am thinking of is when we are unsure of the truth value of information, these social incentives kick in.

What I am particularly interested in, and would complicate questions about perceptions of truth is that reputational judgement is also a truth judgement. It's me saying that if I engage in behavior X, people will perceive me a certain way. If I exhibit a certain behavior, my friends will like me more and that's a judgement which has a truth value which could be inaccurate. So it's not only the judgement of the truthfulness of a piece of straightforward first order information. There is also a second order judgement of if I do X, Y will happen which can be wrong. So in this context, you have an interaction of multiple pieces of information which may or may not be wrong.

But some of these reputational concerns are also accurate, jumping contexts for instance in hiring and negotiation across gender lines. There is a lot of research that in negotiation, particularly for pay raises, women negotiate less than men. There is also a lot of research that says women have concerns that negotiating will make them look bad, that there are gendered norms which say they shouldn't engage in this behavior, shouldn't be assertive, and this concern is actually accurate. There is also a lot of research to show that when a woman does act like an average man in this context, they are socially penalized. So that is a context where people have accurate second order concerns about how their behavior will be perceived by others which then drives behavior. Where my interest really lies is not how do we process a piece of static information in the world, what we see in an article, but how we process and make truth judgements at the second order level. How we think others in our social network will perceive us, how we think about the incentives and whether or not those perceptions of the social environment are true or false. And also looking at how those true or false judgements in the social environment drive behavior.

**Your views on social factors influencing sharing of misinformation seem true if you look at contemporary political campaigns and the way they build social consensus in favor of specific candidates. That brings us to an important question, how do we address that? An intervention to correct for social influences seems to be significantly harder to design than for purely individual cognitive processing.**

That's the million dollar question. The straightforward answer, which is the low hanging fruit is what we might call fact checking across domains. Whether it is in the traditional fact checking in the sense that newspaper A is saying an article from newspaper B is inaccurate. Or it is fact checking in a scientific intervention sense where I have measured your belief in X and that I am going to tell you that it is inaccurate and provide you with true information. While it is not the most effective, on an average this does work fairly well even though it doesn't get you over the hurdle.

I think one of the reasons it doesn't get you over the hurdle is the global point from before, which is that it is not a one-to-one relationship between believing something being true or false, and behavior. There are all these other mediator processes, even if you can beat someone over the head, metaphorically speaking, to get them to accept true facts, it doesn't always lead to behavior change. For me it is more to do with thinking about the larger consolation of social processes.

Going back to my previous example about a partisan who sees a piece of information with an ambiguous truth value, what I laid out to you is that there are actually two different judgements that may be true that ultimately contribute to behavior. One of them has social incentives, the other doesn't. The point I've been making in my work is that we need to be focused on the social incentives aspect. The classic fact checking paradigm would say, Jeff you've laid out this context but someone doesn't know whether a piece of information is true or false about an article, let's tell them it's true. The problem with that is even if you tell them it's true, all those other social incentives are still at play and that second order judgement might still be wrong. So one argument that I make is that we need to intervene in beliefs where there is some sort of social motive. If you tell someone "That article you are reading is inaccurate", that might change behavior. I would suggest instead to tell them that their judgement about the positive reputation benefit they will garner from spreading this information is wrong and overestimated, if they hit share, their friends are potentially going to perceive them worse. That's also a fact check in the sense that I am providing somebody with true information that updates a false belief, but it is in the domain where there are other social incentives and that changes the social incentive calculation on their part which I believe is much more effective. That's what I've found in my research, when you simply tell people that these second order judgements are inaccurate, you see a much bigger effect on behavior than you do in the traditional first order fact checking literature. Simply believing that in your head doesn't have any positive or negative consequences, but if it is embedded in a social context where certain incentives are at play, then actually intervening on that inaccurate belief can have a much larger effect on behavior.

To your broader point on stopping misinformation, one thing that I stress on is that it needs to be domain specific. This is the thing about social media, the way that misinformation spreads on Facebook is fundamentally different than how it spreads on Twitter, TikTok, Instagram, or even WhatsApp. Those social networks are structured so differently and the companies themselves have different incentives, they move information within their networks differently and we can't make a blanket statement about how disinformation spreads on social media because each platform is different. A student from India that I am collaborating with is investigating fake news in India. Something we've talked about is how several instances of misinformation are directly tied to instances of Mob violence and that these messages are not spread on Facebook or Twitter, they spread instead on private WhatsApp groups. This is not the same as Russian trolls spewing out public disinformation on Twitter, these are private networks which are much more difficult to intervene in from the outside that drive this behavior. So you already have a fundamentally different process by which misinformation spreads which then leads to these terrible outcomes. We were talking about how we should think about this based on existing research and it dawned on us that we cannot, because everything we know about how misinformation spreads on Facebook doesn't apply to a private WhatsApp channel. Because it is so fundamentally different, you can't intervene on it directly. Arguably, people might perceive these private chat networks as more trustworthy than a random piece of information they come across online. If that is the case, it is possibly much more effective because people are much more likely to trust it than they would with random pieces of information they came across on Twitter. To summarize, once you get anywhere above the individual level and are thinking about the structure of social networks and how information spreads, I think the best way to think about stopping misinformation is to broadly think about the following; Which is to focus on debiasing the environment and not the individual. Because those domains are mediated by the structure of the environment. You can try all day to debias the individual but then Facebook could change its algorithms and all those efforts will be for nothing.

**What are some other factors beyond social incentives that influence how open an individual is to losing their biases or being persuaded a certain information is indeed false? For instance, most of our available strategies might not work effectively in a highly polarized environment, so then are there potential factors beyond individual cognition and social incentives that contribute to how receptive an individual or group is to being debiased?**

To clarify what I said before relating to debiasing the individual versus environment, that depends on where you, the individual are at the debiasing level.

If you are interacting with a family member directly then there is a chance that you are able to change their mind. That is largely because of your existing relationship with them. This comes from the view that we are not purely information processors. If both you and I are in a room with your loved one who believes something false, both of us could say the same thing to them but the fact that it is coming from you will have a much stronger impact because of all the other social ways that we process information.

To the other question, there are definitely individual differences that make people more or less sensitive to the truth value of information in the first place and then the willingness to actually change. Things like standardized measures of intelligence and information processing are very good predictors. People who are more educated are less likely to fall for misinformation and more likely to change their opinions. Then there are a host of personality predictors, things like openness to experience and conscientiousness would predict sensitivity to misinformation. There are also negative predictors, like this construct in psychology called reactance. It is this idea that any time you think someone is telling you to do something and you respond by saying no. People who are high on reactance are pretty difficult to intervene on.

These are domain general explanations, but there are also domain specific explanations which are applicable when a piece of information is aligned with your ideology. If you believe the world is flat, that is not going to affect your behavior in the same way as in the Indian context, if you're a Hindu and believe your neighbors are secretly killing cows. Believing the earth is flat is unlikely to lead you to kill anybody for your belief, but people have died due to the latter belief. It becomes harder, from a psychological point of view, to think where being polarized and having this sectarian view fit in. Is it that it prevents people from processing information or is it that people can process the information just fine but this overrides objectivity. Maybe it is the case that most the people who see a video think it's false, but due to the other social incentives, they might think that if it is true and deserving punishment, we must investigate. Another way to look at it would be from the perspective of classic obedience and group polarization, which is that I as an individual don't know whether something is true or false, but everyone around me thinks it's true and seem really angry, therefore, if I question them I might be the target of their violence. Then there is this concept in psychology we called pluralistic ignorance, where we perceive everybody else as an extremist, and because everyone thinks everybody else is an extremist, they are afraid of them and conform to the most extreme behavior when there aren't actually any real extremists. These are other second order inaccurate beliefs which can drive behavior. So in any given context, it is about understanding that it might not be that a certain individual is really polarized or that their level of polarization is preventing them from processing basic facts about the world related to their ideology, it may instead be these other factors. We possibly tend to overestimate the extent to which ideology prevents people from correctly processing information, ideology distorts their view of the world but it doesn't turn their brains off. It might be the case that ideology induces pluralistic ignorance and the way we intervene on that would look very different than if it were just that ideology obscures information processing.



**You propose two possibilities, one being that the individual sharing misinformation perhaps is aware the content is false and choose to share it regardless. In what way does the notion of cognitive dissonance relate to how an individual or group navigates this behavior?**

One thing I really like about the cognitive dissonance framework despite it being an old idea is that it touches upon some of the things I have said before. Cognitive dissonance is about how people process ambiguous information. While often in any classical rationalist model, we tend to think people make hard judgements, the core insight of what makes cognitive dissonance such a revolutionary theory is that it really challenged the classic idea that actions are the outgrowth of beliefs; I believe X which leads to behavior Y. Cognitive dissonance said this relationship is backwards, that in actuality, actions influence beliefs. According to cognitive dissonance, when those two things are in conflict, when our actions and our beliefs are dissonant, we don't change our behavior to match our beliefs, we tend to change our beliefs to match our behavior. So applying this is something that I have thought about as somewhat of an end goal of the work that I am doing. Going back again to that example I gave previously of a partisan sharing a piece of information, I would hypothesize and am actively trying to test that if they share that information, the act of sharing information makes them more likely to believe it in this cognitive dissonance sense. You've engaged in the behavior, and that behavior feeds back into your attitudes, taking the social risk of signaling to my in-group that I think this is true, and that is actually going to make me think it's more true than before I engaged in the behavior. I don't think people in the disinformation space have taken that view, and I'm glad that you brought up cognitive dissonance in this conversation. We've known about cognitive dissonance for a long time but we haven't really applied it to the domain of misinformation yet.

**What is fascinating is that despite belonging to different ethnic, religious, national groups, populist movements tend to share solidarity between them. Most of these models tend to explain solidarities built on either class, religious, ethnic or other commonality shared between groups. What is the glue that connects various far-right populist movements?**

The commonality depends on the level of analysis, I think there is a similarity. For instance, I would describe Trump's and Modi's brand, the types of policies they promote, the kind of people who support them as populist right-wingers who are skeptical of democracy, demagogic and have an ethno-nationalist agenda. They use the position to exalt majority group members and demonize minority group members. In the United States that means white evangelical Christian and in India it means Hindus of high social class. As a psychologist, to me those are very general dynamics that we in political psychology call the template of right-wing authoritarianism. This is in contrast to left wing authoritarians who tend to look different. If we can think about general social cognitive template, I see that connection between not only Trump and Modi, but also Orbán in Hungary, Erdoğan in Turkey, arguably even Putin, a nationalist who after taking over as President reversed a century long trend of downplaying the orthodox Christian church in Russia. I think you're asking the right question, which is if this is happening all at the same time in all these different places around the world, is this maybe a larger global phenomena? That's the point at which my expertise is gone, I am not a sociologist even though I have my own personal political opinions about it that mostly relate to inequality and reactions to capitalism.

**We've talked about the Cognitive aspects, political aspects, but not Technology itself. Why have Technology companies seen very limited success in curbing misinformation, including Twitter that seems to have made serious attempts at addressing the issue?**

We are not sure how much a lot of these interventions work. When you look at American social media companies, they have done much better in 2020 than in 2016. They are much more prepared, they are working much more closely with academics and law enforcement. When it comes to low hanging fruit like deactivating accounts that have been identified as run by Russian security forces, they have gotten much better at that. Back to my earlier point about how these contexts are fundamentally different and how it changes across platforms, Twitter has been much better than Facebook and that is likely in part because of the people who run Twitter versus people who run Facebook, but it is also the incentives for these companies. Facebook makes money when misinformation spreads in a way that Twitter doesn't. Hence, Twitter has been a lot less hesitant to squash misinformation. You also see this with YouTube. There is a lot of talk about right-wing radicalization on the platform and several mass shooters have been radicalized through watching videos there. YouTube and Facebook are domains where the algorithms are designed for you to click on things, which functionally means they feed you increasingly more extreme content. Twitter doesn't have that since it is structured fundamentally differently. While social media companies are getting better, the extent to which that betterness is due to them genuinely caring versus doing it due to social pressure is hard to say. In my mind, the variance of their effectiveness correlates highly with monetary incentives. Those who are making money effectively off of disinformation are the least likely to do anything about disinformation.

My ultimate point and the only way I think these problems are eventually going to be resolved is interdisciplinary work. It needs to not be a bunch of psychologists in one place studying misinformation, a group of sociologists elsewhere and a group of computer scientists in a third place where they maybe talk to each other once in a year. It needs to be the reverse, it needs to be a group of people, all with their primary interest in the phenomena, who happen to have a different disciplinary perspective; rather than being core disciplinary scholars who only occasionally talk to other people. Specially since this cuts through so many lines across ideology, psychology, social media and different levels of analysis, the only way there is ever going to be concrete progress made in this domain is through interdisciplinary research.