

Rationally Speaking episode 135: Robin Hanson on, “Most human behavior is driven by signaling”

Julia: Welcome to Rationally Speaking, the podcast where we explore the borderlands between reason and nonsense.

With me today is our special guest, professor Robin Hanson, who is a professor of economics at George Mason University. He's published extensively in the field of economics, but also in fields like artificial intelligence, cosmology, evolutionary biology, physics, politics and more. Robin is also the author of the popular blog, *Overcoming Bias*. Robin, welcome to the show.

Robin: Great to be here. Let's hope we're closer to rationality than nonsense, but I guess we'll find out.

Julia: It's always a toss up, really.

Robin: Yeah. Well, I want to know what the vote is at the end, that's for sure.

Julia: The topic I was really hoping to focus on today is something that you've written a lot about on *Overcoming Bias* -- in fact, there's a whole tag devoted to it on your blog -- and that's the idea of signaling.

I'm going to ask you to explain what that is in a second. But I think that the best way to communicate to our listeners, who aren't already familiar, just how linked the concepts of Robin Hanson and signaling are in the minds of your audience is to tell them of the existence of this meme, which -- Robin, hopefully you are aware of?

Robin: I am well aware!

Julia: There's this meme that was going around some corners of the Internet which has a photo of you and various words, transposed or juxtaposed with the image.

For example, one of the memes has your face and the phrase, “Most people think X is about Y. X isn't about Y.” Then there's different variations, like this one that says “School isn't about learning,” and one that says “Toasters aren't about toast.” And one that just has the phrase, “Look at you, sitting there, signaling.”

Robin, maybe you could explain why this meme exists. What do you mean when you say things like “School isn't about learning?”

Robin: Right. There's a lot of subtleties to try to get across, but first just the very basic concept:

Whenever what you do influences what other people think about you, then you're tempted to take that into account in your actions. If you comb your hair and then when people see your hair combed, they think differently about you than when your hair isn't combed, that will influence the chances that you choose to comb your hair, and how carefully and how much energy you put into it.

When you respond to those incentives and you react to how people perceive you and then they still draw conclusions from what you do, then it has to be because different people who look like you, but are different, really do different things. So that they look at an appearance and they can draw an inference about what's behind the appearance, because there really are differences behind the scenes.

One of those differences can be how easy it is for you to do something and how expensive it is for you to do something. In the classic story of, say, school -- when you look at somebody who goes to school and you think of them as smarter or more knowledgeable, or quicker, or better on the job, then they respond to those incentives. They decide they want to go to school and get a degree, so that you will think better of them.

When lots of people realize that, and they all decide, well, I guess it would be better if I went to school because people would think better of me -- not all of them achieve the same total amount of school nonetheless. So that means [00:04:00] there must really be differences behind the scenes in who is better able to go to school, or to survive the process.

It makes sense, then, to draw conclusions about people based on whether they went to school. Because in fact there are these differences behind the scenes, that make it easier for some people to actually go through the process. Signaling is just this situation where there's a connection between appearances, and conclusions you draw about appearances, and then the effort of what people do in order to make the appearances look good.

Julia: So: the signaling framework, it's not just about the fact that people draw conclusions from outward appearances, but also about the fact that we consciously or unconsciously make choices that will result in giving a certain appearance to someone, in hopes of signaling something, like that we're smart.

Robin: No one's been around the world very long without realizing that people are very aware of how other people think about them, and they're very aware of what they think would make other people think better of them. And they are in fact putting substantial efforts into doing the things that make other people think well of them.

There's no doubt that there's a lot of that going out on the world. In that sense, signaling is trivial -- of course it's happening, duh. Of course that's not what people were making fun of me for. They're making fun of me --

Julia: -- good naturedly!

Robin: Of course.

Julia: Affectionately.

Robin: Affectionately. I'm honored by getting such attention.

Julia: Yeah. I don't have any memes out there, so clearly I haven't made it yet.

Robin: You still have time. You're young. Many years you have to go to accumulate these things.

Julia: (That was a masked hint to my listeners.)

Robin: They may well pick it up. You'll have to do your part by saying something that sounds a little silly, at least.

Julia: Well, I'll work on that. Okay, Robin, you were saying that in one sense signaling seems trivial?

Robin: Right. Now if we set that aside [00:06:00] and take a bigger picture of human behavior: Just look at all the different things people do, and compare that to all the different things other animals do. Most of the things that most animals do make a lot of sense in terms of raw survival and reproduction. They eat. They hunt. They find warm places to sit. They run away from predators. They protect territory.

And humans do those things as well. But then humans do a wide range of other behaviors that are just strange from the point of view of most animals. We have an entire civilization here, of course, and some of the civilization is oriented around keeping us warm, keeping us fed, protecting us from predators and accidents, things like that.

Then there's all these other things we do that don't seem very tied to pure survival and reproduction. We do art. We do music. We tell stories. We have podcasts like this. We talk about abstract topics. We go to school. We go to the doctor. We invest money in financial hedge funds. We build enormous billion-dollar projects.

A lot of these things we do, right on the surface just seem strange. That is, why would an animal be doing that? Then it turns out for a number of the other things that we do, we have simple stories to explain why we do them, and we usually think they're adequate. In fact, those stories don't work very well at all.

The story that we go to school "to learn things" doesn't work very well, because we don't really learn very much at school. What we learn, we don't remember. And what we do remember isn't very useful.

We go to the doctor, we say, to get healthier. But in fact, there's very little correlation between people being healthy and people going to the doctor.

People invest enormous amounts in financial remediation, invest in investment funds, [00:08:00] and buy insurance, and a wide range of things. And they're spending 8% of GDP on that, and it's just actually hard to understand why most of that is useful.

We just go down the line. There's a whole bunch of things that people do, that just seem strange from the point of view of ordinary animals. And then there's a whole bunch of other things where people *think* they understand why they do them.

In fact, economists and social scientists like myself say: Actually, we don't understand why people would do any of these things. Because they don't seem to actually achieve functions they say they're trying to achieve.

Julia: Before we delve into how thoroughly the signaling hypothesis explains these choices, maybe we can just spend a little more time on why you think the standard stories fail. Because, for example, I suspect a lot of our listeners will have felt confused at your statements that we go to the doctor and healthcare doesn't actually help us very much. Can you explain more why you say that?

Robin: Well, each one of these topics we can spend an entire lecture on. They're counterintuitive, and that's partly why people enjoy making fun of the "X is not about Y" thing.

The first data point is to say school can't really be about learning. Medicine can't really be about health. Investment can't really be about returns.

Because we have all these pieces of data related to them that say, the way people seem to be trying to achieve these goals are just so inefficient and ineffective, that it's just hard to believe that this is what they're really doing.

My colleague, Bryan Caplan -- who you should have on sometime, of course -- is writing a book about education, and he'll tell you in great detail why the standard story about education just does not fly.

Julia: Well, I think people might be more board with the idea that we don't really remember most of what we learn in school. But they hear your claim about medicine, and they'll think, like, what are you talking about? Is penicillin not effective? Do vaccines not work? What do you mean? So maybe you could just say a few more sentences about that.

Robin: Of course. The first thing to say is that we're talking about the main explanation, not the entire explanation. In the case of school, of course people are learning some things in school. Some of the things they learn are useful. But most of the things they learn, they don't remember, and they're not very useful. We're looking for some main explanation for the rest of what they're doing.

Same for medicine. Of course, some parts of medicine are in fact affecting health, and some of the reason people go to the doctor is because they're sick. But that isn't very satisfactory for most of the behavior because we just have a lot of data that isn't very consistent with that.

We have data about geographic variation, for example. There are places where people just spend more on medicine, like counties or hospitals in the United States. When we look at places where they spend more on medicine, and then we look at how healthy people are there, in terms of how often they die -- controlling for age, and all the other things that we use to predict people living or dying -- we just find very little to explain in terms of the medicine.

We know a lot of things that influence health. We know there's exercise, sleep, air pollution, age, of course gender, smoking. We just know a lot of things that influence health, and people are actually not very interested in those, surprisingly. When they get to medicine, they're really obsessed with medicine. Yet we find very little correlation between medicine and health, once we have controlled for the other things.

We actually have some randomized experiments, where we took some people and paid them, basically by giving them a lower price for medicine. And the people who had a cheap medicine price, they consumed a lot more medicine, but they were not any healthier.

This is a consistent result in the literature for a long time. And it's of course a puzzle, from the point of view that people go to the doctor to be healthier.

Other related data is that people are actually remarkably uninterested in information about the quality of medicine. There have been many experiments where they offer people more information about the track record of their surgeon, or something like that, even when they're about to undergo surgery [00:12:00] and it turns out they just don't want the information. If you give it to them, they just don't use it. Again, that's a puzzle from the point of view of the "Of course people are trying to get healthier, that's why they go to the doctor" point of view.

Julia: A question that's often in the back of my mind when I read these signaling stories is: how much we can reasonably expect people to *know* that the method that they're pursuing -- to get healthier, or to get the benefits of an education -- how much do they know that that's ineffective?

It's not clear. The fact that the studies you cited are so counterintuitive indicates to me that we shouldn't actually expect people to know that what they're doing is ineffective. And that therefore it's not clear...

Robin: Let's go back to the ground, ground and basic.

We've got humans. We have human behavior. And we're trying to explain it. There are a number of very simple explanations for any kind of human behavior that people often resort to when they're looking for very simple explanations. In a sense, these more complicated explanations should be compared to the simpler explanations. And you shouldn't adopt a more complicated explanation, when the simple ones seem to work well enough.

The simplest explanation for almost anything anywhere is randomness. In fact, almost always, whatever we explain, we usually explain with some systematic theory *plus* randomness. We're always adding in some degree of randomness when we explain almost any data set we have. So one very simple explanation for anything is just to crank down the systematic part and crank up the noise and say "It's all noise, it's all random."

So you could say, well, it's a complicated world. And people don't really know how it works. Therefore, they're just acting at random, relative to what works. That's, of course, a very simple, straightforward theory. And we can apply that to everything. They can be acting at random at work, and acting at random when they drive, at mating, and everywhere else.

Another very simple common theory that is somewhat related [00:14:00] is the social conformity theory. Why do people do things? Because the people around them do the same things. So they're just doing whatever people around them

are doing. Now that theory explains some correlations. It explains why people in the same area, and time and place do a similar thing, rather than random things.

If we just have the randomness theory, the randomness theory doesn't really explain why there are these correlations. Of course, we could create new randomness theories that just put in correlations by hand.

This social conformity theory does explain conformity in some sense, that people end up doing similar things. But from the point of view of *which* similar things they do at any one time and place, it's really a randomness theory that says, "Well, there were just some random processes that ended up producing whatever people in some time and place thought was the thing to do, and just everybody does that thing now."

We're putting these right up front because these are very common theories that people appeal to often, and use to explain almost everything that goes on around them. Their main failing is that they just don't explain patterns very well. If there are any consistent tendencies, these theories don't really have a way to predict them.

Whenever we're seeing more consistent tendencies, we will have to be drawn to some other kinds of explanation to make sense of them. Because it's not enough just that people do things randomly, or that they do the same things as each other, in order to explain why they consistently, around the world, do the same thing.

Julia: I like your point about simplicity. And it reminds me of a nice quote from a philosopher named Alfred North Whitehead which you may be familiar with.

Robin: Sure.

Julia: He said the guiding motto for a natural philosopher, or a scientist, should be "Seek simplicity and distrust it." I don't know exactly what he meant by this, but my sense --

Robin: Einstein's quote was: "You should find the simplest theory, [00:16:00] but no simpler."

Julia: Yes, excellent. Okay. So my intuition about why this is good advice is at least two-fold. First, that simpler explanations are more useful, because they have more predictive bang for their buck.

But also because in the process of forming simple theories, we make it easier to notice when those theories are wrong. So we can't just keep on adding

parameters and parameters, and epicycles and epicycles, and guard ourselves against falsification, or against contrary evidence.

That's where the "distrust simplicity" part of the advice comes in. What I'm genuinely confused about is: the signaling framework does seem to be seeking simplicity, but where does the distrust come in? What kind of evidence would make you accept that, okay, the signaling theory is too simplistic in this case, it doesn't actually explain what's going on?

Robin: Well, the theories will make some concrete predictions. And you want to be comparing the actual predictions with what the theory predicts. And you want to be asking how many free parameters you're allowing yourself. I think we can apply those, but I think we have to explain the theory a bit more before we can explain to listeners how we're going to do that thing.

One empirical test to keep in mind, straightforwardly for signaling theories, is: Do people draw conclusions about other people on the basis of these behaviors?

Julia: Yes.

Robin: And can people do anything to influence those appearances? If people draw conclusions about what they see, and people can influence those appearances, then there has to be some element of signaling going on. If those two parts are coherent with each other, people are actually [00:18:00] somewhat knowledgeable about how people see them, and somewhat knowledgeable about how to influence those appearances --

Julia: -- And care about how people see them.

Robin: Right. And if other people have some idea of what people can do to influence the appearances, then we're in the signaling world. It's a more a matter of how *much* signaling is going on, and at what level -- not so much, is there *any* signaling going on. There has to be signaling going on in that case.

I think for an awful lot of these behaviors, it's really easy to check and assure yourself. Just ask people around you some random questions, to see people do in fact draw conclusions about other people from any of these behaviors.

We can set ourselves some hypotheticals here, if you'd like, when we get to that point. Just walk through what conclusions people draw. So that's one set of data that we can use to check signaling theories: what conclusions do people draw, and what efforts are people doing? What appearances do people prefer to have, and how much and what kinds of effort are they putting in to look better?



The other broader context that I started to allude to before is to step back and just ask: What kinds of behaviors would you *expect* humans to be doing in the first place? Just on an a priori level. What are the kinds of behaviors you expect to have evolved?

Now some people, yeah, even some of my colleagues at lunch today, favor the randomness theory of, "Humans are so complicated and the modern world is so at variance with our ancestral environment, that you should just shouldn't be able to expect, or predict, much of anything that people do. Because it's just all so complicated."

There's some sense to the idea that our world is complicated and we should be conscious about being overconfident there. Nevertheless, I think we know some things about the kinds of things we expect our ancestors would have evolved to do, and therefore, some of the ways those would have been projected into the modern world.

Just to review some basics about human history and human nature, chimpanzees and other primates are very social creatures. [00:20:00] They have very large brains, compared to other animals, because they are very social. They have very complicated social lives, in which most of their environment that matters are other creatures like themselves.

And they're paying a lot of cognitive energy and attention to thinking about what other people are doing, what other people think about them. Thinking about and creating alliances, deceiving each other. And they basically create these political coalitions that support each other, and try to take out other coalitions.

This is the social brain theory, which I think is roughly correct. That in fact, for our ancestors, the most important thing in their environment was other creatures like themselves, their social world. And most of their brain was actually devoted to the social world around them.

Then there's another standard theory about how humans differed from other primates. And the key differences between humans and other primates according to this story (you can see the book *Hierarchy in the Forest* by Christopher Boehm) is that humans had language and tools, and that allowed coalitions of large groups against anyone who would try to set themselves up above.

So they had strong egalitarian norms. They had egalitarian norms for sharing, against bragging, against dominance. And these norms are enforced by the fact that we could watch each other's behavior. And that if we saw somebody violating the norm, we could say it in words. We could tell other people. They

could exchange it, they could verify it by watching. We could coordinate to talk about what to do, how to respond to this norm violation, what punishment is appropriate.

And we could implement that punishment, because we had tools, and everybody has to sleep. And everyone who's sleeping is vulnerable.

Chimpanzees, basically, since their tools are pretty weak, if you punch one while they're asleep, you got the first punch in, but that's all. And then there's going to be a long fight. For humans, our tools are strong enough that you can kill somebody in their sleep. [00:22:00]

We have tools that act from a distance. So that even when they're awake, a large group of us could stand off at a distance and kill somebody. That meant that core groups of humans could in fact kill other humans, when they thought they were violating norms enough.

The bottom line is humans had this very different world of norms and enforcing norms, but we also had this strong social background of being strongly social creatures with very complicated reasoning about each other.

In this context, your prior about what strange human behavior should be about, should be a social prior. You should be thinking that most behaviors that you don't know what they're about probably have something to do with social interactions with other humans. Not all social interaction is signaling, of course, but a lot of social interaction is about not just coordinating to do things, but deciding what we believe, and what to do to who.

In coalition politics, in other primates and in humans, you have to decide who can I trust? Who are my friends? Who are good friends, who are powerful friends? Who would be useful to have as friends, and who are pretending to be loyal to me but actually would betray me? Who are betraying me behind my back? What's going on that I can't see? Are those two people whispering, are they looking at me, are they talking about me, et cetera.

I think we should just have strong priors that these sorts of complicated social strategies are a lot of what we should expect strange behaviors to be about. Certainly a big component of a lot of those social behaviors is looking at people and trying to draw our conclusions about them -- is this person a good mate, is this person a good ally, is this person loyal to me?

Whenever you're trying to draw those information inferences from things you see, then the flip side of that is people trying to do things to make themselves look good to you. [00:24:00] They're trying to decide what would make myself

seem loyal to you. What would make me seem like a good mate. What would make me seem like a good ally?

Julia: I agree that we should have a fair amount of weight on the prior that we do things to make other people think well of us, and respect us, and form coalitions, and everything that you said. I'm just confused about how to weight that prior against other priors that I think we also have good reason to put weight on.

In the medicine case, for example, when you said people are surprisingly unwilling to pay for information about whether the hospital that they're at will have a higher chance of successfully completing their surgery than a different hospital...

You would think would be information that they would want to know if they were actually purchasing the surgery in hopes of having a successful outcome. That is a rational decision making strategy, which I agree with.

But cognitive science and just our experience as humans living in the world, has generated a bunch of other, I think pretty solid, facts which should be part of our prior. Like the fact that humans are bad at probability and are pretty scope insensitive, and don't really feel the difference between a 5% chance of failure versus an 8% chance of failure. Also the fact that humans are superstitious thinkers, that on some level, it feels like if we don't think about risks, they can't hurt us, or something like that.

Not even knowing the result of that study, if I think about a priori, what I would expect to see... I would have put (it feels like anyway, hindsight bias, et cetera) It feels like I would have put a significant amount of weight even in the absence of signaling caring, that people would fail [00:26:00] to purchase that useful information because of these facts that we now have about the human brain.

Robin: One of the things we're obviously we're going to do in data analysis here is do comparative analysis. We know a lot of different industries in the world. People buy cars and repair them. They buy houses and repair them. They have human bodies and they repair them. We can compare people's behavior across different industries and we can look for differences to explain.

Any explanation that's just about generic human psychology is an explanation that would apply equally to all these different industries. When we see different behaviors in different industries, these common explanations aren't going to work very well, unless they're tied to some other difference between those industries that we can point to.

As an economist, I've spent a lot of time studying medicine and how it's different from lots of other industries in the world. Those differences are the starting point for these explanatory exercises. You're looking for theories that can explain some of these differences.

We typically get our car repaired at auto repair shops, and in fact there is a substantial correlation between cars being repaired and people paying for car repairs. Similarly for housing repairs. It's not that people just are so bad at knowing how to buy repairs. There's no correlation between the pair of expenditures and general and things getting repaired in general.

Medicine is different from these other areas. We have something different to explain there. I think also, in many other industries, we do pay a lot of attention to quality signals and much more than in medicine.

It's not just that we are just bad at paying attention to quality signals in general, and that therefore, we just can't figure out how to pay attention to quality. We act differently about medicine than we do in other areas where we also have difficulty understanding things.

Not to say that the simplest model goes well. I think in most areas we do spend [00:28:00] less on information about quality than you might have expected. That's a wide puzzle. But we still have these big differences between areas.

We can go through many other sets of ways in which medicine is different from other kinds of behaviors. And of course, if you introspect, I'm sure you can realize that we do treat medicine quite differently. We're much more eager to regulate medicine. We're much more eager to get medicine purchased communally, and through communal channels. We are much more eager to get medicine associated with prestigious sources. These are all ways in which medicine is different.

Now, there's of course a whole academic industry of trying to explain these ways in which medicine is different. I suggest a particular way of explaining it. The same is true for education. As Brian can tell you: Education is different from many other kinds of industries that you might think it would be similar to, and those are the puzzles that we try to explain about education, why people treat education differently than others in other industries.

Julia: Actually, if we can just go back to medicine real quick. I realized I'm not sure either of us explicitly stated what the signaling hypothesis *is* in the case of medicine. We explained how it's not about health, but what is it about?

Robin: Right, no, we didn't. So one plausible hypothesis -- it's a tentative hypothesis -- is the idea that medicine evolved as a way to show that we care. That is, we are reassured when other people spend money on medicine to take care of us. We like to reassure people around us that we care about them, via spending money on medicine to take care of them.

They respond positively to that reassurance. There's what's called the placebo effect, whereby people just, in a sense, seem to get better directly as a result of the appearance of care, even when there's not actually a physical effect, a direct physical effect.

The idea [00:30:00] is that our distant ancestors, of course, needed to take care of each other. They had things that went wrong, and they needed some help from each other. Because there was all those complicated coalition politics, they were afraid -- rightly so -- that the moment they got injured or sick was exactly when their supposed allies would betray them and abandon them. Because that in fact is the right time to do so, if you're going to do it. They are weak and much more expensive.

Being sick or injured is usually something that took a while to get better from, and so your friends or allies would have to actually use a lot of resources on you to take care of you, and they should only really do that if they expect you to stay as a friend, an ally to them for a long time. So it became a credible signal of a long term intention of allegiance, to take care of somebody when they're sick or injured.

Because of that, we were all eager to send that signal. To reassure allies that we are in fact intending to stay with them. And that we intend to take care of them and be their allies for a long time. We were eager to send those signals of caring and attention, and when we did so, they responded exactly as we hoped, by feeling reassured and cared for and loved.

We all even say, tears come to our eyes, just imagining the scenario of ourselves being sick or injured, and our friends and allies coming to our aid and taking care of us. It's very touching. Rightly so -- because that's exactly the sort of thing we needed the signal of.

This whole story can be consistent with medicine actually working a lot. You could have medicine that was very useful, and then a layering onto it this reassurance of help. That when medicine isn't in fact useful, or it becomes much less useful, you still might continue on with the same reassurance behavior even if the medicine itself wasn't as effective.

As we've gotten rich in the last few centuries we've actually spent a lot more on medicine as a percentage of income than we used to. You might think that the first amount we're spending was useful. But after a while we have more useful things to spend it on, but we just kept spending anyway. [00:32:00]

Julia: What about medical care that people purchase for themselves? How is that signaling?

Robin: Well, you've heard of Valentine's Day.

Julia: Rings a bell, yes.

Robin: There's a tradition of giving somebody a gift of chocolates or roses or something on Valentine's Day. Did you know that some people who do not get gifts of chocolates or roses buy it for themselves on Valentine's Day? Why do they do that? Well, can you imagine why that would feel nice?

Julia: I can imagine why it would feel sad.

Robin: It also could feel better than the alternative of not getting anything.

Julia: Yeah...?

Robin: Okay. You want the appearance of being cared of. Especially, you might do this if you could fool your co-workers into thinking that somebody had bought them for you.

Julia: *That* I get.

Robin: There you go. Okay. You want the appearance of people caring for you, even if you have to pay for it yourself. You want to --

Julia: But surely, our drive to survive, and avoid pain, is just as evolutionarily significant as our drive to signal caring. Why wouldn't we --

Robin: Well, there are actually interesting theories about pain itself. I was just reading an article about ways in which pain can be involved in something like signaling.

For example, most other animals do not have very much pain during birth compared to humans. Human females of course suffer far more pain than other animals during birth, yet in fact the physical damage and risks are not substantially larger in human females than in many other animals where they feel much less pain.

One of the stories which seems plausible is that human female pain during birth is there in order to induce other people to come take care of the woman, to help her give birth. Since humans have a lot of helpers along at birth, traditionally, then this is a way to make [00:34:00] that happen. It's a credible signal of needing and wanting to help this pain. There are senses in which our pain could be produced via an evolutionary process that has signaling in mind.

This brings up an issue I wanted to make sure I mention, because some people have brought it up several times. Which is that the signaling theory isn't necessarily a theory about your conscious intentions or plans.

Julia: Right.

Robin: There's a whole chain of causation between your evolutionary heritage and your doing any one thing. Like going to school, or going to the doctor, or having a chocolate on Valentine's. Signaling theories can apply at many different levels here.

For example, you might consciously decide, I'm going to go to school so I will look good. But it didn't have to be that way. Some young men decide they consciously want to be a rock star because it'll attract women. And many young men do that, but of course many other young men decide they want to be a rock star because they love rock music.

Now, either way works evolutionarily. It's the behavior that produces the outcomes and not necessarily your rationalization. For some kinds of behavior, evolution can give you a conscious desire to be seen, to look good. And then you consciously make a plan to achieve that looking good.

For other kinds of behaviors, evolution may just give you a desire for something. Like the comfort of people taking care of you. Or just the joy of being a rock musician. And you may not realize its evolutionary functions, but some of those functions can be signaling.

I don't specifically mean to accuse people of consciously planning these things. I think that's part of the resistance to the signaling hypothesis. That it basically says that we're manipulative, and overly concerned with appearance. And most of us like to think that, even though sometimes we pay attention to appearance, mostly we're driven by [00:36:00] our basic desires for liking certain kinds of things.

We like to think we don't really care what other people think. That even though we do care what other people think, we'd still do the same thing, even if they

didn't like it. Because, by god, we're just that principled, dedicated person liking the things we like.

Julia: Right, so in the case of buying medical care to signal caring: Probably most people would say that what it feels like on the inside to buy medical care to try to save their dying mother's life, is that they *actually* care. They actually want her to continue living.

And I could imagine a plausible evolutionary explanation for this. That we care about people who share our genetic material. Because genes that had that property, or caused that property, were more likely to propagate than genes that didn't. So does your theory distinguish between genuinely caring, versus wanting to show caring, or are those both part of your theory?

Robin: They're both a part. The fact of the matter is there's in fact a big complicated inter-mixed land between our evolutionary motives and our conscious attention.

We're even thinking about school. Some people just are very conscious about wanting to go to school because they want the degree and they don't really care about what they want. Other people think they want to learn. Other people are just mixed-up, in between, and on Tuesdays they're one way, and on Wednesdays the other way.

Same way for lots of these other sorts of motivations. Even in the medical case, you take people who think they care about their grandmother, you explain to them, "Look, most end of life care is really not very helpful, and puts them in more pain, so just quit and send them to a hospice," and it's really hard to do that.

It's in part because they realize they might be accused of not caring. That is, there's often this suspicion and concern that the reason you're not spending more on sending them to regular doctors is because that costs more. [00:38:00] That the real reason you want to go to hospice is because it's cheaper. People are very actually conscious, and concerned, about that accusation. And that actually drives a lot of people to continue to push their parents into more painful, less healthy but more expensive, care at the end of life.

Julia: That does ring true. That actually matches the "feeling on the inside" explanation that I was talking about.

Robin: You have to look inside to see that feeling. On the surface of your mind you say to yourself, "Of course I just care about grandma." The question is where are your real feelings? Or where is your real motivation? And that's actually spread out all over the place.



Julia: What types of things do you think have the biggest gap between people's feelings about their motivations, and the strength of the signaling explanation?

Robin: Well, I think the basic fact is, the basic puzzle is we do a lot of things that we don't really know why. It's perfectly plausible to understand why creatures like humans might signal a lot, because we're very social, there's lots of things to signal. The puzzle might be why don't we know that we're signaling.

The basic explanation there is there are some motives that are high, and some motives that are low. And we like to try to appear to have high motives rather than low. There are many ways in which we have actually low motives, motives that we don't like to admit to, but that are really there. But we prefer to substitute high motives. And also for these social norms, we prefer to act like we're following the usual social norms even when we aren't, and we're violating them.

So the difference between what we think we're doing and why we're doing it, and why we're really doing, is highest exactly when our real motives are low and things we don't want to admit to, and the motives we're pretending to have are great and noble, and the sort of things that everybody admires.

Julia: Because it's easier to give the appearance of someone with noble motives if you actually have -- or you've actually [00:40:00] have deceived yourself in thinking you have noble motives.

Robin: Right. We all leak in so many different ways. At some level ...

Julia: Figuratively speaking.

Robin: Right. Our motives leak, our feelings leak in many ways, we just have all these different channels by which our feelings come out. In our eyes, and our tone of voice, and our tremors and all sorts of things. So if somewhere deeper inside, we have a sincere belief in something, then that makes it look more consistent, like that's our motive. But of course we have many, many layers below those layers which can be really driving what we do.

Julia: We're almost out of time but I just can't resist: Speaking of people leaking, there is a quote that I posted on my Facebook wall, I guess last October. It was from the eminent statistician Andrew Gelman. He was actually complaining about economic models in general. He said it's as if you went into a bathroom in a bar and saw a guy pissing on his shoes, and instead of thinking he has some problem with his aim, you suppose he has a positive utility for getting his shoes wet.

Robin: I think that brings us back to the randomness theory. If you just saw one guy pissing on his shoe, yeah, you should think that. If you consistently go in the bathroom and they are consistently all pissing on their left shoe, then I do think you have to start to wonder if something else is going on.

Julia: Yeah, that's what I was going to ask you, is what you think the actual parallel is to the signaling case? It's still not clear to me that the pissing on the shoe is systematic and reliable enough, consistent to what we would expect from signaling, as opposed to other things.

Robin: The less you know about human behavior, the more you can maintain a simple randomness theory that everything is random. I think as you learn more about what people are actually doing, you actually tend to see a lot of very consistent patterns on what they do, and that forces you to come up with theories that are capable of explaining consistent patterns, and randomness [00:42:00] and even social conformity just don't do very well for those.

Julia: Well, I think you're using the word randomness to refer to a whole collection of things. The thing that you're calling randomness, I would call a whole collection of tendencies and cognitive biases and heuristics. So they're just a more complex explanation. Which is unfortunate, they just seems like that's in fact what we have on our hands scientifically.

Robin: Of course, that's what all randomness is really. Complexity.

You see a pile of rocks on the ground. There were very specific forces that put each rock there in its place. If you don't know what those are, you tend to summarize it in a simple randomness theory. Which is adequate if you don't know those details.

Similarly, if people's decisions are made for a thousand little biases and a thousand little cognitive errors and hacks -- if those aren't in any particular pattern relative to each other, or structure, then the net effect to them will tend to be random. Unless somewhere a systematic process comes in and makes it be systematic.

Julia: Right. Interesting. Okay, well, I'm sure we could have 10 more episodes of signaling alone --

Robin: Okay, we'll do it! I'm ready.

Julia: -- but we're out of time. So let's wrap up this section of the podcast and move on to the Rationally Speaking picks.

[music]

Julia: Welcome back. Every episode, we invite our guest to pick the Rationally Speaking pick of the episode: A book, or website, or a movie, or something else that tickles his or her rational fancy. Robin, what's your pick?

Robin: I think I have to give a nod to Geoffrey Miller's *The Mating Mind*. It was a book that very much influenced me 15 years ago when I was first trying to make sense of a wide range of strange things people were doing. [00:44:00] He offered a grand ambitious hypothesis that tried to explain a wide range of puzzling human behavior, and that was inspiring to me to try to take on that level of ambition and scope.

In the end, I'm not sure I'm that persuaded by the particular hypothesis he came up with. It has some plausibility. But I was more just drawn to the attempt and the spirit of it, of trying to make sense of a lot of strange things people do. Just realizing that a lot of what people do is strange, that's right on the face of it. Once you see it, it's kind of obvious, but we just miss it unless somebody points it out.

Julia: Great. Well, we will link to *The Mating Mind* on the podcast website. And I encourage all of our listeners to check out Robin's blog *Overcoming Bias* where you can read all about the signaling theory as well as many, many other interesting topics.

Also keep an eye out for a forthcoming book. Robin do you know when your book on signaling will come out?

Robin: I'll have a book on signaling coming out in the next two years, but I'm afraid I have to wait a little bit before I can tell you the title.

Julia: Okay. Great. Well, keep an eye out, I'm sure it'll be fascinating.

Robin, thank you so much for joining us. I really enjoyed this conversation and found it pretty enlightening.

Robin: It's great talking to you Julia. [00:46:00]

Julia: All right. Well, this concludes another episode of Rationally Speaking. Join us next time for more explorations on the borderlands between reason and nonsense.