Rationally Speaking #235: Tage Rai on "Why people think their violence is morally justified"

Julia Galef: Welcome to Rationally Speaking, the podcast where we explore the borderlands between reason and nonsense. I'm your host, Julia Galef, and my guest today is Tage Rai.

                 Tage is a social sciences editor at Science Magazine and a research associate at the Sloan School of Management at MIT. His background is in psychology and he's the co-author, with anthropologist Alan Fiske, of a recent book titled 'Virtuous Violence'.

                 Which, even if you haven't come across the book, you might have seen some of the articles that have been written about it in the last few years. The thesis basically is that most acts of violence that people commit are motivated by moral feelings; that people feel, when they commit violence, most of the time, that what they're doing is defending morality. Their violence is righteous, basically.

                 And this model of violence seems to conflict interestingly with a lot of other common sense notions of violence and academic theories of violence as well. So that's what we're going to be talking about today.

                 Tage, welcome to Rationally Speaking.

Tage Rai: Thanks for having me.

Julia Galef: Tage, what was your prior going into doing this research? What was your starting rough model, or assumption, about what causes violence?

Tage Rai: It depends on what time of your life you're talking about, but if you grow up in a relatively peaceful part of America or something, it doesn't seem like all violence is particularly morally motivated.

                 But then you start to realize ... I moved around a lot, once I started to go to college and stuff, and you start to realize that people's attitudes about right and wrong are actually way more different than I would've guessed. And a lot of the things where I

thought, "Yeah. I'm really pissed off and angry and I kind of do want to hurt somebody," it turned out ... I mean, I'm from the South. It turned out that a lot of my friends from the North didn't have any feelings like that at all...

My background research-wise was: I was studying moral psychology, and you start to see the sort of diversity and breadth of moral practices around the world, across cultures, throughout history. And it just really became clear that we see a very very narrow slice of that in sort of Western liberal academic America.

Julia Galef: Well, I guess I'm wondering whether when you and Alan launched this research project together, did you start off with the suspicion that morality was going to be the root ... the common thread? Or was it purely exploratory - you didn't know what you were going to find, you just wanted to see if there was some pattern that had been missed?

Tage Rai: So before this project I had already sort of made the argument that there's probably a lot more violence that's morally motivated than we'd like to acknowledge or think about. But I don't think we really would have guessed the extent that we found. That we'd see it everywhere, that perpetrators would sort of readily admit it to the degree that they do, that sort of thing.

Julia Galef: Got it. And this was qualitative research, right? You weren't running randomized controlled trials?

Tage Rai: No, no. So the book is almost all scholarly. So we're drawing on other sources. First person accounts, historical documents, ethnographies, newspaper accounts, everything like that. We do some interviews ourselves on some of the topics, but the majority of it's scholarly. There's no quantitative experimental data that we did ourselves.

I did some experimental work [later], but that's what I've been doing the last couple of years.

Julia Galef: Okay, let's dive into some concrete examples. Could you give one or two examples of violence that you studied, that you wouldn't necessarily have expected to find would be morally motivated, before you started? But that now seem quite plausibly moral to you?

Tage Rai: Sure. So, I think examples that I would come up with would be… intimate partner violence might be a good example. And tied to that, things like sexual assault.

                  So you know, I think a standard approach in the literature might have considered that, "Well, really what's driving that is an instrumental desire for sexual satisfaction," or something. And in fact you might think that if the perpetrator doesn't even necessarily recognize this victim as a human being and therefore they're morally disengaged from the act, they're just trying to get sexual satisfaction or something like that.

                  Instead, that's actually not what we find at all. What we find is that perpetrators really do ascribe a lot of mental and emotional states to their victims, a lot of moral considerations into their actions.

                  And what they're trying to do is rectify what they see as a violated social relationship between them and the victim. Or between them and the victim's social category in general. So these men are trying to get back at women, or they're trying to get back at this particular woman. And they're trying to create what they believe to be the morally correct relationship between the two individuals.

                  And if you took that away, actually, if you did sort of strip away the humanness of the victim then actually the satisfaction of it would go away too, so the morality is totally tied to the act.

Julia Galef: How do you know that if you took away the humanity, from the perpetrator's perspective, that their desire to commit the act would go down? Is that a ... You're inferring that, or you've observed that?

Tage Rai: I think we'd have to infer that, from the data in the book. So what we can look at is what the perpetrators are saying and doing. And then you know, we have this separate paper in PNAS that was an experimental paper, looking at what drives perpetrators if they're ascribing ... If we kind of manipulate the levels of humanity and mental states and things like that, that perpetrators ascribe to people, does that increase or decrease

aggression? And it's supportive of these kinds of ideas -- but yeah, you'd have to extrapolate it, so I over-spoke there.

Tage Rai: Another example, something where I really wouldn't have guessed, would be something like robbery, okay?

Where physical violence and assault ends up getting involved. It turns out that when we look at robbery... here's something where if I was really just thinking about it from an instrumental point of view I might have thought, "Oh, well, I should be robbing strangers out of convenience," let's say. So whoever's the easiest target, and I don't know them or anything like that.

And that's just not really what we see. What we see is that actually a lot of times the robbery victims are people that you know. Which, there could be multiple reasons for that. But even when they're not, there are moral motives that come into play.

Both in terms of kind of what perpetrators do -- so, they use more violence than is necessary. And then they also claim that they're really trying to get back at this particular social group, and also to establish dominance over that group. And really again, get across what they believe to be the morally correct relationship.

Trying to reestablish that in some way. Establish some sort of equity, some sort of fairness over a perceived injustice.

Julia Galef: So I'm sure you're familiar with the research about how people come up with justifications, virtuous-sounding justifications for self-interested behavior.

Like... Well, just in experiments, people who are given an opportunity to do something selfish, like take more than their fair share of some pool of money, will come up with a definition of fair that makes that okay. Or they'll denigrate the person that they're taking the money away from, so that they don't have to feel guilty. Stuff like that.

And maybe the conscious reason that they feel they have for that act is moral. But clearly that reason was not causally upstream of their act. Because we have this randomized experiment showing that people's definitions of morality, or their view of their

partner, depends on whether they have the option to cheat that partner, basically.

And that is easier to prove in experiments. But there seem to be so many real world cases that I would put in that same category. Like, well, slavery, you know? Slave holders in the South would claim that slavery was good for the slaves, right? They weren't smart enough to run their own lives so, "They need us to take care of them!" And I'm just pretty dubious that that reason preceded slavery, rather than the other way around.

So ... I mean, my question is: Do you have any reason to doubt that the moral justifications that robbers or domestic abusers give for their actions are rationalizations, as opposed to the actual causes of their acts?

Tage Rai: Yeah, I mean, I think there are a few ways to respond to this.

The most basic I usually think about is: In the absence of any other evidence then what should we rely on? And I think in the absence of really tightly controlled experiments on violent crime and real physical aggression, then we should as a starting point use what people are saying.

The second, though -- I would say that it's true that I can't really get inside the mind of a perpetrator. I can only go off of what they say and what other people say and their actions. But if I were to think about it and I were to think, well, usually when we say, "Hey, that's just a post hoc justification," the implication is that it's an excuse, correct?

Julia Galef: Right, yeah.

Tage Rai: That's meant to excuse the behavior. And what I would say is if the goal of someone is to excuse their behavior and to mitigate blame, moral justifications are really bad in this context, right? Violent perpetrators shouldn't be making moral claims. They should be saying, "Oh, I didn't do it. It was an accident. I wouldn't do it again. I'm sorry."

That's not what we see. We see they're owning their actions and saying, "Yes, I did it. I'd do it again. That person deserved it." These are all things that are going to make their situation worse

not better. And so the underlying logic for the post hoc justification isn't totally clear to me.

And then the third thing is, again, absolutely true, I can't get inside the mind of a perpetrator. But if the purpose of a post hoc justification is to appeal to observers, then I think I am learning something about the moral attitudes of the social group.

And so that is very interesting and usable data. It's telling me something about what kinds of moral standards other people have, such that these kinds of moral justifications would actually help conceivably.

Julia Galef: Yeah. Interesting…

Tage Rai: Does that make sense? So, I might not learn about that particular individual, but I do think I learn about the social group, through looking at the justifications that the individual perpetrators provide.

Julia Galef: Yeah. So I think you learn something about the social group, and what kinds of justifications would be acceptable... I think my view, as of now, is that we learn less than you might think we learn. For a couple of reasons.

One, because the perpetrator might be able to go through a bunch of kind of contorted rationalizations, to make their situation fit a legitimate reason. So if you were honestly applying your society's or your subculture's rules, about what counts as a morally justified act of violence, then it would apply to a much smaller set of instances than the actual set of instances people try to apply it to.

And then the other thing I was going to say is that ... I've been reading a lot about signaling lately. Where we're motivated to adopt beliefs or preferences, aesthetic tastes even, basically just because we unconsciously want to influence the way other people see us. So we're motivated to adopt sophisticated beliefs, because we want people to see us as sophisticated, that kind of thing.

And the goal in most of those instances of signaling is not to convince other people that the belief you hold is true. It's to convince them that you believe it's true.

So if that's what's happening here, then I, the perpetrator of a violent act, don't necessarily need to convince you that you should share my justification for what I did. I just need to convince you that I genuinely believe my own justification. So that you believe that I didn't act knowingly immorally.

Does that make sense?

Tage Rai: Yes. I think so. I would still push back and suggest that if I wanted to reduce the likelihood of me getting in trouble or something like that, it would be better for me to signal to you that I had no idea that it was wrong, or that it was an accident, or that I would never do it again. As opposed to signaling to you that, "Oh, I was aware of this and I was committed to it."

Obviously the worst signal would be for me to signal to you, "Yes, I knew it was wrong and I did it anyway, because I'm a psychopath."

Julia Galef: Right. That's not a very frequently sent signal! Or, intentionally-sent signal.

Tage Rai: Right. That's not a signal that gets sent.

And it's actually a really interesting sort of philosophical legal question. Because the law is sort of based as if that was the signal, right? The concept of mens rea is this idea that, "Oh, I knew it was wrong and I did it anyway." That's not really what we see. What we see is that people say, "I don't think what I did was wrong and I did it."

Julia Galef: Well, I don't know much about the law or the definition of mens rea but -- is it more about, the person has to know that their act was wrong by their standards? Or that they have to know that the act was wrong by society's standards?

Tage Rai: That's true. You're right... So I would only make the claim that a perpetrator is trying to appeal to their sort of local peer group, not to some sort of national culture. And oftentimes there's

going to be disagreement between the overarching national group and the local group.

But I think, stepping back -- I'm totally willing to acknowledge that relying on perpetrator accounts of this sort introduces all sorts of potential biases, and it has to be taken with a major grain of salt. And so the claim that most violence is motivated by moral sentiments is not one that I want to pick as the hill that I would die on, right?

But what I would hope is that we could agree that there's going to be some sizable chunk of violence. Whether that chunk is 80% or 20% isn't that important to me. Or whether it's 51 or 49 isn't that important to me.

But as long as we can agree that there's some sizable chunk of violence that is motivated by moral sentiments, then we're left with the question of, "Well, how do we explain that? Do our theories really do a good job of capturing that kind of violence?" And I would argue that they don't.

Julia Galef: Good, yes. Okay. So this is getting to the next question I wanted to ask you, which is: what are the other leading theories to explain violence that you and Alan are diverging from?

Tage Rai: Sure. I would say if we try and think about the literature on just aggression generally, the sorts of explanations that you find for why someone would hurt or kill another person are that, well, they didn't feel the right emotions. Or they couldn't control their emotions. Or they were a psychopath. Or they didn't see their victim as a fellow human being, or they morally disengaged from their act. Or they knew it was wrong, but they gave into some sort of selfish temptation.

You know, you can block these sort of explanations under things like instrumental theories of aggression, or impulse control theories of aggression. Or the sort of disengagement, psychological-type theories of aggression.

But what I would argue is implicit in all of those approaches is this basic core assumption, which is that violence happens when something went wrong. There was some sort of breakdown in the psychology of the perpetrator. They didn't feel the right

things, they couldn't control themselves, they didn't know what they were doing.

And I think from a certain kind of Western liberal point of view that might make intuitive sense, on introspection -- if we were to think, "Well, what would lead me to kill somebody? Oh, it must be I went crazy or something."

But as somebody who is studying moral practices across cultures, that really wasn't satisfying. It wasn't capturing what we saw, where you have cases where you would have to assume that everybody's a psychopath, or you would have to assume that everything a person is saying they don't believe.

If you were to apply a claim of moral disengagement, would it really make sense with all of the felt emotions that the person was expressing?

And so it became this question of, well, we have a lot of theories about how inhibitions against engaging in violence break down, thereby allowing violence to break through… we don't really have a lot of work on the other side of the equation. What makes you want to be violent in the first place? And that was really what we wanted to get at.

And then the question becomes, I think: If you shift your perspective towards seeing violence as not about moral failure but actually about moral activity, moral performance, then how does that change the relationships between violence and various kinds of psychological processes like dehumanization, like self control, like rationality?

Julia Galef: Okay. So maybe the best approach would be to go through some of the patterns that we see in the data, or patterns in the world, and talk about how they fit with the virtuous violence theory?

So, one pattern ... And feel free to tell me I'm wrong about the existence of any of these patterns. I'm going off of common wisdom here. But one pattern is that as people get wealthier they're less inclined to commit a lot of kinds of violent crime, like armed robbery. Presumably because they have less to gain from acquiring additional wealth. Maybe they have more to lose, their

current life is pretty good now. And I think this is true at the individual level and at kind of the society level over time.

Doesn't that pattern fit better with a rational or instrumental theory of violence than a morally motivated theory?

Tage Rai: So, I think what we can say is if you're coming from an instrumental view, what you're going to say is, "Hey, look. It turns out that as we reduce the cost to engaging in violence and increase the benefits, then you get more violence," or something like that.

But then you're faced with this problem that so much violence is clearly not beneficial to the perpetrator. You know, wars are on both sides. People get caught a lot, and people who engage in violence face huge material consequences for their actions, oftentimes consequences that they clearly saw going in.

And so the question is: Yeah, it might be true that if we reduce the costs and increase the benefits, violence goes up. But it's also true that if we look at the overall cost-benefit, it's really negative in a lot of cases. And yet people are still engaging in violence.

Julia Galef: Yeah, I guess I should've disambiguated earlier between a rational theory of violence, where people are acting in what they perceive to be their own self interest, and... what I might call an "adaptive" theory of violence. Where we have drives to act in ways that maybe were in the self interest of our ancestors, or in the self interest of our genes in the ancestral environment. But are not necessarily in our rational self interest now, as individuals in the year 2019.

So maybe for our ancestors, if somebody disrespected us or pushed us or something, the rational thing to do would be to fight back, and defend our status and honor or something. And so we have that drive now. Or men, more so, have that drive now. And it's not actually the most rational move in today's world, because you can go to jail and it doesn't actually matter if a stranger in a bar disrespects you. But we still have these drives and so we can't help but act on them.

Is that an alternate theory to yours, or is that consistent with yours?

Tage Rai:	No, no, that's consistent, and I'm glad you pointed that out. Because I think that captures something I should have made clear. What I'm really talking about is proximate motivational psychology – so, what is the person feeling? The assumption is that there are ultimate causes to this, that are driven by adaptive reasons. But that in the proximate psychology of the perpetrator they're not necessarily doing a cost-benefit calculation of their actions. Which, the instrumental approaches to aggression tend to assume that that's what's going on.
	But I don't think anything I'm arguing for here is out of line with standard evolutionary approaches to cooperation, for example, or third party punishment, or something along those lines.
Julia Galef:	Now that we're talking about evolutionary patterns: there's the pattern that animals, especially male animals, seem to have a built-in drive to dominate each other. Or to try to dominate each other. Which would make sense evolutionarily, because that's how they get more access to food and mates and so on.
	So it seems reasonable to surmise that humans also have such a drive. And if so, then shouldn't we just expect that to be the cause of a lot of violence, just a priori? Maybe not all, but shouldn't that be our dominant theory, that we have a drive to dominate each other?
Tage Rai:	Yeah, and there is work making exactly that sort of argument. And I'm totally down with that work.
	I think where I differ from a lot of moral psychologists is that I have a really much broader view of morality than a lot of other psychologists. And so I think those behaviors are things that I would count as sort of proto-moral things. That what we do is just sort of extensions of that. Built in morals.
Julia Galef:	So, yeah -- what someone else might call a "desire to dominate" you might call a "sense that you're morally entitled to dominate"?
Tage Rai:	Yeah, I would say that people are motivated toward authority relations. And they find hierarchy rewarding, and that they think it's natural and good and just, and they feel a motivation to maintain that.

Julia Galef: I see. Well, in that case, do you think ... How much of this apparent disagreement in models of violence is a real disagreement? And how much of it do you think is having different definitions, for morality versus self interest?

Tage Rai: Yeah, so I try as best I can to really avoid semantic arguments. I find it super frustrating when people are like, "Oh, but that's not moral. Or that's not whatever," and you know, "That's not dehumanization, or whatever that is, right?"

                And I tend to just want to be like: Okay, that's fine, but then let's get concrete about the particular psychological processes and I don't care what we call it.

                I think in this space, it's not so much that a lot of the theories are absolutely incompatible. I think in some cases they sort of are. So if we think about the instrumental ones, it's really hard to come up with a theory that is incompatible with an instrumental theory that says you do things that are good for you in some broad sense. Instead it comes down to, what are the parameters that people are actually assuming?

                And traditionally in instrumental theories of aggression, they tend to assume that those parameters are material utilities. So, material costs and benefits in the world. But actually a lot of the utilities are non-material, social and moral utilities. And if you think of it that way, then my work is perfectly compatible with an instrumental view. It just assumes really different parameters, and really different goods, to the instrumental theories.

                The other point about bullying, though, and those sorts of acts, is I would really push forward to kind of say... You know, it's not clear to me how, if somebody's sort of honor is offended or something like that, and they're sort of reduced in status.... And then they engage in some sort of violence to restore their honor -- we say, "Oh, yeah, that was moral."

                When really, concretely, what they've done is their status was lowered, and then they've engaged in aggression to raise their status back up, to some sort of initial default level. But then when somebody engages in aggression just from the default

level to go up, we say, "Well, that's not moral. That's not moral at all. That's just status-seeking."

To me, I think that those are really equivalent. And the only reason we view them differently has more to do with these sort of moral biases… things that have to do with omission and commission, or defaults, or something like that.

Julia: Well, I mean… I definitely have that intuition that you're describing, that I'm more willing to say, "Okay, I can see how that's morally motivated," if someone has been beaten down by society and now they're trying to restore their rightful place, or something.

I've observed people who do really antisocial things to individuals who have not harmed them at all. And what it has always felt like to me is that they're parsing *society* as the moral unit. They're not thinking about individuals as being separate moral units. They're just thinking about, "Well, society has done harm to me," or, "Society has taken something away from me, therefore I am entitled to take something back from society." And who they take it back from is not really that relevant.

Which is very alien to my, I guess, modern, individualistic, post-Enlightenment, utilitarian sentiments, or something. But I could imagine a version of morality, that's maybe an older one, that does not distinguish between the different relevant units of moral calculus in society.

Tage Rai: Yeah, and I think, really, it sort of depends on the cultural rules as far as what's appropriate or not. In some cases, it's saying, "Hey, you really need to target the specific person that aggrieved you." In other cases, it's saying, "Well, you know, you can target anybody in that person's social group."

In other cases, it's just pure displaced aggression. If you were harmed, then you have the right to basically go out and hurt whatever stranger you happen to meet first. And that's kind of the classic example of these weird head-hunting cases, of dubious ethnography.

But this idea that you can have this sort of range, I think, is real.

Julia: When I read accounts of bullies who enjoy humiliating their victims -- or to go a little farther afield, I also like reading about trolls. Self-identified internet trolls. People ask them, "Why do you do it?" And a common response is something like:

"Well, you know, it's fun, it's amusing to me, and anyway, people shouldn't take it so seriously. The rules don't apply to the internet. "The rules of how you're supposed to treat people don't apply to the internet; it's your fault anyway if you let it bother you."

And I guess what I'm wondering is: Could your theory of morally motivated violence just be stated as, "People, when they commit acts of violence -- or figurative violence, in the case of trolling -- is it not so much that they think they're being actively *virtuous*, as that they think what they're doing is morally *permissible*?"

Those are two very different things that I think you sometimes kind of slip between, or back and forth between, in your discussion of the theory.

Tage Rai: Yeah, I think that's a fair critique. I think part of why I slip between them, though, is because I honestly see them as sort of on a continuum, as opposed to being distinct.

But what I would say about your example is, you know, I do think it's true in all sorts of contexts that we could have bullies who basically do seem motivated by a sense of righteousness. A sense of, "Yeah, this is how it should be."

And then we could also have bullies who are doing what you're saying, that they're drawing a sort of boundary beyond which moral concerns don't actually apply. And if they feel that they're within that, in those bounds, then by my sort of approach to morality, they wouldn't necessarily be acting morally.

And here's where it gets complicated. When I talked about, "Oh, well, are you defining things scientifically or sort of phenomenologically?" Hopefully in those cases that they match, but they don't always.

But I tend to argue that what morality's really about is about regulating different kinds of relationships. But if you don't

actually perceive a relationship, or you don't perceive the relationships as being relevant or motivating in any way, then your actions aren't really going to necessarily be morally motivated. And that's where you are going to find just more instrumental kinds of stuff.

Julia: Can we look at what interventions should, or do, reduce violence, as a way to determine what the right causal theory or causal explanation of violence should be?

           What would your model, of what causes violence -- would it suggest that there are certain interventions that we should be doing, that other theories of violence wouldn't?

Tage Rai: It seems to me like you asked two kind of different questions there. The first one was, if my [theory] suggests different interventions, and those interventions work, does that cycle back to disambiguate which theory is right? And I'd have to think about that.

           The second one is easier, I think. Which is, does my sort of approach make different sorts of policy implications or something like that, than the other, different ones.

Julia: Different predictions, yeah.

Tage Rai: And I think, yeah, it does. At its core, our approach would argue that, yeah, you know, if you really want to change violence and reduce violence, then you have to really focus on the social norms surrounding it. As opposed to, let's say, the cost and benefits, the material costs and benefits involved.

           So, for example, I have this new project where... it's very much tied to signaling. Where we do this third-party punishment stuff. And it turns out that if you pay people to engage in third-party punishment -- punishing somebody who has made an unfair offer -- if you try to pay them some additional benefits on top, they're actually less likely to do it than if you didn't pay them anything at all. That we essentially crowd out the behavior.

           And it turns out that the reason you crowd it out is because you've now corrupted the sort of signaling value of the punishment. And so, it turns out that people want to signal to

other people that they're a good person. And in the context of moralistic punishment, the way to signal to other people that you're a good person is by hurting another person.

Julia: Yeah.

Tage Rai: Does that make sense?

Julia: Without benefit to yourself, you mean. Yeah.

Tage Rai: Yeah… we can kind of think about some real-world examples, too. A real world example might be: In Chicago and Baltimore and some other places, there were these kind of Cure Violence or Ceasefire-type programs. And initially, those sorts of programs were really geared around carrots and sticks. Violent gang members in the community would basically be told that, hey, they need to stop engaging in violent crime and if they don't, there's going to be serious consequences. If they play ball, they're going to get various kinds of help with jobs and other kinds of structural things.

               And on its own, that wasn't really successful enough. Only when it started to get paired with them being brought face-to-face with community members who were respected, and who were telling them that this was wrong, did things start to really change.

               Now, that doesn't mean that the morals worked on their own, either; they didn't. You had to pair it with material consequences, too. But really leveraging the sort of social ties within the communities did have this sort of effect. I think this is part of some of the other kind of community policing kinds of things that we're seeing.

               On the flip side, our theory sort of predicts that things like body cameras aren't going to necessarily be straightforward for changing moral behavior. Because the implicit assumption there is that people are less likely to do things that are wrong if they know they're being watched. But if people don't actually believe that what they're doing is wrong, then we don't have any reason to think that being observed is going to change anything.

Tage Rai: Does that make sense?

Julia:	Yeah, good point.
Tage Rai:	So, you get these kind of weird situations like that. Another space where I think our theory makes some different predictions has to do with the sort of broad approach to saying, "Oh, well, we just need to start to recognize the humanity of our enemies," or something like that. Build empathy, don't dehumanize, et cetera.
	My argument really suggests that a lot of violence is directed toward people we recognize as fully human. And that's actually part of why we desire to hurt them, because we want to hurt people who deserve it, and who can suffer, and who will understand why we're doing it to them.
	And so, that sort of makes different predictions about whether you should really be building up humanity. If you think about things going on at the Southern border now, I think a standard sort of dehumanization approach might be like, "Oh, well, we really just need to emphasize the humanity of these children in particular," or something like that.
	But our approach would actually say, "Well, no, if the moral motives are geared toward thinking that what these people are doing is wrong, then people are simply going to reframe that. They're going to see those pictures of children and they're just going to blame their parents and get even more angry."
	Does that follow?
Julia:	Yeah, I mean, I guess it depends on whether you think... because your theory does allow that some violence is instrumental, just not nearly as much as people think, right?
	So, if you thought that the efforts to keep refugees out, or to imprison refugees were instrumental, because the main goal was just preserving America's low refugee status or something, then maybe making people feel guilty about the effects of their actions on children would help.
	But if it's motivated by moral outrage at these refugees who dare to violate our laws, then I guess your theory would predict, no, this is not going to work. Is that right?

Tage Rai: Yeah, I think so. Again, getting back to this kind of initial framing of what's motivating versus what's sort of inhibiting, and those things coming into contact and being kind of two separate processes... I think when we think about American attitudes and opposition to kind of immigrants and refugees and stuff, there are sort of two distinct psychological processes.

One really might be a process of dehumanization, where we don't see victims as fellow human beings, and therefore we don't care about their welfare. And that creates a kind of sense of apathy, such that we're not willing to help them.

But that's very different from another psychological process, which is sort of active antipathy toward these people based on kind of moral motivation, where we actually want to hurt them.

And we can look at other kinds of cases. My work would suggest that a lot of things where we've said, "Oh, it was about dehumanization," it's not totally clear. The people who engaged in direct violence, if we think about ethnic cleansings or genocides or something like that, but people who are really kind of the most ideologically motivated -- my argument would be that they weren't necessarily dehumanizing their targets.

Or at least that dehumanization wasn't really driving their actions on the front end...

Julia: Right. The confusing thing, I think -- and this is something that reading some of your papers helped clarify for me -- the confusing thing is that verbal acts of dehumanization, like referring to groups of people as vermin, rats, or snakes, or brutes, or animals, things like that... that looks like dehumanization.

But the reason that people want to do that is because they believe that the targets of that language can suffer, and will find it humiliating and degrading to be called vermin. Whereas if they just saw the targets of the language as unfeeling, unthinking brutes, then the language would lose a lot of power.

It's a confusing... When we talk about, "Are they dehumanizing their victims?", it depends on whether we mean, "Are they

thinking of them as un-human," or "Are they intentionally trying to make them *feel* un-human?"

Tage Rai: Yeah. That's right. Yeah. Are they actually conceptualizing them as less human, or... I don't even know if they were trying to make them feel less human, so much as they're just trying to really kind of make them feel humiliated and degraded through the comparison.

Julia: Well, this was another case in which I thought there might be a semantic disagreement happening. I know you don't like to have semantic debates, but my perception was that there might be two different senses of dehumanization being used.

Where one notion is, "This person has no mental states or feelings, at least none that are salient to me," and then a different one is, "This person has no inherent rights or dignity. They're not someone that I have to treat like a person."

And so, maybe the people who disagree with you, and are arguing that dehumanization leads to violence, they're using the second definition. Whereas you're using the first.

Tage Rai: I think that's totally possible. When I talk to people in this field, what they tell me is, like, "Well, no, what dehumanization is, is ascribing all these sort of evil aspects to the person. They have evil desires, and they want to hurt people..." and et cetera, et cetera.

And I'm like, "Well, okay, that's a really weird definition-"

Julia: Yeah, I don't endorse that. That seems like confusing cause and effect, or something.

Tage Rai: What I kind of tell them is, I'm like, "You know, again, I don't really care what terms we use for a lot of these things." If people want to say some of the stuff that I'm categorizing as moral isn't really moral, or something like that, that's fine.

But if we can agree that there might be distinct psychological processes, processes where we're sort of attributing various kinds of mental processes to people, and that's causing one set of phenomenonological processes. Where we're stripping away,

actively stripping away those kinds of mental phenomena, feelings and emotions and thoughts, then that's causing a different kind of thing. Then, as long as we agree on that, then I don't really care beyond that.

Julia: Got it. I guess, second-to-last question: we were talking a few minutes ago about what predictions your theory makes, but we could also talk about retro-dictions -- What would your theory have predicted would have happened in the past?

And so, if you look at the *Better Angels of Our Nature* story that Steven Pinker tells, violence has gone down steadily over the centuries. At least if we're talking about interpersonal violence and not state-sponsored violence.

Is that what your model would have predicted? Given that, I don't think that, for most of the period of that decline in violence, there were campaigns to change people's moral values.

Tage Rai: Yeah. I don't necessarily think... I don't think, if it was 1500, that I would have necessarily predicted some giant decline in violence.

Pinker's story, I don't disagree with him about the decline. I know that's kind of contentious in some circles, but it seems straightforward to me. Where I would probably disagree -- or at least, question a lot -- would be on the claim that the source of the decline has to do with kind of rational enlightenment and stuff like that.

Instead, what I would say is, just because a lot of violence is morally motivated, that doesn't sort of imply the inverse, that moral motivation automatically requires violence.

And that really, what's happened over the last several hundred years is that people got non-violent options to satisfy their moral motives. And if we're talking about how to make the world less violent policy-wise, that's going to be sort of the key, is giving people non-violent means to satisfy their morality.

Again, getting back to these sort of bigger policy issues, I think we've often had this approach of material benefits and costs, you

know? If we want to reduce violence, then we should just increase the severity of punishment, or something like that.

But there's actually this longstanding, interesting puzzle in criminology, which is that, "Oh, well, it turns out perpetrators aren't that sensitive to increasing the severity of punishment." And the normal explanation for that has been like, "Well, criminals are just dumb. They have bad [utility] functions. They can't weigh the costs and benefits appropriately."

But actually, maybe it's the case that, especially in violent crime conditions, they're motivated by moral sentiments that are actually relatively insensitive to material costs and benefits, and are more sensitive to sort of social pressures. In which case, leveraging material costs and consequences isn't going to get you what you want.

Julia: Would an example of giving people an alternate outlet for their moral feelings be, just having the state be the one to take justice into its hands, so that you don't have to? And so the thing you do when you're outraged at what someone else has done, is get the police on it?

Tage Rai: Yeah, absolutely. I mean, I think this is where, structurally, it does match up with some of the Enlightenment claims. That we basically outsource the law… What I would say is that there actually is a ton of moral violence. We just don't do it ourselves. It's just transformed into structural violence. We just outsourced it to institutions.

Julia: Yeah. Yeah, that makes sense. I actually don't know if Pinker would disagree. I don't remember Better Angels well enough, but it seems very sensible.

Tage Rai: I don't know whether he'd disagree with that point or not. What I've seen is that some of these disagreements come out like, "Well, yeah, maybe fewer people are killing each other, but there's way more incarceration," or something. You know, how do you weigh those things against each other? And we don't really have a good criteria.

Julia: Yeah. I mean, when the question is, "How do you weigh the value of one type of violence versus the other?", that becomes very

tricky. But if we're trying to answer a causal question, then that's… only slightly less impossible.

Tage Rai: Yeah. So, outsourcing violence to a third party institution is one way. Another way is, if the purpose of violence in a lot of these contexts is to achieve some sort of legitimacy within your local community, then doing the kinds of structural sorts of things [we're doing] are going to have an effect, right? So, if you increase the ability of people to earn a living or earn respect through those sorts of means, as opposed to through violence, then that's going to have an effect, as well.

Julia: Okay, great. Well, that's probably a good place to stop, but Tage, before I let you go, I wanted to ask if there's any book or blog or other source that has had a big influence on your life or your career.

Tage Rai: There are two that come to mind, is that okay?

Julia: Yeah. Definitely.

Tage Rai: Okay. Probably the biggest sort of academic influence directly on my research would be Durkheim's *Division of Labor in Society*. I think psychology in particular is hyper-focused on the individual, and what does the individual mind perceive about the social world?

And what we actually need is something that would be more akin to the kind of distinction in economics between micro and macro. We need macro-psychology. And Durkheim, I think, really gets that.

Julia: Do you think that social psychology doesn't fit the bill?

Tage Rai: No, because… I think it might have used to, but actually… There was a kind of cognitive turn in social psychology around the late '70s and '80s, that went from being a psychology of groups, and how do groups work, towards, "What inferences about social groups do individual minds make?"

And that's really what a lot of social psychology is now, especially in America. Less so in Europe.

Julia: Oh, interesting.

Tage Rai: The other book that comes to mind is from earlier in my life, which would actually be The Iliad.

Julia: Really, huh?

Tage Rai: I think The Iliad actually really did have this kind of effect on me. As a modern reader of The Iliad...

              We have this character, Achilles, who, basically, he's the greatest fighter, he's half-god. And the leader of the Greek army, Agamemnon, takes away his slave girl, right? And Achilles basically sulks in his tent and refuses to fight because his slave girl was taken away. He was given the second-best slave girl, when he thought he deserved the best slave girl. That's your main hero on the Greek side.

              And the main hero on the Trojan side is this guy, Hector. And Hector is the brother of Paris. Paris is the one who kidnapped Helen and brought her to Troy, and caused this war to happen in the first place. And Hector's basically the older brother, and he's just trying to clean up this mess. And all Hector cares about is his family. He cares about his wife and his son or something, and he cares about his dad, and he cares about Troy and saving Troy and all of his people. And he doesn't have any special powers. He's just this really brave leader of the Trojans.

              And as a modern reader, reading this, who are you rooting for here?

Julia: I mean, I don't remember who I was rooting for when I read it back in Freshman year of college. If I actually read it, instead of BS-ing my way through the class discussions. But from your description -- it's very hard to root for the sulking man-baby in the tent with the second-best slave girl.

Tage Rai: Yeah, you would be rooting for Hector, right?

Julia: Yeah.

Tage Rai: At the end of the book, you get to this sort of climactic battle where basically, what ends up happening is Hector and Achilles

are facing off, right outside the walls of Troy. And it's just them, one on one. And I think there's a sense that if Hector can just defeat Achilles in this battle, the Greeks will be so defeated that they'll turn back. You know? Hector knows he can win the war, he can save his family, he can save everyone if he just does this.

And in that pivotal moment, what happens is Hector turns, gets scared, and he runs away!

Julia: Oh. I must not have read it, because I feel like I would have remembered that disappointing moment.

Tage Rai: Well, who knows? Maybe I'm misremembering. But this is my memory of it from around the same time-

Julia: No, I trust you more than Freshman-year-Julia. Go on.

Tage Rai: He starts running around the walls of Troy and Achilles chases him. I think he shoots him in the back.

Julia: Ugh! This is all just so disappointing.

Tage Rai: And then Hector is dead, and then Achilles ties him up to his chariot and just drags his body around Troy again and again, all day, so that all the people in Troy can see their hero being dragged on the ground by a chariot, and his body being basically desecrated.

And that's sort of like... Eventually, they give Hector's body back, but that's the core of the book.

Julia: Who are we supposed to root for? From the perspective of Homer, was it? Who's the hero?

Tage Rai: But this is... Exactly! This is what I love about this book, and why I think it kind of really had a big effect on me. It's like, as a modern Western reader, we read this book as a tragedy. We say, "Oh, man! That really good guy, Hector, he lost, he died!"

But I actually think that 2,000 years ago, or 2,500 years ago, whatever, the ancient Greeks cheered in that moment. Achilles was the hero. Because he had sort of divine right and he was just naturally morally good.

It was a completely different set of moral standards, that essentially makes Achilles the hero. And it was right for him to sulk in his tent. Of course he should have! He deserved that slave girl! And it would have been wrong for him to fight, because he had been offended.

Julia: Oh, my God.

Tage Rai: The guest/host relationship had been violated. The honor code had been violated.

Julia: It's so crazy, because in a modern movie, if your first introduction to a character is he's sulking in his tent because he didn't get the first-best of something... that's your hit-with-an-anvil-over-the-head cue that you're supposed to dislike that person. And they're not going to triumph in the end. Yeah.

Tage Rai: Exactly! That's why I think that book does it for me so much. It shows this is a perfect example of how wildly different moral forms and values can be throughout cultures.

Julia: I increasingly feel, the more I read history, that the two main lessons I draw from it are, one, "Wow, people in the past were so similar to me. I didn't realize that people in the past could think and feel the same things I think and feel."

                 And the other lesson being, "God, these are like aliens! How are these humans? I don't relate to these people at all!"

                 It's really a wide range.

Tage Rai: Well, I mean, I actually think that probably the case is that, a lot of modern Westerners, there is a part of them that actually feels the same way as the Greeks. We just have a lot of other explicit morality laid on that also competes with that.

Julia: Well, anyway, I'm going to not think about that. Thank you so much.

Tage Rai: I'm sure there are scholars of ancient Greek literature that disagree with my reading of The Iliad...

Julia: Well, again, it's been a long time, but wasn't it written by people who were on the same side as Achilles?

Tage Rai: Yeah.

Julia: That definitely points to, we would have been expected to see Achilles as the hero, for sure. Wow.

Tage Rai: Yeah! They go in and they slaughter everybody! And they hunt down Hector's family, and they throw his baby from the top of the rafters, and they do the worst things possible. But I think that they're the good guys, for the Greeks.

Julia: I feel the same thing every time we read the Haggadah, on... I don't know if you've been to a Seder, a Jewish Seder on Passover. But there's this passage in the Haggadah, which is the thing we all read, where we're thanking God for all the things he did for us, including killing the firstborn children of the Egyptians. We're like, "Oh, it would have been enough if you had only done these things, but also you killed all these innocent babies! Thank you so much! That was such a lovely thing to do!"

                And, yeah, it's very jarring as a modern reader.

                Well, thank you so much for that stimulating conversation and debate and storytelling. And I was just trying to think about what photo of you I'm going to use for the podcast website, and I think I'm going to go with this... I don't know if it was an author photo, an interview photo or something, but there's this great pic of you and, I guess, your co-author, Alan, looking very grim-

Tage Rai: Oh, don't use that!

Julia: In a graveyard. Which I guess was the...

Tage Rai: ... they had us take this horrendous picture in front of a graveyard...

Julia: Was that their attempt to illustrate that you guys study violence, I guess?

Tage Rai: Yeah, that was the idea.

Julia: I mean, out of all the background illustrations of violence that they could have chosen, a graveyard is one of the less bad ones, I guess.

All right, well, thanks again! Tage, it's been a pleasure.

Tage Rai: Same here, Julia.

Julia: This concludes another episode of Rationally Speaking. Join us next time for more explorations on the borderlands between reason and nonsense.