

Transcript: Zvi Mowshowitz / Robin Hanson Debate “Should Blackmail be Legal?”

LessWrong post with highlights is [here](#). Video is [here](#). This transcript has been significantly edited for readability.

Ben Pace: I think we'll start off with opening statements from Robin and Zvi, about 10 minutes each, then we'll talk for about an hour and have a bunch of Q and A stuff after that. So I think that's everything. In that case, I'm looking forward to hearing Robin's statement on checkmate on blackmail. Robin?

Robin Hanson: Hello, everyone. I broke my arm a week and a half ago, so I'm a little slow so if there's any way in which I seem a little slow compared to usual, that's probably it, but I can't type very well, but I can talk. So this is great. Let's talk.

Robin Hanson: The topic really is gossip and regulation of gossip. Blackmail is one of the many ways to regulate gossip. Gossip is, of course, simply talking about each other. But it is also one of the ways that we as humans enforce norms. When we disapprove of someone's actions, we use shaming to hurt their reputation and change their incentives. That's gossip.

We generally allow gossip in our society, though we limit and regulate it in some ways. In some instances, we require gossip, like when you witness a crime. In other instances, we ban it, such as sharing nude pictures of someone else or the pin number to their bank account. We also ban false gossip through slander, libel, and defamation laws. In general, it makes sense to regulate gossip in these ways, so keep in mind that it both banning and requiring gossip are potential solutions to particular problems

Laws against blackmail are essentially an attempt to keep money out of gossip. For context, there are many other things we generally don't allow to be transactional. We allow people to choose to have sex with consensual adults and to have children. But we have laws against prostitution, and though you can adopt a child or put your own up for adoption, it is generally illegal to profit from it in anyway. We do, however, allow for indirect forms of gift exchange. So you can trade money for sex only if you do it indirectly over time, but not directly. This is all relevant because the laws against blackmail are mainly laws against money and gossip.

With respect to gossip, we allow people to have money involved through third parties, like when a newspaper pays someone to tell their story. Furthermore, we have nondisclosure agreements (NDAs), which are very common. NDAs are a contract

between two parties with one offering the other some form of compensation not to disclose something. This is legal. But we do not allow one party to explicitly demand compensation from the other for refraining from disclosing something.

What is difficult to justify, I think, isn't having regulation for or against gossip, or even banning money in respect to gossip. The real puzzle is why we allow people to make deals where they get paid not to say something only if the person with the money makes the offer, but not the other way around.

Of course, once you realize that these are the rules, you can get around them pretty easily. It's actually relatively straightforward to indirectly suggest that you might be open to such a deal, but not explicitly offer the deal. And then, the other side can offer the deal and then it's all legal. For example, typically when you inform someone that you're going to pursue divorce proceedings implicitly, there is the threat, "We've got information on you, and if you let us go ahead, we're going to air all these embarrassing things in court. So if you don't want that embarrassing stuff in court, then you better settle with us on this divorce agreement. Otherwise, we're going to tell."

That's basically the threat of blackmail, but it's legal because it's sufficiently indirect. There are many other legal ways of indirectly suggesting a deal. Again, it is only illegal to explicitly suggest receiving money to keep a secret.

The key idea is that when intelligent people understand this, and use lawyers as intermediaries, they can legally blackmail people if they choose. The law mainly just requires you use a lawyer and makes you go through hoops in terms of how you make the suggestion.

Sometimes people use phrasing here in terms of harm. They say, "Well, most gossip isn't harmful, but this gossip is harmful and this deal is harmful. And so, the problem here with gossip is that it's harmful, but of course harmful depends on the reference point for what would have happened without the deal. So if you have the reference point, that without this deal, I wouldn't have said anything. And then I come in and say, "Hey, I'm going to say something unless you make the deal" then I seem to be harming you, but if I was going to say it anyway, but I'm just letting you pay me not to say it well now, relative to that reference, I'm helping you. I'm making you better off relative to the reference. I think you have to be careful how you phrase talking about harm. Cause it depends on the reference.

I conducted a poll and three against one were in support of blackmail remaining illegal. But this was in reference to more extreme versions involving cash payments. In other situations, people seemed to be okay with blackmail. For example, say that you know something about me that I want to remain a secret. But I also know something about you that you want to remain a secret. So I blackmail you by threatening to reveal your secret if you reveal mine. According to my poll, 12 to one of the respondents were okay with this. If you think somebody is treating you unfairly, three out of two respondents said it was okay to threaten to reveal a secret about them in order to be treated more fairly.

Also If you have a crisis where your kid needs medical care and you're trying to get someone to take care of them, one-to-one are in favor are threatening to reveal a secret about that person in order to get the care.

So people are not very consistent about which situations they think blackmail is okay. It is only the situation involving extorting someone for cash that they disapprove of it. But, as we have already seen, people disapprove of the exchange of cash in many situations, like prostitution. So perhaps this is more about feeling awkward about letting cash be directly associated with things than it is about the actual outcomes.

Again, people are okay with indirect payments for sex, as long as it's not explicit prostitution: a nice date and dinner in exchange for sex would be okay in many people's minds, as long as it's not an explicit agreement. With blackmail, they don't approve of an explicit deal for money, but seem to find it okay with other forms of compensation.

My argument basically is that people have various aversions to money being associated with certain things. And blackmail is an example. It's not that we don't actually want the outcomes associated with this sort of thing. We just feel uncomfortable about the signals it would send if you explicitly allowed it.

Ben Pace: Thank you, Robin. Zvi, do you want to give your opening statement?

Zvi Mowshowitz: Yes. I think that Robin addressed some of the reasons that I strongly believe that blackmail should be illegal and strongly discouraged as much as possible. Similarly, I think he has hit upon some of the reasons why other people generally dislike it, but not the central reasons.

One particular thing is that the idea of money versus not money? I definitely feel like the most dangerous situations regarding blackmail are ones where what you don't want

money, but where what you want is to essentially corrupt another person, appropriate their power, and to get them to do things they would not otherwise do.

For example, when you are asking for better treatment and medical care, people approved of that, because it seemed like all you are doing is securing medical care for your family. But what you are actually doing is redirecting resources to you by threatening someone, which takes resources away from anybody who is not willing to threaten another person to secure those resources.

In general, people without money, which is most of us who aren't that vulnerable to money blackmail, are highly vulnerable to other forms of blackmail, and many people who have access to a lot of power can be convinced to use that power in other ways. If you can blackmail a judge, you can get blackmail on a Senator. You can get blackmail on a journalist. You can get black on any number of people. You can get very valuable resources out of them.

Framing it as gossip is one way of putting it. But the distinction between gossip motivated by harming a particular person and gossip motivated by the desire to seem useful or the desire to spread information is very large. There should be a presumption that sharing information because you feel people would find it useful is probably a positive effect. And we should, if the information is true, probably favor it. If your motives are to increase your leverage, to maximize harm to the person, I believe we should assume that the information is actually net harmful and will do damage.

Robin brought up the issue of divorce. I think that most of us can agree that when all of a couple's dirty laundry is aired in divorce court, this is not a net benefit to society, like useful gossip. This is mostly disruptive and leaves people destroyed and miserable. The parents hate each other and the kids end up traumatized and estranged. It's a very bad scene. And the fact that we usually favor whoever is willing to be madder and more destructive in a divorce settlement, and that this hangs over every marriage, seems like a very, very bad thing.

And I think that if we could arrange for divorce courts to be more private, that would actually be a big win. The question is whether we should ban things directly or indirectly. Prostitution is interesting case because I think both Robin and I would agree that prostitution should be legal, though, to some extent, exchanging money for sexual favors should be against norms.

Yes, you can get around it. Obviously, plenty of people have nice dinners and then go home together who would not have otherwise if the dinner had not been purchased and an exchange made, but there is definitely a big distinction in terms of the lived experience of everybody involved. And in terms of the degree to which the exchange depends on money specifically is pretty big. Blackmail is even worse though. If you get to go to somebody directly and threaten them explicitly that unless you pay X I will do Y, the exchange becomes much more explicit, much more direct.

It allows the extraction of much more money and makes it much more likely that information will be revealed if the blackmailer is not paid in order to hold up their end of the deal and reputation. As Scott Alexander said, "Beware trivial of inconveniences." These are far from trivial inconveniences. They require you to go through lawyers, to make indirect suggestions, to have a legal venue to pursue your suggestion at all, in which to request the payment.

All of these things highly discourage the activity. Also, the fact that the activity is considered illegal and against norms itself, makes it very hard to punish norm violations, especially relatively small ones, by going through this process, because the very act of asking for blackmail causes you to be the one who looks worse when things are going good, rather than them.

In one of Robin's articles, he brought up the David Letterman situation. There's also the Jeff visa situation with the national Inquirer. People who come clean in the face of a blackmail will often be forgiven for the thing in which they were caught because the blackmailer looks even worse and should not be rewarded for their behavior. I think that's a very natural instinct and pertains to the dynamics. If the blackmailer was seen as having done nothing wrong, then the person being blackmailed would be in a very difficult spot.

In a world in which we had to fear blackmail, it would act as a tax on activities and situations. It rewards those willing and able to engage in blackmail. And I believe that the shift in resources would favor the powerful, it would favor the people we would not want to have power and resources, and it would be incredibly stressful for our lives. I think it would be a very bad thing.

We want to be able to go about our lives without living in fear, without fear of destructive activities. There's way too many different reasons why I believe this to include in an opening statement. But if we allowed blackmail on a wide scale, as is effectively

happening, and made it even cheaper, more effective, and more accepted, blackmail would become more ubiquitous.

And that society would become increasingly corrupted and distorted in ways that I do not think Robin believes.

Robin also brought up the issue of bribery. I think it's important to note that often the person being bribed will have influence over the allocation of resources, over people's perceptions and opinions, that they cannot easily turn in to money and cannot easily capitalize on. They also often have the ability to do things they would not do for money if they were not under duress or they would not do for an amount of money that people would be willing to pay, but which they could be blackmailed into doing.

So if you are a spy or a policeman trying to figure out a crime, then blackmail can get you to do things for less of the cost than you would normally demand, which implies that you doing this thing is highly inefficient.

A person who is blackmailed into sex or pornographic pictures, for example, is clearly doing something they didn't want to do and that you weren't willing to pay for otherwise. And therefore, it's clearly highly inefficient. On a final note, if we have a duty to report crime, which could be considered gossip, then blackmailing somebody in order to not report crime is specifically saying that in exchange for money, you will commit a crime. And it seems very natural that this should be illegal.

In fact, it'd be strange if this were not illegal. And in general, the fact that you can threaten to do harm and thereby attract resources means that we reward people for gaining the ability to do harm. And this seems like exactly the opposite of what we want to do. I want to see people rewarded for the ability to help people, to make life better for others. And I want to see them engage in free trade in exchange for that basis. I do not want to see people strive to get leverage over people all the time. I think that would be a bad world to live in. I think we live in much more of that world now than we would like; I like to see the roles against, or this type of blackmail strengthened.

Ben Pace: I would be somewhat interested in hearing Robin's responses. One of the last things he mentioned sounded like he thought blackmail was a good thing, as long as there is no monetary exchange. I was a bit confused by that. I'm curious, Robin, if you had opinions on that.

Robin Hanson: We've got a lot of different topics here. If I just pick a few at random, it likely won't be the most important. But to me, the most important point here is whether you think the consequences of blackmail are bad.

If it's harmful, why should we allow it? That, that is the key question here. So the key issue is norm enforcement. In general human societies, for a very long time, have used gossip as one of the main mechanisms of norm enforcement. And I want to argue that's typically a good thing. If you violate a norm, you are at risk of someone gossiping about it and disapproving of you. You will then lose out socially, even if you gained in other ways from the norm you broke, which is what discourages norm violations.

So the key idea here is if you threaten to gossip in exchange for something, it is another way of discouraging whatever it is the person you threatened was thinking of doing. Even if the information doesn't get out, the same function as gossip is performed, that is, discouraging people from doing things that they would have done otherwise. In fact, it could be a stronger incentive because the person who might gossip may be more willing to do so in the hopes of getting paid. So it seems like allowing blackmail turns up the dial a bit, though not enormously, on the discouragement that you face from the people who will gossip about what you do.

In regard to divorce, if someone has been mistreating their spouse, gossip would, and should, make them look bad. People should be afraid of being made to look bad by gossip about how they mistreated their spouse. That's been a core function of gossip for a long time. So I do not grant the point that we would be better off if people couldn't gossip about how they were mistreated in their marriages. I have other points, but it seemed to me the most important is to discuss whether it is good or bad.

The other point I was thinking of bringing up and I'll just lay it out there, which is just look: Obviously, powerful people are the people most likely to be targeted for blackmail because they have money and influence. They are more vulnerable to people disliking and disapproving of them.

In fact, I think that banning blackmail is a way of protecting the powerful, since they are the most vulnerable to it. The powerful people in our society do violate a lot of norms that the rest of us follow, and they get away with it because it's hard for people to call them out on it. Exactly because they have so many ways to retaliate.

Ben Pace: I'll just stop right there. Zvi, would you like to respond?

Zvi Mowshowitz: Yes. So I will briefly respond to that later point by saying that I disagree that those who have power tend to be more vulnerable to, and targeted by, blackmail. I think that power protects you from blackmail. You can threaten retaliation, you can protect yourself against it. You can better hide your secrets. You can be insulated from the consequences. I think we see this in cases where people are destroyed by smear campaigns. Yes, you notice the famous people who are destroyed by smear campaigns. But they tend to be okay. It is the ordinary man on the street who loses his job when they are subject to a cancellation and can't get their lives back.

Robin Hanson: Well, cancellation is a whole different issue. But if you looked at court cases where someone was accused of blackmail, I'm pretty sure you will find a strong slant toward the powerful. And as the law is practiced, it protects mainly the powerful.

Zvi Mowshowitz: I believe that the threat of it being illegal, and the modifications of people's willingness to behave because it's illegal, changes how people behave substantially. And that making blackmail illegal, doesn't just affect the cases in which blackmail is tried in court.

Robin Hanson: It doesn't mainly affect the people who are actually prosecuted for it. Nobody's ever sued for blackmail because they tried to blackmail you or me, that means people like you or me, aren't protected much by laws against blackmail.

Zvi Mowshowitz: On the contrary, nobody's ever tried to blackmail me. And I'm very happy about that.

Robin Hanson: But it's not because there are these laws out there. If cases are about people like us, they just don't show up much. Right?

Zvi Mowshowitz: No, I think it is in part because of that. I think it is because it has not been normalized that blackmail is all right. I think it's because anyone who tried to set up a large-scale blackmail organization would be shut down. I think it is because we aren't okay with this sort of thing, but also because people know that if they tried to extract a lot of resources, it wouldn't work. But I think you're right, the central point is whether blackmail is a net positive when it happens. Is blackmail good or bad in general?

Zvi Mowshowitz: I do think we disagree about whether or not allowing the blackmail causes good or bad things to happen in general. But I think we also have disagreement about concretely when the blackmail happens.

Zvi Mowshowitz: So in scenario one, I have information about someone and I'm allowed to blackmail them. Then scenario two, I have information about someone and I'm not allowed to blackmail them. So if I'm not allowed to blackmail them, then I can choose to either share the information or not share the information. And I will share the information, presumably primarily on the basis of whether or not it helps me to share the information, whether or not I think it is useful. And I'll be rewarded for sharing the information versus you'll be seen as spiteful and get punished.

Robin Hanson: Let's go through a step by step, setting aside blackmail. Do you think that people gossiping about other people and saying negative things about them is, on average, a good thing or not? Should we ban that? I mean, we could in principle ban saying negative things about anybody, Would you support that? Or do you think, on average, it's good to let people say negative things about other people?

Zvi Mowshowitz: I actually don't think these are the same question. I think that you can want to reduce it on the margin, but understand that banning all of it would be horribly bad. I think banning all the negative gossip would be horribly bad.

Robin Hanson: On what margin would you cut it down on? Or on what margin would you ban people saying negative things about other people? What kind of margin would that be? Big things? People you know well? People you know less well? Pick a margin. I mean, if there's no margin, you can only do it on average, right?

Zvi Mowshowitz: No, I mean, I would ban people saying negative things, ideally in ways that are engineered to cause harm, but it's very difficult.

Robin Hanson: What does it mean engineered to cause harm? I know something about you, I tell it, well, how is that engineered to cause harm?

Zvi Mowshowitz: Okay. Basically, if I try to cast you in the worst possible light, if I go through all of your past actions in order to try and present what appears to be a pattern of behavior, if I try to quote people out of context to make what you did look bad.

Robin Hanson: So an example of that would be, say, a divorce proceeding, somebody trying to make the other side look bad. Another example might be a mud raking news article that tries to make something look bad. So you'd ban both of those. No, you can't try to consistently make somebody look bad, that's just not allowed. You always have to

try to show good sides as well as bad sides in any presentation of the people. Is that a rule you'd endorse?

Zvi Mowshowitz: No. I think that is far too strong, but I said I would support making divorce proceedings private.

Robin Hanson: You're not allowed to tell people about how your spouse mistreated you? Are you really going to make a law saying you are not allowed to tell other people how your spouse mistreated you?

Zvi Mowshowitz: I think you should have a law that divorcing someone doesn't allow you to put them on the stand and then ask them about every little embarrassing thing in their lives, forcing you to go on the record.

Robin Hanson: Can I go to the press? Can I go tell my friends? Can I send out a newsletter to my friends? Can I say bad things about my spouse?

Ben Pace: Zvi would happily have the divorce proceedings happen if they were private, and allow people to say those things in public if they wanted to, but didn't want to set up a situation in which they were incentivized to publish naked pictures or things like that. Is that right Zvi?

Zvi Mowshowitz: Yes. But look, I grant Robin's point in general.

Robin Hanson: In a political context, are politicians allowed to say negative things about their opponents, to try to dig up dirt, to paint their opponents in the worst possible light? Is that a thing that people should be allowed to do?

Zvi Mowshowitz: I don't think it should be illegal, but I certainly think that the way that politics is practiced now should be discouraged and be considered against norms.

Robin Hanson: Well, the point I'm trying to establish is that you are, in general, accepting the idea that, on average, it's good to let people gossip about other people if it is true, even when it's negative. I haven't said anything about politics.

Zvi Mowshowitz: Yeah. I agree with that. To be explicit, in general, I agree with that. And I certainly think that clamping down on that hard would be a very bad mistake.

Robin Hanson: Okay. So my argument will be that allowing blackmail just turns up the dial a bit on that. It gives the same kind of incentives, but makes them a bit stronger.

Zvi Mowshowitz:

So, you see, that's where we disagree. I don't think that what happens is that every bit of information is 10% more likely to get out and be revealed to the public.

Robin Hanson: What number would you give? 10 percent? What about 20, 50, or 100 percent?

Zvi Mowshowitz: It wasn't that I disagree on the number. It's that I think that it affects different kinds of information in radically different ways.

Robin Hanson: What difference? We're starting with the situation, remember, where somebody is trying to make somebody look bad. That's our presumption here. Gossip, but not blackmail, is allowed. And now we're adding blackmail. And you're saying, if we take a situation where somebody was trying to make someone look bad and we add the possibility of getting money out of them, that'll make things much worse. And I'm trying to push you on what's the margin there? Which kinds of situations, and in what way, is that worse?

Zvi Mowshowitz: Okay. The typical case of blackmail involves finding something that would in fact cause a lot of private harm to somebody. That's why you are able to blackmail them.

Robin Hanson: But that's also true for the other kinds of gossip we just talked about. When I'm trying to dig up dirt on you to make you look bad as my political opponent, yes, of course you're going to be hurt if I find dirt.

Ben Pace: Can I jump in and just add that I do think it is not only the case that people who are looking to find dirt on you can use money in that situation. It also, I think, would increase substantially the incentive for people to try to find dirt on you. A lot of people who are not trying to find dirt on you would suddenly start trying to find dirt on you because it would be profitable.

Zvi Mowshowitz: They would also want to manufacture dirt on you.

Robin Hanson: That's also true for politicians now. The politician wants to manufacture dirt on you. Your divorce opponent wants to manufacture it on you. We are already in

the situation. The news reporter wants to manufacture dirt on, says Zuckerberg. So if you're a newspaper personnel and you want to make Zuckerberg look bad, and a lot of people do, then they'll be looking for dirt. Same with Bezos. They're already in the situation where they have an incentive to be selective about just showing the bad things and to manufacture dirt if they could. That's already part of the situation we're comparing.

Zvi Mowshowitz: What I'm saying is that blackmail specifically selects for the scenarios in which there is great harm.

Robin Hanson: So does trying to make somebody look bad. That's the thing we're talking about. There are situations now where people are trying to make other people look bad. Either you want to ban these situations, or you have to accept that on average they are good, even though they contain the problems reporting.

Zvi Mowshowitz: No. I think there's an important fallacy there: that I cannot want to ban something but still think it's in general pretty bad. I think the gossip in general is definitely very good. Sharing true information about people is obviously, in general, a good thing. I think that sharing information about people that happens to harm at least one person is still, generally, pretty good. I think sharing information about somebody which is chosen specifically and engineered to harm someone, I think, in general tends to be bad.

Robin Hanson: It's intended to make you look bad relative to an audience, but that highlights that the audience wants to know, and plausibly socially benefits from, finding out. If somebody is trying to hide dirt, other people have legitimate reasons for wanting to know it. For most of these cases, there are good historical reasons.

Zvi Mowshowitz: I don't think that expresses the primary mechanism for which you don't want it to get out, and the primary mechanism by which it is harmful. I think often the reason why you don't want it to get out is because it could get you fired, specific people would retaliate against you because they would not like particular things that you have to say.

Robin Hanson: If somebody's going to fire you if there's something they found out about you, then it's got to be something that they disapprove of.

Zvi Mowshowitz: No. Essentially your argument is that, in today's world, if people online were to find out something about you and decide to cause a lot of trouble in your

life based on this thing, that it must've been a bad thing that the public absolutely needed to know.

Robin Hanson: So you're going to an extreme, deterministic 'always'.

Zvi Mowshowitz: Sorry. It doesn't have to be always.

Robin Hanson: On average, letting people know about other people's dirt is a good thing. On average, the incentives to not have dirt and the incentive to expose dirt, are good things. Yes, they can go wrong in many ways. But on average, they're good. That's my fundamental claim about why gossip is generally good, even when people are trying to find dirt in order to make someone look bad, and blackmail is just upping the incentives on that somewhat.

Ben Pace: There's many situations in which people look for bad material on others, and they try to get negative information about them to punish them or to hurt them. And Robin's take is that these are often healthy and good for society, and we don't want to abandon these, and allowing blackmail would allow this to work a bit more. It would allow financial incentives into this in a bigger way.

Robin Hanson: For most people it's a small effect, but for the powerful it's a much bigger effect. That's the main distinction, I would say, for people who have a lot of money, or people for whom public opinion is very important, they will face a much bigger threat.

Ben Pace: Okay. So Zvi, I think you understand that that's Robin's take, is that it generally allows for social accounting, and finding out bad things about people, to generally just be more efficient. I don't think that's your take.

Zvi Mowshowitz: So I have a list of ways in which I disagree with this. And one of them is that I don't even think that allowing blackmail would increase the number of such things that are revealed. I think it would decrease it.

Ben Pace: It would decrease the number of things that were revealed.

Zvi Mowshowitz: Yes.

Robin Hanson: It's about the discouragement, less about the revealing, right? There are two ways to discourage something. One is to reveal it, and the other is to threaten to reveal it and to make you pay to discourage it.

Ben Pace: Sure. So you think it would decrease it? Because people would certainly be scared and become very secretive about their information.

Zvi Mowshowitz: People would be more secretive. Sometimes when they are blackmailed, they would in fact pay. Other times they would find credible threats of retaliation and that would prevent the information from coming out. Right now, when normally determining whether or not to share gossip, people tend to share net useful gossip more than they tend to share net harmful gossip. Blackmail reverses this incentive and also causes people to look for net harmful gossip, rather than look for net helpful gossip. Most of the time when people are looking for information, they're looking for information the public needs to know.

Robin Hanson: Explain that. I'm open to that, but you need more of an argument. That's just a claim at the moment. Lay it out. What's the argument for that claim?

Zvi Mowshowitz: Okay. I believe that, most of the time, people are not seeking out harmful information. They're primarily looking to benefit themselves, or benefit their friends or their allies in some way. Hurting someone else is mostly a side effect.

Robin Hanson: That's also true with blackmail. The main effect is the money, not the hurt. Their main motivation is to get the compensation.

Zvi Mowshowitz: The main benefit is to gaining power to extract something of value over a person.

Robin Hanson: But that's not the same as hurting them.

Zvi Mowshowitz: But you do that by gaining the ability to hurt someone.

Robin Hanson: I disagree with that whole framing. Verizon, which supplies my internet and my phone and my TV, they want me to really want their product. So the more that they can make me really desperate for their product, then the more I'm willing to pay for it. Of course, they do it, hopefully, by making their product attractive. But that is a way of gaining power over me. In general, all through society, when people are making deals with the other, in anticipation of those deals, they want to be in demand. They want to

want the other party to want them. That's basically the same thing. It's all about wanting to have the other party want to make the deal. That happens in marriages; it happens in jobs; it happens to me with Verizon. Is that harmful? Is the effort that Verizon goes through to make sure that I don't want to be without their service harming me because it makes me more willing to pay for their service?

Zvi Mowshowitz: But doesn't Verizon do that by offering a benefit? Verizon creates a service that makes your life better so that you will be willing to buy it. So if Verizon were to, say, cut the wires of their competition so that their competition couldn't come to your house, and then threatened to cut off your service unless you paid them 10 times as much, that seems like it would not be a good thing.

Robin Hanson: That's what I said in my initial remarks about the reference point. You have in mind, the reference point is, I say nothing. And so I'm harming you by threatening to say something. But what if the reference was, I was going to say it anyway and you pay me not to say it. Well now with respect to that reference point, I'm helping you by letting you pay me not to say it. So it all comes down to what's the reference behavior you thought would have happened instead.

Ben Pace: This has been quite good. I think Oliver Habryka was going to add in something. Oli, do you want to say something from the chat?

Oliver Habryka: Yeah. So just to clarify, my sense is that right now we have an optimization process that isn't perfect. If information has public value, you're encouraged to share it because the public might reward you for it. In a world with blackmail, the thing that would maximize your payoff would be to choose the thing that allows you to cause the most harm for the other person you are talking to, such that you can then extract money. If you can cause someone \$10,000 of damages in spreading some type of information, then you can extract approximately \$10,000. Obviously game theory is hard. But roughly, you're incentivized to maximize the harm for the other person.

Oliver Habryka: To me, it seems pretty clear that those can diverge. The obvious example here is publishing someone's account password. The thing that would damage their private passwords, it's just losing access to all of their accounts and all the indirect damage that would cause. So they can totally use it as blackmail material because it is a direct thing that would cause damage to them.

Robin Hanson: Remember one of the other options is just to ban some kinds of gossip. So I said, we do properly ban telling people's passwords or their naked pictures. If you think there's a kind of gossip that's just harmful, you just ban that kind of gossip. We're talking about allowing that gossip, except when one side makes an offer to pay the other, but not the vice versa. That's the puzzle we're talking about, not just the generic, that some things you might not want to let people gossip about.

Ben Pace: Before I respond to Robin, Zvi, could I get your sense of whether that was an accurate summary of your position, from the things that Oli was saying?

Zvi Mowshowitz: It's a reasonable summary of one of my positions. One recent example of potential blackmail is what if someone were to blackmail Scott Alexander and threatened to reveal his true name?

Robin Hanson: But again, you can have privacy rules about that. If you just want to say, you're not allowed to reveal people's anonymous names, just make that the rule. It would be okay if it was gossip but not if it wasn't, but that's not true in this case. You think it would be bad even if it was revealed without monitoring content.

Ben Pace: So it sounds like Robin's getting a lot of work out of this. There are a couple of bad cases, but you can just outlaw those. Whereas my model of Zvi feels like those bad cases are innumerable and massive.

Zvi Mowshowitz: Both they are innumerable and large, and also, I think that banning the action itself is impractical. You would have way too much splash damage. There are often things which we would like not to happen, but which we know if we banned them, it would have negative effects and therefore we shouldn't do that. I think the best argument against blackmail is the libertarian argument, which is that there will be bad side effects. And I think that holds true. If you ban the NDAs, there are types of business you simply could not do, for example.

Ben Pace: So instead of proposing a bunch of small events, you're proposing a much larger event.

Robin Hanson: I don't think it's about size. I think it's about whether you can identify the kinds of gossip. The simplest policy is either requiring or banning gossip. We're talking about a more complicated thing, which is whether you can say it but not get paid to say it. Or you're only allowed to be paid to say it if *they* make the offer, not you. That's the more complicated version. The simplest one is just to let people to say or not say. We

already have slander and libel laws. Those are laws about what you're allowed to say. We have other sorts of privacy laws. We have and do enforce lots of laws on what you're allowed to say and what you're not allowed to say. If your problem can be solved with those, just solve it with those. Why invoke this more complicated thing?

Ben Pace: Does that seem like it solves the situation for you, Zvi? We can just ban the specific bad outcomes?

Zvi Mowshowitz: No. I do not think we could possibly enumerate a rule that we would want. First of all, sometimes it will be correct to want to share those things because there is enough of a positive need to share them, and you'll want to share them.

Robin Hanson: So regarding passwords, you don't want to make a law against sharing passwords? You want to let me share passwords if I can find it out, you just don't want me to ask for money for it? That's your solution?

Ben Pace: That's the current situation, isn't it?

Zvi Mowshowitz: I don't think so. My claim is not that there are no things we want to outright ban you from sharing. I'm saying there'll be a lot of things that we cannot enumerate and ban you from sharing that we nevertheless want to tax the sharing of.

Ben Pace: Can you say more about why you think that, Zvi?

Robin Hanson: Show us one.

Zvi Mowshowitz: Sharing Alexander's name was supposed to be the obvious, basic example of this. I do think that, in general, sharing private conversations of various types, which cause mobs to come after people, is similar. But let me give a specific example.

Robin Hanson: Give us something concrete. It's really hard to do this without some concrete examples.

Zvi Mowshowitz: Okay.

Robin Hanson: Alexander's passwords define a concrete example as far as I am concerned, but it seems to me either we want to allow people to air those things or not. I would be fine with saying, "You're not allowed to publish the anonymous name unless

you figured it out independently or something.” But you don't want to do that. You want to let people reveal anonymous names, unless they're being paid for it.

Zvi Mowshowitz: No. Sorry. I think you said that wrong. I'm not saying you can't pay me to reveal the information. I'm saying you can't blackmail me with the threat of revealing the information.

Robin Hanson: Oh, I agree. You can't demand from Alexander the payment. You can demand it from a newspaper, *The New York Times*.

Zvi Mowshowitz: I also can't demand it from someone else who would be afraid of it being revealed who isn't Alexander.

Robin Hanson: You're just making claims. Tell us why that's good or bad. Again, the key thing isn't whether certain information should be revealed or not, it's whether you should add this extra complexity. So again, we always have the option to ban gossip or to require it, but what we have is this weird rule where you're allowed to do it in trade for other things, but not for money, though you are allowed to do it for money if one side makes the offer. This is what I find very hard to justify. I can understand why some things should be private and others published, and why some things should be allowed to be said and others not. But why this weird combination of not allowing money unless one side makes the offer but not the other. You haven't addressed this one side versus the other thing at all in your entire conversation here. You haven't addressed the possibility of an NDA, doing all of these examples you don't like.

Ben Pace: All right, Robin, I think I have a clear impression of what you're looking for. And I think Zvi now wanted to take a few seconds to think of a concrete example that you wanted him to give that met some of these conditions.

Zvi Mowshowitz: I can think of a number of examples of information about me that I think should fall in this category, but I obviously don't want to use them as the examples.

Robin Hanson: Let's talk about Mr. Smith. Tell me about some information about Mr. Smith.

Zvi Mowshowitz: Who is totally not me.

Ben Pace: I also mentioned that I don't know if this sides to your points of view. I feel that often the boundary lines around privacy and private information are difficult to

determine. I remember reading that a case of a doxing of someone's home address, and I was like, "Well, that seems bad." And they were like, "Oh no, they were just going to share a photo of a house, but was not going to share the number." I still thought that felt bad.

Zvi Mowshowitz: Typically, five minutes later someone figures out where the house is.

Ben Pace: Yes.

Robin Hanson: It's not easy to decide rules about privacy and rules about what you're allowed to gossip about. I'm just saying that's a separate issue from blackmail. Blackmail doesn't help make that any easier. If you say you can't blackmail about the house, you already have to settle the question whether you can reveal the location of the house. You have to settle the question whether that's appropriate information before you decide whether you want to allow that to form blackmail. Allowing blackmail just adds more complexity.

Zvi Mowshowitz: Okay. So let's use the example of a doxing. I think that's a reasonable example of something that we should discourage.

Robin Hanson: Why not just make it illegal? Why not simply say, you're not allowed to dox. Done.

Zvi Mowshowitz: So you're not allowed to ever reveal where anybody lives.

Robin Hanson: They don't want it.

Zvi Mowshowitz: You have to get enforcements. So you have to get consent of someone to reveal their location or else you're guilty of a crime, is that the proposal?

Robin Hanson: It's that or allow it, but what the hell does blackmail have to do that? That's my point.

Zvi Mowshowitz: Okay. Blackmail shifts the incentives as to whether or not you choose to find out where somebody lives. It shifts the incentive as to whether or not you choose to reveal where somebody's lives. It shifts the incentives of what it means to be someone who doesn't want their house revealed. So if blackmail is not allowed, then there is much less incentive for somebody to go searching for my address and then blackmail me.

Ben Pace: I feel like there's a fairly simple setup here: if somebody wants to hurt me, and is looking for my address, by default, there are not a lot of people who are incentivized to give this bad person my address. Most people who know me and my address won't go looking for this person, and won't go trying to find that deal. Whereas as soon as they can get money out of it, there is a motivation for them to give this bad person my address.

Robin Hanson: But are you willing to make the claim that this should be disallowed only if one side makes the offer, but not the other? Zvi hasn't erased that law.

Ben Pace: What is the other side of that situation?

Robin Hanson: If I say, "Hey, I've got the addresses of people, anybody want to pay me for them?" That's one side, and the other side is somebody saying, "Hey, I'd like to know the address of... Who will pay me to find that out?"

Ben Pace: Should you be able to pay people to find addresses?

Zvi Mowshowitz: You can, actually, normally, but not in every case. But in many cases.

Ben Pace: I feel like it is probably good to be able to opt out of such a system where your address will not be sold by default.

Robin Hanson: So again, the analog with blackmail is we're only making it illegal if one side makes the offer but not the other. That seems to me the hardest part to rationalize. I could see you saying you're not allowed to sell addresses. I could see you saying you are allowed to sell addresses. But why say you're only allowed to sell it if one side makes the offer, but not the other side?

Ben Pace: I feel in part, selecting against the people who can most easily hurt your life. So I feel like the person who wants to hurt me can put up a bounty on me or something similar. But I feel like suddenly everyone who has this information about me will be given an incentive to make a lot of money by hurting me. Whereas most people would not have any incentive otherwise. That was a bit abstract. That was a bit vague. Zvi, I'm curious if you have a similar sense that most people around me are not looking to hurt me?

Zvi Mowshowitz: Okay. Robin is basically saying this can't possibly be a first, best solution.

Ben Pace: Yeah.

Zvi Mowshowitz: There's always going to be a better alternative.

Robin Hanson: I'm arguing for the simpler policies relative to this strange complexity.

Zvi Mowshowitz: Okay. I think that obviously this is a strange complexity to introduce. It is weird. It doesn't appear in very many other places. However, it gets rid of different complexity, which is that we would then have to enumerate exactly.

Robin Hanson: But that's going to be true regardless. Think about the analog with prostitution. Say you say prostitution should be illegal as long as the prostitute makes the offer but not if the John makes the offer. What social problems are solved by saying only if one side makes the offer is it legal, but not if the other side makes the offer? That's what we're doing with blackmail. NDAs are exactly blackmail, but the other side initiates the offer.

Ben Pace: I'm confused. I feel like in a world where blackmail is legalized, you don't expect me to get any blackmail anymore or something. Because you currently think it's one way. You currently think that anyone who has information about me that could hurt me is very willing to sell it, but they're not allowed to blackmail me into getting it. But other people are not allowed to buy it.

Robin Hanson: Just to be clear, if any feel inclined to blackmail Ben Pace, for example, what you do is you go to him and you say, "I know this thing about you. I'm sorry to hear this thing about you. Bad things could happen to you. I sympathize with you. Let me know if I can help in any way." And then you walk away. And it's up to Ben to figure out that you're implicitly threatening him with revealing this information, but he never said anything. So the whole deal can go forward or not, but it's up to you to make the next step. That's what the law actually achieves—this indirect process. It doesn't stop the deal from happening.

Zvi Mowshowitz: It often does.

Ben Pace: Substantial reduction.

Zvi Mowshowitz: Robin. Concretely, right? If we legalize blackmail, do you think it would increase or decrease blackmail?

Robin Hanson: We would see it increase in the realm of people who aren't very smart. People who don't know what they're doing and don't know the law, which of course is limited because they might not even know blackmail is illegal now or know how to do it. So it will just force more people to hire lawyers to do things. It'll make more of a difference in situations where people aren't very smart and don't know the smart way to do things.

Zvi Mowshowitz: I strongly disagree. I think that legalizing blackmail would vastly increase the amount of blackmail that occurs.

Robin Hanson: Among the powerful. I think for ordinary people, it's pretty clear. Most of us can blackmail each other and get away with it, honestly. But we don't. The rich people who have a high public profile, they will be the people who will be most targeted for blackmail.

Zvi Mowshowitz: I actually think we do. This is actually another disagreement. I talked about this in my blackmail post, if you remember. I basically think that this form of blackmail is actually pretty common. You talked about the whole, "If you tell that secret, I'll tell mine, or if you don't do X I'll say Y." I think this is pretty constant throughout everybody's day. I think people are constantly threatened.

Robin Hanson: That's what I mean. The blackmail law isn't changing things much. If it's happening and it's already illegal, then the law against it isn't doing much.

Zvi Mowshowitz: I think that it is still holding it back substantially, and I think it would get much, much worse if made legal.

Robin Hanson: I still think the core thing here is, if it happens more, is that bad? And so, I want to compare it to the situation. You're saying it's because people want to harm other people. I say no; blackmail is people just trying to get money, and they are preferentially looking for things that make you look bad, but we already have people preferentially looking for things that make other people look bad. You didn't want to ban the situations I identified. You wanted to allow these situations where someone is trying to make someone else look bad as part of gossip. Why is adding more situations like that, due to allowing blackmail, on average, bad, while the first situation is good?

Zvi Mowshowitz: No, we can agree that there are things which we would like to discourage and see less of, but we don't want to ban, right?

Robin Hanson: Give me an example, and I'll tell you whether I agree with it. I'm not sure. I need to see.

Zvi Mowshowitz: I'm trying to think of an example that won't get me mobbed.

Robin Hanson: Sexual affairs. People blackmail about affairs, say the Letterman example. Most people, when they're thinking about blackmail, they are probably thinking about affairs.

Zvi Mowshowitz: You would not want to ban people from having affairs, I presume. You think that would be bad. Make it illegal.

Robin Hanson: I don't want to ban people from telling about affairs. If somebody knows that you had an affair and they think less of you, I think they should be allowed to think less of you because you had an affair. And they should be allowed to find out about it, if they find out in a legitimate way.

Robin Hanson: They can't bug you. You can't let them put a camera in your house to find out, but if they find out in some legitimate way, that you had an affair, they should be allowed to say it.

Ben Pace: To be clear, I think there's a question of scale. I am okay with you telling a bunch of people that I did something bad. I would feel bad if you managed to get it on the front page of *The New York Times*.

Robin Hanson: It depends on who you are. If you're an ordinary person, it won't get on the front page of *The New York Times*.

Ben Pace: If I'm a rich person who is not very important, if I just have a lot of resources to be given, even though it is not important about how I use those, then I think blackmail being legal will increase the chance since I have a lot of resources you can extract.

Robin Hanson: Well, why is it bad if *New York Times* readers find out about it, but not if other people find out about it? Why is that something that makes it bad?

Ben Pace: Because I think there's a level of punishment that information and gossiping, should do to you. People sharing gossip when it seems useful feels like it would accomplish this well. But people sharing it for as much resources as they could take would cross the line.

Zvi Mowshowitz: But also note that the thing we are often punishing when we punish a norm violation isn't the fact that the person violated the norm but rather that they got caught. We are explicitly saying that everybody violates these norms all the time, but if you got caught breaking the norm, it means you probably break this norm, and others, all the time.

Robin Hanson: And you think that's the only problem? You don't think that should be a norm. You think it's a mistaken norm. And I'm happy to accept there are mistaken norms, but I just don't think they're typical. I think most norms are productive and useful. Some norms aren't, and these things change over time.

Ben Pace: I think almost no norm violation of mine should be on the front page of *The New York Times*, and if I had enough resources, then it will get there, if you allow that.

Robin Hanson: Well, I think you want your public stance to be that you do follow the norms.

Ben Pace: No. My public stance is that I do sometimes break norms, and my life shouldn't be ruined because of it.

Zvi Mowshowitz: No. Regarding *The New York Times*, it's interesting that when I worked for a certain corporation, which I will not name, we had a principle that we could not put in any written form any statement that we would not want on the front page of *The New York Times*. And so, the very fact that someone might threaten to cause harm to us, or decide to cause harm to us, by sharing this, meant that we had to be much more implicit, keep less records, destroy evidence, be much less rational.

Robin Hanson: That's the general cost of norms. The norm system has cost. It's unique to humans. Other animals don't. But part of our success as a species comes from the fact that we have and enforce norms, though it definitely does have costs. One of them is we sometimes have wrong norms. Sometimes we miss-enforce norms, in that we draw the wrong conclusions about who violated which norms, and we may well punish too much or too little in other situations. But still, on average, norms are good.

Zvi Mowshowitz: Right. I'm not saying norms are bad, or that we shouldn't, in general, enforce norms. However, to be caught in a prominent way for violating a norm weakens the norm itself.

Robin Hanson: Let me make that analog to law enforcement. You couldn't say, "Well, if somebody is found guilty of a crime, they have to be found guilty just by somebody noticing it and reporting it without any financial incentives whatsoever." But we don't do that. We create financial incentives for people to discover and report crimes. That shows that we don't just rely on simple gossip for crime. We want to add and strengthen the incentive to report it.

Zvi Mowshowitz: So, let's compare two possible incentive schemes: Incentive scheme one is where I discover that you've committed a crime, then I can blackmail you and maybe get money from you. And incentive scheme two is I discover you committed a crime, go to the police, and the police give me a reward.

Robin Hanson: Or there's the third one, which is you have the choice between the two. You could do both.

Zvi Mowshowitz: Right. If I get rewarded by the police then to the extent that society considers it important to catch a particular norm violation, I'll be incentivized to catch that particular norm violation. Right? So, if I get a big reward for murder, I'll investigate murders a lot more.

Robin Hanson: Society is via the legislature, but society via gossip will also be the channel that punishes norm violations

Ben Pace: But you would not get punished in proportion to how much society cares about you. You'd get punished in proportion to how much resources the person has. Those two are not the same thing.

Robin Hanson: No, that's not true. The resource sets the limit, but of course, if it's a mildly embarrassing thing, they won't pay very much to pay off the blackmailer. The amount of how much you pay off the blackmailer is proportional to how bad it would be if it got out.

Zvi Mowshowitz: It's based on what would actually happen to you as a result of that, and how much you care about that thing, and the resources that you have, and so on.

Robin Hanson: Right. So these are two different channels. I'm endorsing both channels. We have a formal legal system for crimes, which are formally punished by legislatures deciding the sentence, resources of the police, and fines. But also, we have a system of blackmail, wherein people are paid financially and other ways for their doing things that the audience will disapprove of. Those are two different ways that the larger would disapprove and discourage things.

Ben Pace: I also feel that the level at which society wants to hurt you, or get your private information, is not really often in proportion to how much they think it is just to hurt you. Whereas a lot of tabloid press that will just hurt you because it gets them attention, and they can gossip in a big conversation in a way that I don't endorse, and I think most people don't endorse.

Robin Hanson: Well, then, why not make that illegal then, Ben?

Ben Pace: I think illegal is a strong and silencing thing.

Robin Hanson: Well, that's true for blackmail as well.

Ben Pace: But I don't think we should positively incentivize bad things.

Zvi Mowshowitz: No, I don't think it is. Why does blackmail have a chilling effect?

Robin Hanson: They all have a chilling effect. The question is whether it's too much or not enough.

Ben Pace: But Robin, is your stake that the tabloids shouldn't do that to people, and it should be illegal, or is your stake that, no, we should have this and encourage it more with more money.

Robin Hanson: I'd say that, on average, the tabloids exposing things about people is a good thing. It goes wrong in many particular ways, but I do not want to ban the tabloids from writing exposes.

Ben Pace: I don't want to ban them from writing exposes, but I currently think the situation is kind of like blackmailing in which they will extract way more resources than is proportional, and do massive amounts of damage.

Robin Hanson: I don't see that.

Zvi Mowshowitz: I don't want to ban the *National Choir*, but I do think the world would be better off without it.

Robin Hanson: There's a niche there. If they weren't there, someone else would come and take their place.

Zvi Mowshowitz: Yes, and I don't actually want to get rid of that specific publication because another publication would take its place, but I do believe the world would be better off if that sort of thing was not present, somehow. However, I don't see a way to enact law that would do that without having a very bad secondary effect that I do not want.

Ben Pace: I also think, Robin, you are definitely encouraging more of that stuff. You currently think it'd be better if that whole class of thing, that is, press doing hit pieces on people, was incentivized more.

Robin Hanson: I mean, over the last few decades, technology has made a lot of things very easy—surveillance, publishing, etc. So it is easier to find dirt on people and to publicize it, but I don't think that the way things were 20 years ago was the optimum and therefore we're worse now, right?

Ben Pace: Okay. So you're not sure whether or not this stuff has been bad, and that seems fair. Tabloid press, I feel most confident on, that has gotten worse.

Zvi Mowshowitz: So what I've been trying to get at with these examples is that there is a difference between the things that blackmail encourages and the things that blackmail discourages. There is also a different proportion of harms and benefits, but they include them in very different ways: from the things that are encouraged by other systems of gossip, other incentives to say things. And I think, in a sense, it makes those things much more negative. In particular, the incentive to entrap, the incentive to create negative material, to induce norm violations, is much much stronger under blackmail. And the fear of such things happening, in general, and the cost of navigating blackmail situations, and the fear of having to deal with these things is very bad. I would feel very stressed living under a blackmail legal regime. So would my wife, for example.

But also, I've been trying to draw a distinction between something that I dislike and would want to see less of, and something that I would make illegal, which I think that Robin would, in general, agree with. But we've been sidetracked before I could nail him

on that, as such. But the idea that there are many things that I disapprove of, and I think that if we had less of them the world would be a better place, but there is no law that I can pass that banning it would not have other effects. But I can try to tax it, right? I can make it more costly. I can make it more inconvenient in ways that discourage the behavior, and maybe that's even better than banning it, because now, when it was really worth doing, it happens anyway.

Ben Pace: Okay, so the summary there was there is a distinction between banning things and taxing things.

Zvi Mowshowitz: Right. And by tax, I don't mean literally, "You must pay 5%." I mean, "I want to make this more annoying for you. I'd ideally like to tax it, but because that's not practical, I can make it harder."

Ben Pace: Why is that better than banning?

Robin Hanson: What was wrong with the banning situation?

Zvi Mowshowitz: What's wrong with banning is sometimes there would be good reasons to do the thing and you would be willing to pay the tax, and we don't want to necessarily discourage that too much. When the government bans things, in many cases it has negative side effects. I don't want to get into why I'm a libertarian in this podcast, but I am and I believe Robin is as well.

Ben Pace: I would imagine Robin would agree that sometimes it's better to tax things rather than ban them.

Robin Hanson: I definitely agree. Although it adds another parameter that you have to optimize so you can get it along, but yeah, on average, it's better.

Ben Pace: So let's allow for some final statements before Q and As.

Robin Hanson: My claim was not so much that blackmail is good but that no one had offered concrete, clear, consequentialist arguments for why blackmail should be banned, especially relative to allowing NDAs in terms of who makes the offer. That's my strongest claim. And my claim, especially, is about coming up with explicit reasons and arguments. So it's about the fact that in society, we just have a lot of policies that, if you look at them, you're not sure what their justification is. If you ask people, they give you

various justifications that are contradictory. I think it's a sad situation that we don't have clear justification for most of our policies.

I'd say for this one, we do it, but we don't have a clear reason. Zvi mostly pointed out logical possibilities that it could be bad, but he hasn't actually offered data or concrete evidence that it is bad. He's just said it could be. I've tried to narrow down the difference by comparing it to gossip or asking why it matters who made the offer. In those cases, we have all the problems he pointed to of people faking evidence, exaggerating, or focusing on the most salacious details. These are all problems with gossip as we have it, in terms of allowing people to gossip as much as they want, unless it's false or violates privacy restrictions.

So my strongest claim will be, look, it might happen to be true that this extra category turns out to make things worse, although we allow many similar things. The parameter just happens to switch right there. But you haven't offered concrete evidence that it's actually bad, in the sense of, if we were just trying to collectively make policy on the basis of what our evidence is, we don't have evidence that says that's bad.

Zvi Mowshowitz: So just before we do a Q and A, I'm curious, what would you expect to see if I told you that someone had this kind of concrete evidence? What would you expect it to look like?

Robin Hanson: I would look at evidence on the consequences of gossip in various situations and whether, on average, gossip is good or bad. You could then say that this kind of gossip is bad, including tracking variables like how powerful is the person, whether it's sexual gossip, whether it's financial gossip, et cetera, and the kinds of audience it has, how big the audience is. Any sort of data on the overall harm of gossip, to me, is the key point, because blackmail is just the threat of gossip. If gossip is okay, then the threat to gossip is okay, and if gossip discourages bad things via telling people about bad things, then the threat to not-tell-but-I-have-to-pay also discourages. That's my summary point.

Ben Pace: Thank you very much.

Zvi Mowshowitz: With NDAs, I don't want to give my long spiel, but I do think that NDAs are necessary in certain ways to engage in certain types of activity, especially legal types of activity, and business types of activity, and that therefore we can't ban them outright, but I'm not going to get into details.

Robin Hanson: But as long as the other party makes the offer, blackmail is okay.

Q&A

Ben Pace: Here is Mr. R with a question.

Mr. R: I have a question for Robin. It seems like blackmail is probably a thing that goes on in the political stage, especially the world stage. Do you think that there's any sort of evidence for how many gains have been made using this? Whether or not the global theater is better off for it or not?

Robin Hanson: I don't have strong opinions. One thing that is noteworthy is that powerful organizations that are able to collect blackmail information and gain power. Famously, Edgar Hoover collected blackmail information about senators and other people in government and gained a lot of power over government officials through it. You could say that's a problem with the government and police since they both have the power and discretion to get a lot of information, which could be used as blackmail.

Ben Pace: But as of now, information is generally good because it's information on non-violating stuff that people want to gossip about?

Robin Hanson: On average, yes. But still, it creates an asymmetry in terms of who has the information. If only Hoover has it, then Hoover gains this extra power, so I would rather there be a more competitive market for blackmail. I don't want blackmailing monopolies. I would rather many different people have the ability to collect information by blackmail.

Ben Pace: You're a man of the people, I see.

Zvi Mowshowitz: I think it's very hard to know the order of magnitude of the effect that blackmail has on political action. It's one of the places where we're likely never to find out if the blackmail occurred, whether or not it succeeded. If it failed, we're likely never to find out the person was previously blackmailed before the information was revealed, whether or not they paid interim, but I think we should certainly be concerned.

We don't know what happened with Jeffrey Epstein, for example. We don't know who he did or did not have blackmail material on. Maxine Gladwell may or may not have

blackmail material, and it may or may not be severe. We have no idea their influence over policy, or how much money or policy was redirected. I think we should hope that the answer is not very much.

Zvi Mowshowitz: Similarly, if the Russians have information on Donald Trump or other politicians, and I'm not saying they do, we should hope that they have less rather than more.

Robin Hanson: Politicians, with the threat of blackmail and other powerful individuals do behave better. I think blackmail is making politicians behave better.

Zvi Mowshowitz: I think blackmail causes politicians to behave worse in a much more important sense, in the sense that they are likely to work for the blackmailers, to do what the blackmailers want, rather than what the people want. I think that's much more important than who they're sleeping with.

Ben Pace: Oliver Habryka, it looks like you have a question, perhaps informed by the conversation.

Oliver Habryka: Yes. I think one of the key things that I'm confused by is what concretely would Robin want to see? Maybe I will try stating a model I have, and you can just tell me whether it's right or not.

My current guess at what you actually want is something like: I think blackmail could be good. I think we want to start with repealing the existing blackmail laws that we have. A lot of follow up legal work will need to happen. And then we will need to figure out all the new things we have to outlaw. We will be opening up a bunch of behaviors that were previously illegal, that now are legal, a bunch of stuff is going to happen.

Your model is not: "let's immediately go and pass legislation legalizing blackmail." But rather something like, "let's try to forget all the new legislation and switch toward a new paradigm where blackmail is legal." Yet lots of other new things are now legal, which in total complexity, would hopefully add up to something less complex, or at least less confusing, than the current blackmail situation.

Robin Hanson: If the system could handle it, I would prefer incremental changes to see how they go. I mean, yes, there's obviously fewer chances of things going badly if you can make incremental changes. That is, you could pick, say, 10 areas that blackmail could be applied to. This year, we'll take the first area. The next year the second area,

and do one by one, and see how things change. And of course, the other thing is you may believe that making various sorts of new privacy rules would be a good substitute for blackmail. So you might want to think about, "Okay, if we allow blackmail about relationships, maybe we should just ban gossip about relationships, and maybe that will address the things you're really worried about," and I might be okay with that. I might be okay with new rules banning some kinds of gossip or requiring other kinds of gossip, because that's at least cleaner and more coherent than, again, banning deals for money if one side makes the offer but not the other, and still letting people say whatever they want.

Ben Pace: Cool. Habryka, is that good?

Oliver Habryka: Yeah.

Ben Pace: That makes sense. I guess one of my intuitive responses to Robin is that I can think of a lot of situations where blackmail makes me feel bad toward people with power, which I expect wouldn't have happened if it is generally disincentivized. A lot of people want to get money out of rich people, but they don't actually care about having control over the rich person's job or decisions. I feel like this punishes people specifically for their financial resources, and not for their actions. In proportion, how important it is that their power is good.

Robin Hanson: This might be related to my thinking. You can't trust your officials as much as most people think they do. Most people seem to think that if a judge doesn't have any obvious stake in one side or the other, we can just trust him or her to be neutral, and I just don't think, in general, our officials are as neutral or trustworthy as they are assumed to be in most ordinary cases.

Ben Pace: I don't think they can be assumed to be good in lots of cases, but I don't think that the norms of what will be most embarrassing for them is actually proportional to what is the most bad that they do.

Robin Hanson: Of course, if you have an official like a judge, for whom it's illegal to make the case go one way or another based on some other consideration, then if they're blackmailed and base their judgement on it, that is illegal, right? That is a different way of making that blackmail illegal. It's already illegal for the judge to take money to favor one side, or not to take money but keep their secret. If the judge makes a deal to favor one side of a case because somebody promises them not to tell the

judge's secret, that's already illegal. We don't need an extra blackmail law to make that illegal.

Ben Pace: I didn't quite follow that. Which part is already illegal?

Robin Hanson: Judges are not allowed to use personal considerations to make their decision on a case. Similarly, for other sorts of public officials. It's already illegal for them. So the scenario where the judge makes a decision one way in the case versus another because somebody says, "If you don't make the decision my way, I will reveal your secret." That is an example of blackmail that I'm not willing to legalize, even though I might make blackmail in general legal, but I'm not making it legal for that official to be biased in that case. That's already a law against that. We already have laws for public officials in such things.

Ben Pace: In general, I'm pretty open to creating financial incentives for giving information on public officials that could lead them to being fired. Like, if I am their boss, and I find out they do a bad thing, this financial blackmail is great. I feel like, otherwise, if you're just trying to embarrass them, or just trying to extract as much money as you can out of them.

Robin Hanson: Again, the problem isn't that they get embarrassed. The problem is that then their decisions are corrupted. That was what Zvi was trying to say. He's worried the public officials' decisions will be corrupted because of blackmail. And I am saying that kind of blackmail is already illegal for other reasons. It's illegal for an official to make these decisions based on personal considerations, and we could enforce that law much more strongly if we liked. I'm fine with public officials' lives being recorded and available for law enforcement to check up on.

Ben Pace: I think, overall, you would find that those decisions will get made more for personal reasons if you allow blackmail, because of the amount of pressure on them.

Robin Hanson: I'm happy to turn up on the dial on whatever mechanisms we have for preventing that. If we're not doing enough now, let's do more.

Ben Pace: But you should punish the civilians who become blackmailers, which is the current situation. Whereas if that's not allowed then we don't need to be able to punish people for accepting the blackmail.

Robin Hanson: I mean, they're already violating these rules, and public officials are being paid off in many ways to do things. That's already true.

Ben Pace: Jacob has a question. And then Alyssa Vance has a question.

Jacob Lagerros: It's perhaps more of a confusion than a question. Ordinarily, in markets, we want things to be priced appropriately, which is partially accomplished by there being competition among suppliers to drive the price down and some competition among buyers that pushes the price up until these things meet. But I'm really confused by how similar mechanisms would look like in the case of blackmail. In particular, it seems like my secrets are valuable to me precisely because they are about me. It's not clear who I would compete with for something to ensure I could actually get a fair deal. And it just seems like much more one sided of people just trying to push me to my limit of what I'm able to pay. So is this more of a confusion than a very crisp question, but I'd be curious if Robin could address how blackmail markets will reach equilibrium and why information should be appropriately priced.

Zvi Mowshowitz: I've never been in a blackmail market, but basically there's only one buyer and the hope is there's only one seller. And so, the seller wants to sell the information because, presumably, that information is not that valuable to the seller. What the seller can get for it on the open market might be zero. It might be less than zero. They might have to make an effort to share the information at cost to themselves rather than get a benefit. And therefore, their alternative to a negotiated agreement is very low. They get little or nothing, or even a negative amount. But your alternative to a negotiated agreement is also very bad. You lose because the information gets out and presumably you don't want that. And sometimes there is someone else offering to pay them, at which point their alternative to negotiate an agreement is much higher and therefore the negotiating position is much better. And therefore, the situation is that they have to guess how much you'd be willing to pay. And then you have to decide whether to pay.

One complication, of course, is that in general, there's no way for them to destroy the information credibly. They can't assure you that they will never come back to you for more. And therefore, the amount you can collect based on that is lower than it would be if you could reliably have it be a one-shot. There's also the danger that because of the fact that you paid them, you will be vulnerable to further blackmail on multiple levels. It's evidence of your guilt of something. It's evidence of your willingness to pay. It's evidence of your general susceptibility to blackmail in general. It's evidence of your wealth. It's evidence of your vulnerability. It's evidence of your lack of power.

Basically what's going on is they're trying to guess how much you'd be willing to pay and then ask for an amount that their returns are maximized. And keeping in mind the cost of failure. The cost of failure might be that you try to hurt them. You might even try to kill them, try to defame them, they might have to go jail. So they will choose a number. If blackmail is legal then potentially, there are no consequences to failure on their behalf. And they are incentivized to ask for an amount of money you probably don't have and probably would be unwilling to pay.

Because if you ask for a million dollars instead of \$10,000 and you succeed 3% of the time, you're way ahead of the game. So in the same way that most people don't ask for enough money in most negotiations, chances are the correct answer in an illegal blackmail scheme would be to ask for an amount of money you probably can't pay in many situations or more money than they could ever use in their lives. If you're trying to blackmail Jeff Bezos, you should ask for \$50 million. Because why would you ask for \$5 million? He is not that much less likely to pay you \$50 million. And then you see what happens. And sometimes you overshoot and you face the consequences. Whereas if it's illegal, you want to be much more careful and therefore your position is much worse so you ask for a lot less money; it's much less lucrative. That's all speculation. Maybe Robin has a different point of view.

Robin Hanson: I basically agree with Zvi. It's not a very competitive market, but I don't think that's terribly important. Most legal cases where A sues B, it's not a competitive market; it's more of a bilateral monopoly. Although in the case of a blackmail, as he indicated, you don't want to pay all at once. You need to pay over time. And there's always a possibility that somebody else could find out this information. In which case, there's now two people trying to sell the information and would likely need to take a cut. But I wouldn't be counting on any particular competitiveness. The main point I would make is that when you're relying on law to enforce things, it's just generally very expensive for legal officials to find out things about ordinary people. When the people around you have a financial incentive to find out and enforce rules, the cost of enforcement goes way down because they get this information in the process of doing other things. And so, arguably, blackmail-based enforcement is just much cheaper at discouraging who will do things than formal criminal law or other sorts of legal responses.

Ben Pace: Cool. Thanks. I'd be interested to hear Robin make this whole deal more palatable. Can you just name a few concrete, good things you think will happen if blackmail is legalized?

Robin Hanson: Sure. My favorite movie of the year is called *The Assistant*. It's about a sexual harassment case. And I think those cases would go better if blackmail were legal. In the movie, the woman actually can't do very much about it, unfortunately. But blackmail would allow her better options. I think in the case of, say, David Letterman, who famously was blackmailed for having affairs, if he could have been blackmailed, he would have behaved much more appropriately. And these were employees he was having affairs with, which makes it much more morally questionable. So I think there would just be a lot less sexual harassment if blackmail was legal.

Robin Hanson: Okay. All right. That makes sense.

Zvi Mowshowitz: I haven't seen *The Assistant*, but in preparation for this debate I read the *Wikipedia* summary, because I knew Robin would talk about it. Based on the plot summary, I actually disagree that it would have helped in that particular case.

Robin Hanson: It wouldn't have hurt. It wouldn't make it worse. It might help. It might not.

Zvi Mowshowitz: I don't know if it would have had an effect. I think that she, in fact, attempted to share the information as best she could, and was told that nobody cares.

Robin Hanson: She was discouraged, but if she had monetary incentive, she might well have pursued it further; she could have gotten money out of it.

Zvi Mowshowitz: When faced with the threat that someone might learn about it, he instead yelled at her. He got mad at her and threatened her.

Robin Hanson: Because he had the power and she didn't have a credible threat. If she had the credible threat of blackmail, then she could have used it.

Zvi Mowshowitz: Wait a minute, blackmail isn't a credible threat. The credible threat is to share the information. So she absolutely had a credible threat; she actually tried to share the information,

Robin Hanson: She tried this out with her HR department. She didn't go to the press. She didn't think she'd get much benefit by going to the press.

Zvi Mowshowitz: So how does her failing to have blackmailed her boss changed that?

Robin Hanson: She would threaten to go to the press.

Zvi Mowshowitz: Why didn't she do that anyway?

Robin Hanson: She correctly anticipated that going to the press wouldn't actually get her that much. With blackmail, she'd get a lot more.

Zvi Mowshowitz: My interpretation is her boss would have said, "Well, actually, I'm going to blackmail you because if you don't do this, I'm going to tell every person who calls for reference that you're a terrible person to hire because you caused trouble."

Robin Hanson: Right. He already had that threat. We're just adding another possible threat for her side. So we're moving the negotiation. That's another, that's a concrete example.

Zvi Mowshowitz: I was just trying to say that I actually think blackmail would shift power in his direction, not hers.

Robin Hanson: So there are a lot of powerful people who break a lot of rules in many ways. And then the people around them shut up about it, let them get away with it because they don't feel they actually have a credible threat to report it. So they don't. Blackmail would mean a lot more actual reporting or a discouragement of the things powerful people do that break rules and norms.

Ben Pace: Alyssa, if you have a question, I'd be interested to hear it.

Alyssa Vance: So there are many questions I could ask, but here is a pretty simple one. In general, do you think selling people's addresses should be illegal? Currently, I think it is legal. There are several companies that do it. It seems like it would be a relatively straightforward thing to ban since it's a relatively clear dividing line and they're only a few businesses engaged in it. It also seems mostly to harm people. The benefit to me of being able to find out where any random stranger lives without their consent, seems like something I don't get any benefit from. And then, the idea that anyone could find out where I live without my consent seems like it could potentially be dangerous, which is why I removed my address from all those sites.

Robin Hanson: I'm with Alyssa. I find it hard to think of any benefits to letting people find other people's addresses without consent, and can see a lot of potential harm. I'm open to counter arguments, but banning it seems pretty obvious to me.

Zvi Mowshowitz: I think that there's an argument to be made about letting other people send you physical mail, but I don't see any reason why you need their address to do that. The post office knows your address. So yes, I would support making that an opt in. You should be able to verify allowing people to find your addresses, and get a cut of the money being made from it.

Ben Pace: I feel like we're currently in a situation where some people would like to be able to find other people's addresses, and there's not much of a market for it. But if we legalize blackmail, this sort of market would grow.

Robin Hanson: There were these things called phone books that showed names and addresses. It was a common thing.

Ben Pace: I think Evan wants to ask a question.

Evan: I have a quick question. What happens in 10 or 20 years when everyone's just a little bit more used to limiting their privacy and there are less expectations of privacy? And say some database gets hacked; with future search tools, it really wouldn't be that hard for me to believe that you could just look up a bunch of names in a search query. You could find out who has looked up some like edgy porn along with their names and addresses. And then you just send out a mass email. How do you guys see blackmail in the future when there's just so much more information about people in general?

Robin Hanson: Like I said before, I'm happy to have lots more privacy rules or rules about what you're allowed to gossip about. I think there's a big debate to have about what sort of information we want to prohibit people from sharing about other people, which is separate from whether you have these extra blackmail rules in my mind. I don't see the extra blackmail rules doing that much, but I would be happy to see a lot of potential new privacy rules.

Zvi Mowshowitz: I'm a big fan of free speech and letting people talk about things when they want to. I would be loathe to say that you couldn't share some information, but if you hack into a database and steal someone's information and share it, that seems like an area where it would be reasonable to prevent. And I certainly don't think the public

has a right to know about what adult porn you are viewing. It is not more important than people's needs for privacy.

Ben Pace: What do you think about the current margin, Robin? Would changing the blackmail laws and adding a bunch of informed privacy laws do damage, enhancing the problems of our current lack of privacy laws?

Robin Hanson: No.

Ben Pace: All right. Asked and answered. Dunka, would you like to ask a question?

Dunka: Yes. I have a question for Professor Robin Hanson. Zvi has mentioned that one difference between blackmail gossip is that with blackmail people are incentivized to induce others to violate norms so that they have information to blackmail them with. And because of that, it's not clear to me that making blackmail legal would lead to a net reduction or increase in norm violation.

Robin Hanson: I would say the world of gossip is a scarier and more dangerous place than you might realize. Setting aside blackmail, people who gossip about you are often trying to make you look bad and trying to entice you into doing things that then they can gossip about. That happens in corporations; it happens in social groups; it happens in politics. People already want the gossip about each other and especially like negative gossip. You can witness this in the current cancel-culture problem. Often the things people get canceled for are things that other people induce them to say, like some guy making the okay sign because somebody else was making it and asked him to. So he makes the okay sign on video and gets fired for it. But not because he ever knew what the symbol meant; he was just doing what somebody else wanted him to do. This sort of thing happens today, all through gossip. It's one of the costs of gossip. I don't see blackmail substantially adding to the trade off there.

Zvi Mowshowitz: I think it absolutely can. I think it's about magnitudes, and which things are the more common use cases, but absolutely. You're incentivizing strongly something, you're going to see more of it, relative to other things. What would you expect?

Ben Pace: That makes sense. We don't have time for too many questions. So Zvi, how do you feel about blackmail?

Zvi Mowshowitz: I'm against it, Ben. I'm against it.

Robin Hanson: I think even Zvi will agree that most people, when they give an argument about blackmail, don't give a very good argument. So it's still an example of the sort of thing people have an intuition should be wrong. And then when they give arguments, they make up a bunch of vague stuff that doesn't actually make much sense.

Zvi Mowshowitz: I think most people end up doing a mix of a lot of stuff that makes no sense, and provide weak versions of arguments that would make sense if they thought them through more carefully.

Robin Hanson: That's an interesting, distinct feature about this policy area.

Zvi Mowshowitz: I'm curious, Robin, can you think of that in a policy area where people actually give good arguments?

Robin Hanson: Well, when people argue the need for taxes to provide the government with revenue, I think that works.

Zvi Mowshowitz: But there are arguments about how much to tax people.

Robin Hanson: Of course, they say rich people have more money, so we should tax them, because that's where the money is. And that does work.

Zvi Mowshowitz: I mean it's less bad, and it's less wrong..

Ben Pace: Do any of you want to provide closing statements?

Speaker 1: I regret not doing a before and after poll here.

Zvi Mowshowitz: I trust this crowd to actually remember what they thought before. If we wanted to do a poll now.

Ben Pace: Robin, do you want to make any final statements? I would be interested in both of you providing other thoughts that you have.

Zvi Mowshowitz: I have many more things I could say. I think that we've reached a natural conclusion, but I will stick around if people want to keep talking.

Ben Pace: And that seems great. Robin, any other things you'd like to add?

Robin Hanson: I just want to say that we can come down to this category of various cases, and his claim is that this extra set of cases are different from the first set of cases. I don't think he's really offered a direct argument that it is much better. He's made the claim, and often policy comes down to this. In the end, if the parameter went one way, one side would be right, and if the parameter went the other side, the other would be right. They take all of their time discussing, just getting clear what the parameter is, but in the end you can't resolve it because you don't know the parameter value.

Zvi Mowshowitz: I can offer my models for the mechanisms I believe we are operating on, and how often I think various things are harmful versus helpful. I can go into much detail and I could write that up and we could talk about it. But in terms of what Robin would like to be an actual study about the harms and benefits as such, I can't think of how one would go about that study to find the conclusion.

Ben Pace: Great. Cool, thank you very much. All right. So this is definitely the official end.