

How AI Breakout Harvey is Transforming Legal Services, with CEO Winston Weinberg

54 MIN · YOUTUBE · [HTTPS://WWW.YOUTUBE.COM/WATCH?V=EXK-_YYQDMM](https://www.youtube.com/watch?v=EXK-_YYQDMM)

https://www.youtube.com/watch?v=EXK-_YYQDMM

SUMMARY

Winston Weinberg, co-founder and CEO of Harvey, discusses the importance of trust and prestige in professional services, particularly in the legal industry. He emphasizes that by partnering with prestigious law firms, Harvey can build trust and credibility, which is crucial for success. The company focuses on creating AI-driven workflows that enhance legal processes rather than simply replacing human lawyers.

- Trust and prestige are critical in professional services; earning the trust of top firms leads to broader acceptance.*
- Harvey aims to revolutionize the \$400 billion legal industry by leveraging AI to enhance workflows and productivity.*
- The company prioritizes building specific workflows for complex legal tasks, then integrates them into a cohesive product.*
- Harvey's approach involves continuous expansion and refinement of its product offerings based on user feedback and industry needs.*
- The legal profession is evolving, with a shift towards more strategic, advisory roles for lawyers as lower-end tasks become automated.*
- Harvey is focused on improving access to justice by making legal services more affordable and accessible through AI.*
- The company faces challenges in maintaining speed and adaptability in a rapidly changing AI landscape.*
- Weinberg highlights the importance of hiring young talent and providing them opportunities to innovate and adjust quickly in a dynamic environment.*

Something that's really important in professional services is prestige and trust, right? The reason prestige is so important is because trust is the most important thing in professional services, right? And so the reason we went after the larger firms is if you earn the trust of a few of those firms, the rest of them will trust you and the rest of the firms downstream will definitely trust you, right? And their clients will trust you, right? So I think that something you know something that we thought about doing in the beginning was well just go straight to enterprise right and there are a bunch of problems with that but one of the main reasons is there's just no reason for them to trust you right that you can actually build these systems [Music] Greetings. Today on training data, we have Winston Weinberg, co-founder and CEO of Harvey. Harvey has occupied a special place in the AI ecosystem over the last two or three years, becoming the canonical example of what it means to build an application layer company on top of the foundation models. A couple years ago, when these companies were derided as rappers on top of GPT models, Winston and his co-founder Gabe realized that this is where a lot of the value would actually be created, dealing with the messiness of real world problems, not replacing human beings, but giving them superpowers. And so over the last couple of years, as much of the industry has taken a

tech out approach to build their business, Winston and the team at Harvey have taken a customerback approach and now stand position to revolutionize the legal industry, which is \$400 billion in the US alone, the same size as the global cloud markets.

We hope you enjoy. Winston, welcome to Training Data. Thanks for joining us on the show. Thanks for having me. All right. So, you have built you and your co-founder Gabe and your team have built a company that has become to many people sort of the canonical application layer AI business or sort of the defining application layer AI business. But when you started back in July of 2022, yeah, end of July 2022. So when you start in July of 2022, pre-Chat GPT nonetheless, um this wasn't really a thing. And as soon as it became a thing, people looked at this whole category and they said, "Eh, these are just rappers on top of foundation models. There's no value to be built here." So they're immediately sort of derided. So I guess the question is as a starting point, what gave you the conviction that this was the right business to go build and that there would be value between foundation models and customers? Yeah, I think the simplest answer to this is quite literally that these industries are incredibly messy, right? Um, and so I think that the largest misconception with kind of the the GPT rapper companies was a lot of what you were starting to do didn't have like a massive delta between what the foundation models provided and what your small company could do, right? So one of the largest advantages that we had in the beginning was how much emphasis our product had on citations. like literally just it was very important in the legal industry to be able to site like line by line and the accuracy of those citations, right? There were a bunch of other things as well, but like from day one that was something we put resources into. If that's all our company was, if we just said, "Okay, we're going to be a citation company, right?" Then yeah, I think you're just a GPT rapper and it's going to go away. But if the ambition of the company is to actually partner with an industry to completely transform it, this is a trillion dollar industry. It's incredibly messy. There are tons of different data sets. There's tons of ways that people work. There's tons of different specialized workflows, etc.

Um, and so I think over time it became clearer and clear that these models weren't going to just automate entire industries away, right? They were going to change and like fundamentally serve as kind of like the groundwork for changing an entire industry. But I mean, one way to even look at this is when a new feature comes out, we usually think of that feature as this is a great piece that we can put into part of a workflow.

So an example of this is like OpenAI will release deep research, right? And my first reaction to that is, oh my god, there's all of these capital markets use cases, right? For like studying markets, looking at all of this data that's on the internet, etc. It's not everything a capital markets attorney does, but you could put that into a piece of a hundredstep process in your product and it makes that process better. And so the best way to think of this is literally whenever I see either the foundation models can do something better than they could before or completely new feature.

It just like unlocks part of the TAM or part of the actual process that you can build on top. Let's pull this thread a little bit. So, at time zero, you got the model down here, you got the customers up here, and you got Harvey doing citations in the middle for for a month or so. Yeah. Let's um what came next? Kind of walk us through some of the major leaps in the product or some of the major value creation opportunities you've seized upon along the way. Yeah. So the best way to kind of describe how we've been building the product and this is I think

one of the things that we've gotten the most right and it has also been the hardest to manage which is and from an org perspective and from a kind of like how which pieces do you focus on etc. And and the piece here is at all times you have to basically expand the product and then collapse it back, right? And so what I mean by this is actually if all of these models work perfectly and humans were perfect at communicating with each other, the best interface is literally email. Like that's it. Like there is no interface. It's literally just email and it's perfect. It has all of your contacts. It can read your mind, etc. I mean it could just be a neural link. Yeah, sure. That's perfect. Yeah, that's even better, right? Um but that's not how it works, right? And I don't think that we're just going to randomly oneshot be able to do that, right? And even if we could do that on the model side, we can't do it on the human workflow side. In other words, like the example I give a lot is even if the models could do, you know, somehow chain all the steps together to do a large merger like the Activision and Microsoft merger, the user can't send that to the model and say merge please and then it just like does all those steps, right?

There's all of these communication sides on the UI side as well. But okay, so holding that still, going back to the, you know, you have to constantly expand and collapse the product when you're expanding. What I mean by this is the chat UI doesn't work for every use case, right? And it doesn't work for every use case right now. I don't think it will for the future, right? And so an example of this is if you were trying to build something that does really good case law research, right? There's multiple steps to that. You want to build a system that is very good at doing retrieval over all the cases. You want to build a system that is really good at comparing contrasting all the cases. You want to build a system that is really good at synthesizing the facts in your case to all of the case law, etc. Right? And if you're doing that, the best way to do it is you expand your product by building specific vertical, you could call it agentic systems, whatever you want that do that thing from start to finish, right? And then you collapse them all together. So what does this end up looking like? We will take a bunch of use cases that are very high value.

We'll build out a specific workflow to do that use case and then we will chain them together so that you can complete a task from start to finish. Right? And so the difficult piece of that is if you're trying to sell something that is seats, you have to sell something that is applicable to as many users as possible. And so you have to balance are you building something that is really good for a securities attorney or are you building something that's really good for all of the attorneys, right? And so that's kind of the collapse part where you want to build these specific workflows, agentic workflows, whatever you want to call it, whatever terminology it works, and then you want to combine them into the same service level on the product, right? And so what this looks like eventually is you upload a share purchase agreement onto Harvey, right? And we have the user may might not see this but there's tons of different workflows that are like extract the reps and warranties from that or summarize it or whatever it is right and we've built them all separately and you can use those UIs and actually do that workflow separately or when you just upload an SPA Harvey says would you like to run any of those workflows and that's the collapse version right so you're building these specific solutions and then you collapse them back in how much of the magic do you think is in the models themselves versus what I what I think you just described as you know an agentic system or a cognitive architecture. Yeah. Um and you know how much of it is it the workflow that's hard? Is it training things together that's hard? Is it the combination of all of it that's hard? So I think like I would categorize this in three areas. So the all you need to do for kind of like every single workflow is basically what does the user want? What is the intent right? Um after that it's so what is what do they want? What's the intent?

How do you extract that intent out of them? The second piece is what context do we need? And the third piece is is this right. Right? And so my point is there are different systems that work really well for different versions of those patterns. Right? So routing like using the models to actually like do predictions and routing is really helpful for the first one. Right? And so that's the example of you know either follow-up questions is really important there. routing uh someone's query to the like particular task you think they want to do is very good. Um and so there's kind of like an orchestration element there, right? Context is okay, do you have predefined systems that will search for your internal documents that are relevant to the question and then your external documents that are relevant to the question? A lot of that work is retrieval, right? like most of what you're building there is retrieval and then routing to make sure that you're accessing the external documents when you need to and the internal documents when you need to as well. And then the third one is is this right? Right. So going back to my citations example which is is silly but it's it's actually really important. Um I think there also a lot of the work can be done by the models but you have to make sure that the models are very good at checking for certain things. So let me give you an example of this.

uh there's a thing in legal that is what is market right and the models don't know what market is and there's various versions of what market is there's market for a particular private equity firm like what terms they are used to you know doing on an LBO or a dealer in a sideter etc there are terms that are across all of private equity and then there are just kind of like general M&A terms right and the models don't have access to this data right and so on That third piece, a lot of it is can you build a system that is very good at retrieving all of those different data sets when needed and comparing them against each other. And I'm sure it helps if you have all of those different folks as customers. Yes. Yeah. So I mean that I think is goes to another piece which is the biggest problems that we have and I think I mean is a very interesting problem and I don't think the models are just going to solve this is the process data for a lot of these tasks doesn't exist on the internet right so the process data for how do you do disclosure schedules or what is market right those are not things that are just kind of like on Reddit somewhere right um and so what we do is we actually hire domain experts who sit down and say these are the steps that I would take right and then you just chain the models on top of that right and or if there's a gap between what the model can do then you do fine-tuning right but the best way to do fine-tuning and post-training is task specific it's not here there was a a large kind of ethos maybe in the legal industry and and a bunch of other industries that somehow if you got all of the legal documents and then you just trained a model on than that it would do law, right? Um and and that's kind of like saying read all the, you know, your case load and all of the textbooks in law school and then just put you into like a profession and you somehow know how to do it. And that's not how it works. So much of it is like you actually have to learn how to do the different steps, right? And then the other side of that is not only is the process difficult, but the evaluation is really difficult, right? So you have to hire kind of like mid-level folks to do evaluation for a lot of these things because if the junior folks could do evaluation, they would be mid-level or they'd be senior, right? Like that's the reality. And a lot of what you do in a law firm or professional services is you're actually evaluating the work of junior folks, right? It's incredibly expensive. I'd argue that maybe 20 to 30% of the revenue of these places is doing that. I want to go back to this concept of expanding and collapsing the surface, which I love. Um, what is your view of the ideal end state for how, you know, a lawyer should be interacting with Harvey? Is it kind of that email chat interface and it's just merge merge company A and B and we're good or I don't I don't think that we will get there anytime soon. And that is not for me to say that I'm not bullish on the foundation models getting better. I mean I you know this I am incredibly bullish on the foundation model is getting better and we have designed the

company kind of without a constant push as a driving force. Um I think the biggest problem with that is that doesn't allow the user to have enough kind of exercise enough judgment on the workflow itself. Right? So, I think one thing that when people are talking about agents, they're talking about these tasks that are like quite simple and don't have massively high like economic value, when we're talking about building workflows and putting kind of agents in those, we're talking tasks that cost hundreds of thousands of dollars, right?

I mean, one reason the the legal industry is so good for LLMs is if you can kind of think of this is is the industry textbased and then how valuable is a token, right? And a token is incredibly valuable in legal and professional services. If you look at like a merger agreement, like a 50-page merger agreement, the token in there, like how each piece of a word is worth so much money if you think about how much it costs to produce. So this is all to say I think the end state is you keep building these agents and workflows and then you chain them together as much as you can right and that makes it so that your UI actually looks the same but your suggestion and your routing model gets better and better your orchestration level model gets better so you can kind of think of this maybe like let's go to the law firm I think that we're building out the specialized associates that can do different tasks But it's also incredibly important that you have the partner or the managing partner operating model as well. And I think that as the models allow us to build more and more of these specialized kind of specific associates, you also have to put all of this effort into actually having the orchestration layer that pulls all that together, right? And so I think our UI actually looks similar in the sense that it it does look like a kind of like a text window, right? But the ability for a user to upload a bunch of documents and for just to suggest what you do or the ability to say this is what you did last time. Would you like to do it again? Things like that improve kind of like a colleague that's just like really good at knowing what you want to do. On that note, are you selling software or are you selling work? Is this legal software or is this AI lawyers? Yeah. So, so far I mean software um I think that the best way like most of our clients in the beginning were law firms and professional service providers like PWC.

Um we now have a lot of enterprise customers as well and the way that the product is evolving is you can kind of think of it as actually two products. One is a productivity suite, right? And that is lawyers in the loop at all times etc. Um and you know the ROI on that is is incredibly helpful. It saves you tons of hours a week, etc. Right? The other side of that is you're building these workflows that do part of the work from start to finish, right? And that is closer to selling the work. The way that we're actually approaching this is we are building those with law firms and helping them get more business, right?

So we will get these revenue split agreements with law firms or for professional services and we will combine their domain expertise with our tech and then they go out and sell it to their clients, right? Um and so we are transitioning from just a seatbased company to actually selling the work as well. Uh and I think we'll see a lot of that this year. What and what are the ripple effects of that kind of like inside the building and outside the building? How do you how do you manage that as a company and what change does that require of your customers? Yeah, so managing it internally is all about taking bets. Um because it's it's you know if you're building software that is and every single feature that you're adding is good for your entire user base, you've kind of like reducing your chance of failure, right? because every time you add something to the software, it is hopefully uh increasing the value that you get from every single one of the users. If you focus on building one of these specific workflows and it doesn't work, you've b like it's zero and one like they it either works or it doesn't,

right? And it's for a specific use case. The value of it is really high, but if you can't sell it to people, then you're in trouble, right?

Um, and so the best way that we've handled this internally is we've actually let the law firms do a lot of the discovery for us, right? So we will do joint projects with a large company and their law firm, right? And so we will use them as design partners to make sure that this is something that is repeatable. You can actually do it, right? Like the technology is there to do it. And number three is appetite. So there is a large problem in legal where a lot of kind of what you need is like compliance checks or you're doing the work for insurance etc. And so the the third piece is also really important where you have to make sure that the in-house team is actually okay with AI handling that use case. Right? So that's kind of how we handle it internally where we spend tons of time on discovery, right? Tons of time talking to customers and we have a lot of design partners for the early stages of that.

One more piece about how we handle it internally as well. In order to generalize, what we're doing is we're building like AI patterns. So here are the 15 types of actions in the legal industry that we need, right? So research like case law research is a huge one. Uh regulatory research is another one. Uh clause extraction is another one, right? And so we'll put tons of effort into those like kind of generic or widespread horizontal things and then we'll chain them together kind of like a Ford factory line. So those are the two ways. Uh so on the GTM side it's discovery. On the product side, it it really is kind of building that Ford factory line of of AI uh patterns. Externally, I think this will have a very large impact. So, we're working with a firm that is almost a hundred years old and they became famous from literally advising I think it was King Edward the something on the abdication of the throne. So in other words, these firms are really old. Like some of these firms are like hundreds of years old and they have had the same business model for a hundred years, right? And they're actually willing to work with us and change some of their business model, right? And like take a bet on this and explore it, right? And so as I I think most people know, a lot of professional services and especially legal has been the billable hour forever, right? And so the problem with that is if you're selling them efficiency software, there's only so many hours in the day, right? unless you come up with a software that invents the 25th hour of the day, which by the way would be the best thing you could do to sell to professional services. But if you can't do that, um, efficiency is is a hard sell unless you can convince them that there are ways for them to actually transform their business, right? So, what we're starting with is let's help you take a bunch of these workflows that you normally do at not a very good cost or a loss, right?

um and let's turn those into software and help you kind of spread it out and get more market share in areas that you don't normally have. So, I'll give you a pretty good example of this. In private equity, one of the ways that law firms go and kind of get new private equity business is they'll do things like side compliance or kind of like lower-end work and they'll do that at a loss so that eventually, you know, the private equity firm will pay them for the LBO or or whatever the big M&A deal is or something like that. And so we'll build them software to help them get that work and not operate at a loss to get it, charge for a flat fee, etc. Uh, and then take some of the cost. Is the tech ready? Like how much I guess legal work can already be automated with today's models and if you were to freeze model development. Yeah. You know, how much can be automated? Yeah. Yeah, I mean I think that's a good question in terms of and especially the second one of I guess like how there's the reasoning ability of the models and then the capability re of the models, right?

And if you froze the models reasoning ability, I think we'd be fine actually. Um, so if we somehow kind of like froze the ability for the models to process data, um, that doesn't mean that they know which process to actually like analyze that data, but just their ability to ration over it, um, to make rational decisions over it. We'd be in a really good spot. So in other words, I think we're at like small percentage points right now of like legal and professional services, but if you paused, we would increase pretty high even if the models didn't get better, right? Because I think the reasoning ability of these models is is is there actually I think the bigger problem is like evaluation, improving process um and collecting more data too. The the data isn't there.

Is model development going to freeze here? Um, I I don't think so. No. Um, I mean, I I think we're we're seeing a pretty large evidence of if you throw more compute at the model and you let it make more and more reasoning steps, it just gets better and better and it improves time. Yeah. Exactly. Right. Um, and it's interesting if we if we go all the way back to 2022 in the early 2022, the thing that Gabe and I had done was we were doing a bunch we got access to GBD3. This is public by the at the time.

This is like early 2022. And we the thing that we found was we were just doing chain of thought prompts before anyone was thinking about chain of thought prompts. And the way that we actually started was we were doing that over a bunch of kind of like legal questions and we cold emailed the general counsel of OpenAI and sent him that. And he basically, his name's Jason Quan. Um, and he basically responded, "Oh my god, I had no idea these models were this good at legal." And I think the main reason people weren't looking at it is because they were just doing one model call over a set with like one static prompt and just calling it a day.

and they weren't doing what are the actual process steps like they weren't telling the model these are the steps that you need to take in order to answer this question right and so my point is that's what we saw in like early 2022 and that's like the direction that these models are going right um and I think that it's really good for legal and for kind of these these industries where you have to have these complex decision-m processes for two reasons one the model get better at that that's just increases the ability of them. But two is that process data doesn't exist. Like I said before, like it it's not online. Like how to book a flight is online, right?

Uh there are easy ways to train that. It's harder to do how do you do an LBL? Well, and you and you guys you now deservedly so have an unfair advantage around the process data because you have some of the very best law firms in the world as customers of Harvey, which I think was a contrarian strategic decision that you and Gabe made a couple years ago and you made it with high conviction. And and I I remember a time when there were lots of law firms who wanted to come work with Harvey and you basically said no so you could focus on some of these big prestigious firms. Um can you say a couple words on what gave you the conviction that that was the right strategy? And then maybe more importantly once you determined that that was the strategy, how in the world did you get them to trust you? Yeah.

like this is a very very scary new world and and and you managed to earn their trust. So how did you do that? Yeah. So I think okay so the reason for doing it um there was a a GTM reason for it and a product reason. From the product side the bet was that the models were going to get better and that you want to build systems that the the next generation model cannot do right and so you want to go after the incredibly complex international

merger type work. Right?

you want to do the very complex work and build systems for that because it is the most defensible by far. Right? So that was kind of the the product side of it. Um from the GTM side of it, I think that something that's really important in professional services is prestige and trust, right? The reason prestige is so important is because trust is the most important thing in professional services, right? And so the reason we went after the larger firms is if you earn the trust of a few of those firms, the rest of them will trust you and the rest of the firms downstream will definitely trust you, right? And their clients will trust you, right? So I think that's something, you know, something that we thought about doing in the beginning was well just go straight to enterprise, right? And there are a bunch of problems with that, but one of the main reasons is there's just no reason for them to trust you, right?

That you can actually build these systems, right? Uh, how did we do it? Um, I we did a bunch of things that do not scale at all. Like in in all honesty, um, I think that one of the things that we did really well is I don't think there is any excuse for someone who is building an AI product and trying to sell to not do hyperpersonalized demos. Like there is no excuse. I think it used to be really important to do that. Now it's paramount and it is so easy to do this, right? Um, and so one of the things that we did in the beginning was whenever I would demo to a partner, I would try to use something that they recently worked on.

Um, and then the other thing too is lawyers are argumentative, like very argumentative. Uh, and I mean that in the best way. So just let them fight with the model. Yeah, I'm serious. And I mean that in the best way. So I would sometimes say, you know, was this a good argument? Um, and how would you improve it? And if they were really bored on the demo and then you say that they are reading every single word that comes out of Harvey, like every single word. And you know, it wasn't always the perfect response, but I think it engaged them in a way that they've just never been engaged with software before at all.

Right. Um, and and we found I mean, one thing that's that's interesting is, you know, a lot of the older partners at at firms sometimes we might be their first AI product that they've used, right? And so it's really important to actually like show them, you know, kind of the basics as well, not just exactly what, you know, is is special about your product, too. How have you elicited the behavior change out of your customers and kind of taught them to use AI? Yeah, this has been really hard. Um, I think that it starts with product by far. So when I was talking earlier about the expand and collapse, the most important piece of the collapse is to make it so you don't have the blank page problem, right? Like that is by far the most important thing. So that when you go onto the landing page of Harvey, there are all of these buttons that you can click that will help you get started, right? And that is super important. And now we actually have it to a point where you put in how many years you've been practicing, what kind of lawyer you are, things like that, and it will change what that screen looks like when you start. Right? So that's been super helpful because if you make it so that those very specific systems are exactly what they do day-to-day, it is really easy to get them to start and if they start using it, they'll be creative.

They won't be creative right off the start. It's hard to get people to do that, right? Um so on the product side, I think that's the most important piece is just making it very personalized and making it so the time to value is

really short, like as fast as possible and then they'll go out and explore. on the CS side, we've hired a lot of lawyers. Um, and I think that's what been super helpful is you need to hire domain experts who can say and and we're doing the same too in tax and these other areas. You need domain experts who can come in and say this is how I would use it and this is, you know, I've been doing this for six years. I was in the same shoes you were in. Um, I know exactly what work you do.

Talk about the trust thing. Hallucinations. I think if you look at where AI has worked really well, it's where, you know, where hallucinations are a feature like in the creative industries, right? Hallucinations are a feature. In your industry, I'm assuming hallucinations are an absolute bug. And so, what do you do to make, you know, those argument that lawyers really trust? Yeah. Really trust what the model was saying. So, I'd actually push back on the second piece a little bit in a lot of what lawyers do is creative actually, like a lot of it. Um, so like I'll give you an example for this.

Litigation would not exist if there were just 10 really simple rules and everyone knew which fact pattern fell into the boundaries of those rules. Like there would be no litigation, right? And so I I was a litigator. I mean transactional is also incredibly creative, but litigation I think is just a better example because what you're trying to do is I have these facts. Here are somewhat of the rules of the game. Let's figure it out, right? And that's actually incredibly creative. Like very creative.

Um so but still who like accuracy is more important than than creativity to be clear here. Um and so on the accuracy piece so even given that creativity is actually super useful in some instances and we try to make it so that our product doesn't like get rid of that aspect. Um on the accuracy side there are a lot of things you can do to improve it for one. Um but for two the main thing is law firms are hierarchical. So you like a junior associate if they get a task they do the first draft of it and then a second year reviews that and then a fifth year might review that and then the partner reviews that and it goes out right um and so actually the you know minimal viable quality of your output can be lower for a law firm than it can be for an in-house team actually. So selling to the law firms was also helpful in the beginning because so much of the work gets reviewed, right? And so you aren't selling them something where they kind of click it and forget it, right?

They're actually in the loop at all times. When you're selling to an in-house team, it is better to sell them the specialized versions of tasks that they can kind of see an insanely high accuracy level on. Um, and the more specialized you are building, the higher accuracy you can get because it's easier to fine-tune. It's easier to do evaluation because there are less steps and the surface area is just smaller. Let's talk about the lawyer of the future and the law firm of the future.

Foundation models are going to keep getting better. You guys are going to keep doing valuable stuff on top of those models. What does the job of a lawyer end up being? What does a law firm end up looking like? Let's go like five, 10 years out. Yeah. So, I'll start with the job of a lawyer because I I just care a lot about this. Um I think that it goes back to actually what it used to be. Um so 50 years ago the role of the lawyer was an adviser and a lot of what they did is look around corners give a different angle of advice etc. And and we've actually seen this evolve to a degree where the CLO chief legal officer title is like pretty new and really it's like most of the CLOs's

that I have met they are part of the business like they are business drivers. it is not just the no person, right? Um, and I think that that is going to happen. So, I think what's going to happen over time is kind of this like lower-end work is going to get somewhat commoditized and automated. Um, but then the highle strategic work is actually going to be more valuable, right? And so this is really good for a young lawyer. It's fantastic because you know most people they go to law school and their goal is not to sit in a data room or do discovery or doc labeling for 10 years and then maybe go to trial once. That's not what they want to do, right? They want to give advice to clients. That is why you want to be a lawyer. They want to help people, right?

Uh they want to help people or they want to win. One or the other, right? But sometimes both. Um that's what you want to do. Uh, and it very much is it's like close to like professional athletes where they want to be the best at their craft and you get better by having more hands-on strategic experience than you do just sitting in a data room forever, right? And so I think that experience will be really good. Um, the law firm I think will be I think there will be a lot of transitions in how law firms operate. So I mean I now am a client like I use a lot of law firms, right?

And one thing that I've always got annoyed with is the really high billable hours for very low-end work, like looking at, you know, whether this change of control clause was triggered or not, things like that, right? Um, and as a customer, that can be annoying, right? But actually for the high-end strategy of like, how should I go about buying this business? How should I think about restructuring this part of my org? Things like that, I would pay more than you pay the best lawyers on Earth right now. like the delta between a junior associate and the best partner on earth is like three or 4x, right? Which actually doesn't make sense to me. Um, and so I think what will end up happening is there will be a lot of fixed fees for kind of the part of a transaction, the part of a litigation that is somewhat more commoditized, but the insight, the value, like strategic advice, looking around corners, things like that, I would argue you can charge more for that. I know part of your mission is related to providing better access to justice. What does that mean to you and and kind of draw the line for how Harvey helps us get there? So, a little bit of background info on this.

The average price of a lawyer in the United States is \$352 an hour. So, almost no one can afford a lawyer, right? Um, and I think there's a bunch of arguments about, you know, how much latent demand is there for lawyers, etc. The reality is there is a massive population of folks that do not have access to our justice system, right? Like either way you slice it, that is the case. Um, and even if you had every single lawyer work 20 hours a week on access to justice still wouldn't close that gap still literally in in without any anything else being fixed, right? Um and so I do think that we are going to have a large transition to lawyers using AI um to actually help and increase access to justice. Um there are a bunch of things on kind of like the regulatory uh side that prevent a lot of this but I think a lot of those will change. So the the best example of this is um there are kind of two conflicting rules here. It is unauthorized practice of law, right?

Um so businesses cannot give legal advice. someone who has not passed the bar actually cannot give legal advice etc. Um and the second one is you cannot make an equity investment into a law firm unless you are a lawyer, right? Uh and so that has basically cut out any sort of traditional financing in the legal sphere, right? Um Utah, Arizona, and some now some other states are also thinking about getting rid of these rules or creating sandboxes

to kind of experiment with that and it's been going really well. Um, I do think that there is going to be a lot of change in this in the next couple of years and I think it will be amazing. Like there there really is I think that if you can figure out the way to make sure that people are getting the same quality of legal advice with AI and maybe a lawyer in the loop depending on kind of how that that pans out. This will be incredible for folks.

And the my last piece on this is one of the things that I think people don't think about a lot is most people don't know when their rights have been violated. They do not know, right? And so like I I think like we take for granted that we have all been like, you know, in in this kind of like bubble have been super well educated. We know exactly what when something happens, we know whether we have legal recourse or not to some degree. A lot of folks don't, right? And so they don't know if a position a person in authority is doing something that is illegal or not.

Right? This is a huge problem in housing, huge problem um in kind of like collecting unpaid fees for things that shouldn't even be a fee, etc. Right? So I think that's an area where you can do a lot of not just providing the services but a lot of education um at scale that you couldn't do before. Can we zoom out to the AI market more broadly? Yeah. OpenAI was lucky enough to partner with you in the early days.

Um what what do you make of the recent developments that are happening in the AI ecosystem and what do you think are the developments that are most interesting for you? Yeah. Um I mean costs going down are always good. Um I I think like one one thing to think about here is there are so many use cases that we have gotten to work 70% of the time, right? Um, and the Oer models have been incredibly helpful for us because not even just the O series models themselves, but it unlocked kind of like our product strategy for the next 6 months to a year, right? Where there were so many systems that I don't think we thought that we would be able to build in the next year if it was just the GPT series models, right? And I think that the Oer has massively kind of changed that for us. And so our product roadmap has drastically changed because of that. What's an example of something that you couldn't do before that now you can do? So examples of this are like things that you need to do multi-step reasoning for and pulling from many sources at once. Right? So the thing that the O series models is really good at is orchestrating a plan and also executing on that plan. Right. So if you build a bunch of systems that are really good at basically extracting information from Edgar, then extracting information from case law, then extracting information from all your internal documents, the missing piece for us was what do you do once you extract all that information, right? Like once you have all of those different pieces, how do you combine them into the correct work product? That's what the O series models allow us to do. If Sonia had a magic wand and she were to wave her magic wand and you and Sam Albin all of a sudden switch places and you're now CEO of OpenAI and he's now CEO of Harvey. What would you do different at OpenAI and what might he do different at Harvey?

Yeah. Well, at Harvey he'd raise more money. He is pretty good at that. He's really good at that. Um maybe I'll start with what I think OpenAI has done an incredible job of and I would maybe even double down on it more. I think that they have done an incredible job of capturing the consumer zeitgeist. So I am not from Silicon Valley and I don't have I mean now I have tons of colleagues here and tons of friends but I didn't beforehand right and almost none of my friends know what any AI tool is other than chatbt that's it right.

Um, and I the thing that I think I would do differently and to be honest I think they are going this route anyway is just like put way more effort into productionizing that for consumers as much as possible. Right? So and and that's not just model performance at all. Right? So I think that still one of the biggest problems with chatbt and just kind of like all of these tools in general is they're looking so much on the performance from the model side and not at how do you make the experience easier for the user, right? How do you make it so that it extracts more information from the user? How do you help the user figure out what it can do and what it can't do? How do you help the user figure out how it combines with different pieces of information, etc.

So, I think on on this side, it would just be putting more and more effort into understanding the consumer behaviors and how they use AI right now and making that easier for them. What you guys have been partner with OpenAI from the very beginning. what has it been like and how has it evolved over time as it's gone from kind of undiscovered to being you know the center of the universe in many ways? Yeah, I mean so I think the thing that we have kept up that has been really awesome for the engineers um at at Harvey but but also at OpenAI I think is we have always been working on things that I think a lot of companies aren't.

And so what I mean by this is OpenAI would give us models all the time and say it performs way better on all of their benchmarks and we would respond and say sorry but it doesn't on ours like it's actually not better for us. Um and I do think that the reason our relationship has been so strong is we keep saying okay here's a model that can do XYZ. Uh I would like it to do XYZ and the rest of the alphabet too. Right? And that's what we're trying to do. Um, and so I think we've actually helped them a lot too with at least like applied use cases and how they think about post training and how they think about what are some of the things that companies like are really pushing to try to do, right? Um, and so I think that has been really really good for the relationship overall. Do you have any hot takes on Microsoft Open AI?

Um, yeah. Um, one of the hardest things about being an application layer company is you have to bet on model providers. Um, and it is not like we build our systems so that you could kind of just pull out one model and replace it for another, but you could do pieces of that, right? So like we don't just use one model for the entire system. You'll build a piece that's really good at this and then a piece that's really good at this and then you'll chain them all together, right? Um, but the reality is, you know, who's in the lead changes so often. Um, and different models are good at different things, right? And you have another problem too where all of your customers or most of them I said the partners probably aren't using a bunch of AI tools, but the associates for sure are and the associates are using all these different AI tools, right? And so what I'm trying to say is we I think have done a good job where we work with all the model providers and we are constantly testing out models with all of them. And I think that Microsoft, which has, you know, such a massive customer base, they have to figure that out too, right? They have to figure out, oh wow, we have customers that think that Claude is better at certain things.

What do we do about that, right? I mean, they can't really do anything about that, but they're like, you know, Mistral or whatever it is, right? And so I think that I think that that relationship is evolving as it naturally would. If the opportunity for Harvey is to revolutionize a trillion dollar industry and provide better access to justice, what is the threat? What are the biggest threats to Harvey? Um, I think not moving fast enough. I mean, I I I tell

this to my team a lot.

Um, and I think it's becoming very obvious in the past couple of months, too, that we're really living in a time when all of your timelines are compressed. I would argue this is in all of human history the most compressed timeline in terms of what you can change in the world, right? Um, and I think that you have to move so incredibly fast in order to keep up with that. And the speed is compounding in a couple different ways.

If you are constantly moving fast, you are required to constantly test everything around you. Pay attention to every single every single part of your industry. How your every single customer is using your product, how your subprocessors are changing their models, everything, right? And if you do not do that fast enough, you'll make a bunch of mistakes, I think. And you I mean you constantly will make mistakes, right, by moving very quickly. But the scariest mistake is you move too slowly and you miss a massive thing. Right? So there is part of a feature that someone releases something and you have steps 12 through 13 complete on that feature. But that 13th step you just cannot get the models to do it and somebody releases something that unlocks that you need to put that in the product immediately right because you need to start testing it. You need to see actually did it solve it etc. And if you aren't moving quickly and you just say, "Ah, that's something new.

We'll try it out eventually." I think it's a huge problem. Speaking of moving fast, good segue into our lightning round. Since starting this company two and a half, maybe three years ago, how many days have you taken off? It's a loaded question. Uh, I have not taken a day off. I probably should though to be to be in in all honesty. I think like I've definitely a little bit worn the you know in like full transparency I've worn the badge of honor of like you know don't take any time off obsession etc. And I think on on one side I actually very strongly believe in that. So going back to my point about the timelines being incredibly compressed you need to be obsessed right like you massively need to be obsessed. Um but I do also think the other side of that is like you need to transition how you are a leader like that needs to change. Um and you know our company last year we started the year with around 40 people. We have 260 right now and you just need to change how you are spending your time right um and I think that I've definitely learned how to spend my time differently. And there's also stuff I've held on to and I actually like deeply believe in. So, one of the things I I deeply believe in is I actually think you should do part for a little bit of every single job at your startup for a little bit. Like I actually I did I do it too much and I did it for too long. But it is all most of my hiring mistakes have been I didn't understand what the role did like at all. And I a bunch of people told me what it was, but you just it doesn't help you hire if you don't understand what the actual role is. Um, and so that one I I feel really strongly about, but at the same time, I also took way too long to hire. Um, and I probably I was probably doing too many low-level things for too long, right? Um, and so it's a combination of both of those. For those 260 people who work at Harvey today, what is the best thing about working at Harvey and what is the worst thing about working at Harvey? Yeah. Um I think the best thing is every day something new is happening right um from the product side from the GTM side uh you'll see you know a bunch of changes with the model providers that you want to integrate very quickly it is ex like it is stimulating above all else like for sure um and I think we're seeing like a lot of impact too and it it changes too and I think that's if you ask folks that have been here for a year or a year and a half and there aren't tons that have been for a year and a half. Um the thing that I think is most fun for them is how much our market has changed, right?

Where I mean we used to it it was brutal in the beginning for kind of how you do your sales process, the requirements people had and things like that. And now most of our customers like they partnering with us, right? And they're letting us do their onboarding. They're doing all of these things where it really has seemed like we're kind of this wave of all of us together. Doing it together. Yeah. Doing it together. And that that changed a lot. It there was a lot of push back in the beginning.

Like a massive amount of push and there's still some, but I think that has changed pretty drastically. Um the worst thing is the expectations are really high. Um, and you know, we had a really good year last year and I think we we had our offsite recently and I went up to the offsite and I basically said, "Hey, we did a really, you know, we had a great year last year. This year needs to be like significantly better and we need to raise the bar." Um, and you know, I think that's not always what people want to hear. Sometimes people are like, "Wow, we did such a great job last year. Now I take it easy." Yeah, let's take it easy. Let's just do the or do the same as we did last year. and then it'll be fine, right? And the reality is, again, going back to those timelines being so compressed, you can't do that. And I think that the the main thing I I ask people is basically, look, like I don't know what your goal is. I don't know if your goal is to make the, you know, a large industry change to learn as much as you can, make money, whatever it is, your options to do that in the next decade is the best it will ever be. like it it just is right. What you will make more impact than you ever will have the chance to in your life.

How have you changed as a CEO and what prior have you updated the most? Yeah. Um uh hopefully I've changed somewhat. Um I I think the so the the prior that I've updated the most is teach not do. Like I'm I'm bad at that. um because I I kind of want things done so quickly. I have a massive problem of whenever I start seeing friction, I just go, "Okay, I'm going to go do it." Right? And and that's actually really really bad. Um I think it it makes it so other folks can't learn, it makes it so that it is kind of like too top down, right? And it's something I'm working on a lot and I think I've gotten better at it hopefully. Uh you can ask my direct reports. I'm not sure if I have, but I I think I have. And I think that that's the area that I want to keep getting much better at is taking kind of, you know, slowing down in some instances and actually like setting oursel up to scale. Um, instead of everything being like go fix that, fix that, fix that, fix that. Um, yeah. Who's a better athlete, you or your co-founder Gabe?

Oh, Gabe is a much better athlete than I am. I mean, it's it's really unfair. So, he he played professional soccer. Um, and he is just a much much better athlete. But he had a knee injury recently and I will say I do a little bit like we live together and every morning like I'll get up to go to go to the gym and I definitely slam the door like a little bit loudly just so he knows that I'm going to the gym and he can't quite yet. I got it like we we met each other uh before this and we and we were we like became best friends and had no plans of doing a startup. Um, and he was just like always a better athlete than me. And so this is my like revenge a little bit. It's not gonna last long.

He's healing right now. And so he's probably right now working out and getting better. All right, last question. I'm going to steal the last question from Guy Ros. I don't know if you ever listen to his uh his show How I Built This, but he has the same last question every time, which is how much of your success has been luck and how much of your success has been skill? Um, it depends how you define luck. Um we have been in a place where we

have the options to apply skill and we have the leverage to apply skill and if you apply it correctly enough times to actually have a large impact from that.

And so the luck is the timing. The luck is do you actually have the option to make a difference to make an an impact. And actually I think you were we were kind of asking earlier about what have I learned as a as a CEO. One of the things that I I think I've actually tripled down on, quadrupled down on is young talent, like by far. Um, and that goes back to giving them the luck or the opportunity to actually try something they've never done before. It works out really well and like they don't get it right every time and but I don't either, right? But you adjust really fast and their ability to adjust I think is better than a lot of folks that have been doing this for a long time. And so maybe the the best way to phrase that is their skill is their ability to adjust to luck and seize luck opportunities more than anything else. I like it Winston. Thank you so much. Yeah, thank you.

[Music] [Music]