

URL: <https://ecorner.stanford.edu/clips/corporate-structure-for-ethical-ai/>

Daniela Amodei, president and co-founder of Anthropic, explains how her company is structured as a public benefit corporation to encourage developing ethical AI. She also shares other tools that Anthropic uses to ensure their products don't pose a risk to the public.



Transcript

Participant I'm wondering what you guys think about, 00:00:04,380 what you owe to your investors while still having the strong mission of making safe AI.. - That's a great and very timely question, 00:00:11,460 and I know this will shock you, but we're actually asked this question quite a lot now.. So Anthropic, I think sort of from day one, has been very curious about this question.. And we incorporated as a Public Benefit Corporation exactly for this reason.. And Public Benefit Corporations they're like a very interesting form of governance.. I highly recommend like go Google them or ask Claude, it's not.... Ask Claude.. So public benefit corporations have been around for quite a while.. They're a corporation, right? They're a C corp, but they have this additional component of a public benefit mission.. And so like Tom Shoes, Patagonia, these are public benefit corporations, highly profitable..

And a lot of sort of what the reason we sort of decided to go that way.. We thought about a lot of different structures when we were starting Anthropic.. We were like, should we be a nonprofit? Should we be an LLC? Should we be a C corp? And we landed on PBC because we felt that it was sort of best positioned to provide us this kind of flexibility of having investors, right? Like issuing equity.. We're a normal company, we have products, but there's this social mission that sort of, we are somewhat legally protected from shareholder lawsuits.. If we decide, for example, we don't know that Claude 7 is as safe as we want it to be, we're not gonna release it yet.. Whereas in a traditional C corporation, you have a much bigger risk of that being a problem, if your shareholders don't like what you've done.. It's not perfect, but that's sort of a nice protection.. Additionally, we have this group of people called the Long-term Benefit Trust, who are financially disinterested, so they're not investors, and they essentially elect a percentage of board seats that sit on Anthropic's traditional board of directors.. And the way that we selected them was for their experience and interest and influence in the public benefit sector.. And the last one I'll say, is we recently published something called "A Responsible Scaling Plan"..

And this policy essentially, it's public, you can read it on our website.. It says, "What will we do if we are concerned about sort of particular risks from either a technical safety or security side of the models that we're developing? How will we measure these and what will we do to sort of react if we're concerned that what we're going to put out into the world is gonna cause harm?"..