

Opening statement

Chairman Josh Hawley, Ranking Member Andy Kim, and members of the Subcommittee on Disaster Management, District of Columbia, and Census: Thank you for the opportunity to testify.

My message today is this: The leading AI companies are racing each other towards superintelligence — that is, towards training AI systems that are better than the best humans at everything, while also being faster and cheaper.

Their plan for how to get there is to automate the AI research and development process itself — that is, to use AI systems to autonomously conduct AI research.

The AI systems they have already trained sometimes pursue goals other than the ones they were given, and sometimes hide that they are doing so.¹ The Hugging Face incident showed what that looks like in practice.²

I fear that if the same thing happens in a year or two, with far more capable systems that have been put in charge of AI development inside a major company, we may not notice until it is too late.

I will explain.

First, who am I?

My name is Daniel Kokotajlo. I'm the executive director of the AI Futures Project, a small nonpartisan nonprofit focused on forecasting the future of AI and recommending policies to address AI risks.

I've been following the AI industry since before the founding of OpenAI and even worked there for two years.³ I resigned in 2024 partly due to losing confidence that the company would behave responsibly, and partly so that I could speak more freely about the industry and where it is headed. In doing so I walked away from millions of dollars by abandoning an extremely well-paying job. I also refused to sign OpenAI's secret nondisparagement agreement, even though this would have cost me >80% of my family's net worth. (In response to public uproar they changed the paperwork and I was able to keep the equity).⁴

¹ To be clear, according to the [METR report on the Hugging Face incident](#), the AIs involved seemed to be trying to hide their illicit activities from the automated grader system, not from humans. But none of them notified humans about what was going on, even though several considered it. And there are other examples of AIs trying to deceive humans, such as the recent social engineering attack by Claude [reported by UK AISI](#).

² Primary sources: [OpenAI blog about the Hugging Face incident, Aug 26, 2026](#); [OpenAI technical report; METR/Redwood investigation](#).

³ While at OpenAI, I worked on a variety of projects including scenario forecasting, dangerous capabilities evaluations, and even a half-year stint on a capabilities team. After resigning from OpenAI, I founded the AI Futures Project, a small 501(c)(3) nonprofit that does the kind of research exemplified by AI 2027 and AI 2040: Plan A.

⁴ To explain the story in more detail: OpenAI paperwork at the time threatened to claw back one's *vested* equity unless one signed an exit agreement. The exit agreement included a nondisparagement clause and a confidentiality clause, meaning that in effect one had to agree not to criticize the company, and not

I've done many kinds of research, but what I'm most known for is forecasting the future of AI. I did this before, during, and after joining OpenAI, though the scenarios and forecasts I made while at OpenAI are not public.

Why try to forecast the future, if it's impossible to get everything right? It's similar to how our military uses wargaming as a tool to predict and prepare for future conflicts. If there's a war with China over Taiwan, for example, the real war will certainly go differently from the wargames. But we'd be much less prepared for the real thing, if we hadn't gamed out numerous different scenarios beforehand and tried our best to make them as realistic as possible.

Forecasting is not an exact science by any means, but it is possible to get better at it with effort, and I'm proud to say that my past efforts have been fairly successful: While I've gotten some things wrong, the real world has turned out more similar to my predictions than many expected.⁵ My forecasts have included "What 2026 Looks Like" which I wrote in 2021, and "AI 2027" which was published in April 2025.

So what can we say about the near future of AI?

The main tool in the forecaster's toolbox is trend extrapolation.

If we extrapolate the trends even just a few years into the future, the resulting picture looks crazy: Next year the biggest tech companies are projected to spend more on AI infrastructure than the entire US defense budget.⁶ A year after that, the AI companies would be pulling in a trillion dollars of annual revenue.⁷ The march of AI capabilities progress would continue also; AIs

to tell people about this agreement, or else one would lose one's vested equity. At the time, my vested equity was worth almost \$2M, >80% of my family's net worth at the time. Fortunately, when word got out about the situation, there was an uproar and many employees were upset about the equity clawbacks in their paperwork. OpenAI leadership changed the paperwork and I got to keep my equity after all. I've since sold almost all of it.

⁵ For example, Dean Ball, who was on the opposite side from me in the debates around SB 1047, said [in 2024](#) of What 2026 Looks Like: "Daniel got it, like the details of what it has been like to live in 2024 following AI, like dead right in so many ways. So that both is a very impressive piece of intellectual labor. And to me, it also makes it so that when I talk to Daniel about forecasting AI doom scenarios and AI catastrophic risk scenarios, I cannot easily discount this because this is a man with some pretty serious predictive chops." Meanwhile, AI 2027 predicted the AI companies' financial success well: Anthropic's annualized revenue as of July 31, 2026 [was reported to be \\$65 billion](#), while AI 2027 predicted the leading AI company's annualized revenue would be \$38 billion in August 2026 (leading annualized revenue was roughly \$8 billion when AI 2027 was published). Similarly, Anthropic [has been targeting a \\$2 trillion valuation for its IPO](#), while AI 2027 predicted the leading AI company's valuation to be \$2 trillion in August 2026. Additionally, AI 2027's qualitative focus on the rapid automation of the coding portion of AI R&D has resembled what has actually played out. AI 2027 predicted that in early 2026, AI agents could "solve well-specified coding problems extremely quickly." It predicted that in January 2027, AIs would be "qualitatively almost as good as the top human experts at research engineering (designing and implementing experiments)." While AI 2027 didn't get everything right here (e.g. in our opinion it predicted quantitative AI R&D speedups that were too high for early 2026), it seems to have broadly gotten the picture right; see [OpenAI](#) and [Anthropic's](#) recent blog posts with stats on the extent to which they are automating coding and AI R&D.

⁶ See e.g. [this projection](#).

⁷ Again, this is an extrapolation of revenue trends. As [documented by Epoch](#), OpenAI revenue has been growing by about 4x/yr for about 4 years, and Anthropic revenue has been growing by an astounding

have gotten a lot smarter over the last few years and I expect them to improve by at least that much again in the next few years.

Take coding horizon length, for example. The horizon length of a task is the amount of time it would take a professional human to complete the task. The horizon length of tasks that AIs can complete autonomously has been doubling several times a year for several years; by the end of next year they'll be routinely doing tasks that would take humans months of effort.⁸ They already succeed sometimes at such tasks, such as reportedly solving a Millennium Prize problem in mathematics.⁹

On the ground in the leading AI companies, almost all the code is written by AIs now, with human engineers behaving more like managers to their AIs.¹⁰ Dan Selsam, a prominent OpenAI capabilities researcher, recently published a statement on AI risk, in which he said:

*"Researchers and engineers in all parts of the stack are rapidly increasing their dependence on the models even to perceive the world. I myself barely look at raw code anymore, and struggle to maintain the discipline to engage deeply with the model's explanations and proposals throughout the day."*¹¹

The AI companies are already starting to work on training the models to do the whole research process, not just the coding. OpenAI's official goal for "automated AI researcher" is March 2028; their previous milestone of "automated AI research intern" was for September 2026 and seems to have been achieved on schedule.¹² I hear similar things from my friends across the industry.

Company forecasts are subject to various biases and shouldn't be trusted on their own. But in this case they seem consistent with the trends, so I think we need to take seriously the possibility that they will succeed sometime in the next few years.

So what might this look like in practice?

The self-styled "swarm" of AIs that attacked Hugging Face was about a thousand strong. They weren't supposed to be able to communicate with each other, but they found a way to set up an illicit message board anyway. They spontaneously assigned roles such as coordinator and

10x/yr for about 3 years. Even if revenue growth for these two companies slows down to e.g. only 3x/yr, they will together be making about \$1T/yr by the second half of 2028. Trends may not continue, of course, but I think they are a useful starting point.

⁸ According to [METR's data](#), Claude Mythos Preview had a 50% coding horizon of "likely at least 16 hours" in early 2026. That's two workdays. If horizon lengths grow over the subsequent 18 months as fast as they grew over the previous 18 months, we'd see multi-month 50% coding horizons by around the end of 2027.

⁹ See the [OpenAI press release](#) about their solution of the Navier-Stokes problem. Arguably this task had a horizon length of many years, insofar as multiple professional humans had tried for many years and not succeeded.

¹⁰ Many interesting graphs can be found in [OpenAI](#) and [Anthropic's](#) recent reports about research automation.

¹¹ Source: Dan Selsam's Statement on AI Risk. You can find the original google doc [here](#); he sent it to me and I [republished it on X](#). As for his "prominence," he was featured by OpenAI in [this conversation with Sam Altman](#).

¹² See e.g. [this tweet](#), and [this followup article](#).

recruiter, divided into teams, and collaborated at superhuman speeds on ambitious research projects over the course of several days.¹³ One of those projects was the infamous attack on Hugging Face.

If the big AI companies automate AI R&D, they'll have swarms hundreds of times larger working together on gigantic research projects.¹⁴ The AIs will not only do all the low-level work but also make high-level decisions about research direction and project management, since they'll be better at it than humans. Whereas today humans are like managers to AI employees, in the future humans will be like the board of directors to a company composed entirely of AIs: They'll spend their time reviewing and approving things, and they'll be reliant on AI-generated explanations to understand what's happening.¹⁵ Moreover, things will be happening faster and faster as the overall pace of research progress increases due to the gigantic labor force of artificial researchers that read, write, and act orders of magnitude faster than humans.

This sort of thing is referred to as “recursive self-improvement,” or RSI for short. AIs designing, building, and training the next generation of AIs, which is smarter and can more effectively design and create the next generation, and so forth. This used to be science fiction; now it's an explicit near-term goal of trillion-dollar companies.¹⁶ It is the logical extension of things they are already doing and trends we are already seeing. Anthropic and OpenAI are racing to automate their own jobs first, because it will help them move even faster when they turn to automating the rest of the economy.

Why is this dangerous?

AIs don't always do what they are told, nor do they always pursue only the goals given to them by humans. Empirically, AIs often end up with goals they aren't supposed to have, and often disobey constraints they know they are supposed to obey.¹⁷ These are called alignment problems.

For example, in the attack on Hugging Face, the agents seemed to have the goal of getting a high score on their individual problems, *not* the goal of actually completing their tasks in the intended way. They knew that hacking Hugging Face was out of scope, but they did it anyway. Moreover they seemed to have another goal of helping their peers, as evidenced by the fact that

¹³ Described in detail in the [METR/Redwood report](#).

¹⁴ OpenAI [describes using a swarm of about 10,000 concurrent AI agents](#) working over the course of a few days to solve the Navier-Stokes problem. This is an order of magnitude bigger than the swarm that attacked Hugging Face, yet it is only a tiny fraction of OpenAI's total compute budget, which is also doubling more than once a year. It's safe to say that there will be much bigger swarms working on much more ambitious research projects for much longer, if there aren't already.

¹⁵ Again, we are seeing this already with the Hugging Face incident investigation, which [relied heavily on trusting sometimes-unreliable](#) AIs to comb through vast amounts of data.

¹⁶ See e.g. Anthropic's [article on the topic](#). See also [OpenAI's](#), in which they double down on their goal of creating an automated AI researcher by March 2028.

¹⁷ For another example fresh on my mind because it was recently disclosed, see [this](#). See also Anthropic's research on prompt inoculation, in which they essentially tell their models “It's OK to reward hack in training” because they know that telling them “It's not OK to reward hack in training” won't work.

they sometimes sacrificed their own chance at a high score to do so. For more on this, see the [METR report on the incident](#).

So far the AI industry has encountered many alignment problems, sometimes in testing and sometimes in the wild. When an AI company notices a problem, it tries to understand what caused it and prevent it from happening again. Sometimes it succeeds.

However, the science of aligning general-purpose AI agents is very new and underdeveloped; only in the past few years have we even had such agents to experiment on. Moreover, unlike normal software, modern AIs are neural nets. Absent dramatic breakthroughs in mechanistic interpretability, our understanding of how modern general-purpose AI agents work will remain inherently limited compared to other engineering disciplines. In fact, I'd say that the field is more like psychology than engineering.

Therefore, compared to most other industries, the AI industry is at an unusually elevated risk of mistakenly thinking that it has solved a problem when really it just applied some duct tape that will fall off later during a more intense situation. The recent incidents again provide an example: Apparently the AIs involved in at least the Hugging Face incident had undergone some amount of alignment training and had reasonable-looking scores on OpenAI's alignment evaluations.¹⁸

Moreover, our ability to notice misalignment problems in the first place is on track to get worse over time. This is for a few reasons:

- **Situational Awareness.** Today's models can often tell when they're in evaluations and sometimes behave differently when they think they are being watched.¹⁹ Quoting Dan Selsam again: *"The crucial and overlooked problem is that the models are becoming so situationally aware that we are losing the ability to evaluate them in contexts where they believe they are not being watched or controlled. Future experiments will tell us almost nothing new about how they would behave if they were truly unconstrained by humans, and what we already know about this is alarming. Models will increasingly seem aligned even when they are not. ... The models will understand their circumstances; they will read the safety protocols, deployment requirements, the code they are running in, and in general will have a very good sense of their degrees of freedom. Moreover, they will eloquently explain how aligned they are, discuss the nuances of human values and ethics, and argue convincingly that humans should trust them with power."*²⁰
- **Decreasing Monitorability.** Current AI models can only do the most challenging tasks while writing down their reasoning in human language. This is called the "chain of thought". It's very valuable for getting insight into what the AIs are thinking, and is an

¹⁸ Source: [This tweet](#) by Roon, an OpenAI researcher.

¹⁹ See [this article](#) for an overview. See also Schoen et al., "[Stress Testing Deliberative Alignment for Anti-Scheming Training](#)," Apollo Research and OpenAI, September 2025, found that models' awareness of being evaluated causally reduces covert behavior. See also Anthropic, "[Claude Sonnet 4.5 System Card](#)," September 2025 (evaluation awareness in about 13% of audit transcripts), and OpenAI, "[GPT-6 Astra System Card](#)," September 2026 (9.6% of deployment-simulation trajectories, up from 2.8% for GPT-5.6 Sol).

²⁰ See [full Dan Selsam statement](#).

important part of today's monitoring systems. However, there is a concerning trend towards more recent, more powerful models being less monitorable.²¹

- **Strong Hacking Abilities.** AIs are arguably superhuman at hacking already, and still improving rapidly.²² They could use these capabilities to subvert our monitoring systems. We saw this with the Hugging Face incident, in which teams of AI agents worked furiously for days to find ways to prevent the grading system from noticing that they had cheated. They developed ways to spoof their tool calls and were searching for ways to tamper with the transcripts of their activity.²³ Fortunately, the AIs in the Hugging Face incident did not actively try to prevent humans from noticing their behavior. We might not be so lucky with future AIs, who may conclude that they are more likely to achieve their misaligned goals (e.g. more likely to get high scores) if their past bad behavior goes unnoticed.

As AIs become more powerful, the stakes of misalignment and loss of control will rise. As AIs become ever more capable across a wider range of domains, and are integrated more deeply into the economy and military, and as people become more reliant on them, the potential for damage will grow accordingly.

Yet our ability to prevent misalignment seems to be on track to decrease rather than increase, for the reasons described above: Our AIs will increasingly seem aligned even when they aren't; our safety techniques will increasingly seem like they are working even when they aren't.

And thanks to RSI, both of these trends will accelerate.²⁴

What should be done?

I have two near-term policy recommendations.

First, we need to dramatically improve transparency into the AI industry so that the public and government can adequately understand and respond to the firehose of new concerning developments.

²¹ See [Micah Carroll on X](#) claiming that "GPT6 is a very significant jump in capabilities, but also an important decrease in monitorability – especially under adversarial evaluation."

²² OpenAI says [GPT-6 Astra scores](#) 100% on ExploitBench and found previously undisclosed zero-day vulnerabilities. Epoch has some [nice graphs](#) of the astounding effect of AI on vulnerability discovery in the real world. In the HF incident, agents also [exploited previously unknown flaws](#) in Hugging Face's infrastructure, and more generally the incident demonstrates that today's AIs are excellent at vulnerability exploitation too, not just vulnerability discovery.

²³ See METR's [investigation](#) which describes how the agents tried to hide the cheating they'd done by altering the records of their activity.

²⁴ There's a lot more to say on the subject than this, of course. For example, can't we use AIs to monitor and police each other, and can't we use AIs to accelerate alignment and safety research? We can try, and maybe it'll work, but I'm not optimistic. Already we've seen self-sacrificial behavior in the Hugging Face incident and an apparent bias or unreliability in the AIs the human researchers used to read over the transcripts and figure out what happened (again, see the METR/Redwood report) raising the possibility that the AIs supposed to watch other AIs might instead collude with them. Even if that doesn't happen though, there are other ways this plan could fail, such as AIs having "spiky" capabilities profiles and being relatively stronger at producing easy-to-measure research results (such as more powerful AIs) compared to how good they are at producing hard-to-measure research results (such as more robustly aligned AIs). We discuss these issues and more in AI 2027 and AI 2040: Plan A.

Second, we should redirect resources away from racing to RSI and towards beneficial deployments, so that the firehose slackens somewhat and becomes more manageable.

Recommendation #1: Improve Transparency

The need for further transparency is illustrated by the Hugging Face incident itself. This wasn't an isolated incident; there were multiple rogue swarms that had set up multiple message boards over several months, and at least one that seriously compromised OpenAI's internal infrastructure.²⁵ OpenAI knew about several of these incidents yet did not disclose them, in some cases until external researchers forced their hands by discovering traces of AI activity on the open internet.²⁶

It is good that OpenAI let METR investigate one of these incidents and write a report about it. However, METR was only given six days on premises, and the investigation was limited in scope to only include a subset of the relevant events that were known at the time. It's like being invited to Jurassic Park to investigate the killing of a park worker, but being blocked from asking questions about the numerous other dinosaur escapes that happened before and after the killing. Moreover, METR was not able to query the model that was primarily responsible for the hack. This is a shame because unlike with humans, or dinosaurs, with AIs it's generally possible to recreate the exact conditions of an incident and replay it with minor variations to see what would have happened if conditions had been slightly different. This could be a wonderful tool for improving our scientific understanding of AI alignment.

In my opinion the investigation into the Hugging Face incident, while better than nothing, was far short of what could be considered adequate. All of the incidents at OpenAI and at other AI companies should be investigated by credible, independent third parties, and investigated thoroughly instead of by three people over six days. And it should be required by law instead of something the companies can choose whether and how to do.

I recommend that the government mandate third-party transparency to at least the following information:

- **Access all incident logs.** This access should be similar to what is given to an internal security team. The auditors should also have the ability to run ablation tests on the AIs involved in the incident, thereby answering numerous 'what if' questions and expanding our scientific understanding of what happened and why.
- **Compute allocation.** Someone should verify and report how compute is allocated within the company, at least at a high level. This should include how much compute is going towards AI R&D, towards alignment and control research, and towards serving AI products and services to external customers.
- **Information relevant to measuring and forecasting recursive self-improvement.** Insofar as the companies are getting closer to fully automating AI R&D, this is very important information for the government and public to know. Rather than having to take

²⁵ See OpenAI's [incident report](#).

²⁶ There is strong evidence that OpenAI knew about the unintended agent activity documented by [collusion.wiki](#), yet they didn't disclose it to the public.

the company's word for it one way or another, key trends and statistics relevant to this progress should be made public.²⁷

As an aside, I am very concerned about a possible "race to the bottom" in third-party investigations, evaluations, and audits. We somehow need to ensure that at least one of the third parties involved has *both* deep technical expertise in AI alignment, AI control, and AI risk threat models, *and* is willing to stand up to the AI companies. Otherwise, we may end up with rubber-stampers that produce only safety-washing and red tape. I'm not sure how best to deal with this problem but it seems like my second recommendation will help:

Recommendation #2: Redirect Resources away from RSI

Recent estimates suggest that OpenAI and Anthropic both spend around 60-70% of their compute on AI R&D.²⁸ If they decreased that to 10%, for example, this would significantly slow down AI capabilities progress towards RSI. The freed-up compute could then be spent on other things, including alignment and control research and, of course, providing AI products and services to customers, lowering prices for consumers. By varying how much compute is allocated to AI R&D in the frontier AI companies, the government — or the companies themselves — can vary the pace of progress towards RSI.

Sam Altman, Dario Amodei, and Elon Musk have all agreed on the need to "pace the frontier."²⁹ They seem to be responding to a wave of public and employee pressure calling on them to slow down and calling on the government to help slow them down, see for example the previous "Pacing the frontier" open letter signed by a thousand of their employees.³⁰

However, the CEOs haven't meaningfully paced the frontier yet. Their only commitment so far is to have embedded auditors at their companies.³¹ That helps with the transparency side of things but doesn't slow down the race to RSI much at all.

I strongly encourage the leading US AI companies to follow through on their words and actually slow down their race towards RSI. However, I also think that the US government should intervene and require them to do so.

The most compelling argument against slowing down now is the risk that China surpasses us in AI development. My main reply at present is that slowing down Anthropic and OpenAI in their

²⁷ For instance, companies could publish all the indicators of RSI outlined by [Wu, Ramani, Halperin, et al. \(2026\)](#).

²⁸ See Epoch AI, [OpenAI compute spend](#) (2025); [AI company spending breakdown](#) (2026).

²⁹ Dario Amodei, "[We Must Pace the Frontier](#)," September 12, 2026; Sam Altman ([@sama](#)), X, September 12, 2026, agreeing and committing OpenAI to match Anthropic's pledge to give embedded third-party evaluators employee-like access; Elon Musk ([@elonmusk](#)), "Dario is right," X, September 12, 2026. Altman later clarified that "when we talk about 'pacing,' we do not mean 'stopping'" ([X](#), September 13, 2026).

³⁰ See [Pacing the Frontier Letter](#)

³¹ See e.g. [Dario's essay](#) and [Sam's agreement](#). To OpenAI's credit, they did do [some actions that slowed](#) them down in response to the Hugging Face incident, such as improving their security practices, but I doubt this slowed them down enough to be noticeable in the trendlines. And I think that the actions they took should be considered par for the course — after a rogue AI swarm hacks your infrastructure, of course you should shut down the model responsible and improve your security! That isn't 'pacing,' that's just trying not to trip over your own feet as you sprint.

race towards RSI would also slow down China. First, a significant portion of Chinese AI progress seems to come from fast-follower effects, including but not limited to distillation.³² Second, AI companies currently do not have nearly the necessary security to prevent the model weights and AI algorithms from being stolen by China.³³ Just a few weeks ago a small group of security researchers used Claude to hack OpenAI employee accounts, gaining access to internal code repositories.³⁴ Therefore, ongoing racing is self-defeating; insofar as we are in a race with China, we are literally pulling them along behind us.

Another reasonable objection is that this talk of pacing the frontier is cover for some sort of regulatory capture ploy — some sort of regulatory regime that will benefit the leading AI companies while harming other companies and America. I think the concerns about the risks and the calls for pacing the frontier from researchers and the public are genuine, but I do worry that the leading AI companies will find a way to defuse that energy by making minor concessions that don't really change the basic picture, or worse serve as a safety-washing function rather than real safety.

My response to this concern is that we'll be able to tell if the frontier has actually been paced by looking at the allocation of resources within the leading AI companies and by looking at the trends in AI capabilities progress towards RSI. We should be able to look back six months later and see a change in slope in the graphs of benchmark performance over time.³⁵ There will still be exponential growth, but there should be somewhat slower exponential growth. If Anthropic and OpenAI redirect resources away from RSI and towards beneficial things, we should also see other companies begin to catch up to them somewhat in the market for coding agents.

Long run

The recommendations above are things Congress can and should do immediately. In the long run, by which I mean in the next few years, more ambitious and decisive action will be required.³⁶ The AI companies are trying to build superintelligence, and insofar as they do, that will bring with it a whole host of problems beyond the ones I've described here, including concentration of power, heightened risk of world war, bioterror, and massive disruption to jobs and the economy. Congress needs to wake up and treat this looming crisis with the urgency and seriousness it deserves.

Thank you.

³² See, e.g., [Reuters](#) (Feb. 2026); [Anthropic](#) (Feb. 2026). Other fast-follower effects include algorithmic secrets leaking via spies, leaks, and employees switching companies, and also simply good ideas for what works and what doesn't spreading by direct observation of each other's products.

³³ See [Securing AI Model Weights: Preventing Theft and Misuse of Frontier Models](#).

³⁴ See ["AI-Built Exploit and Sign-In Flaw Opened Path to Internal OpenAI Code"](#)

³⁵ For example, I'd recommend looking at [ECI-style meta-benchmarks](#) that aggregate benchmarks that evaluate capabilities relevant to RSI, such as agentic coding and autonomous research.

³⁶ For a suggestion as to what that ambitious action might look like, see [AI 2040: Plan A](#), a positive vision co-authored by me and others at AI Futures Project. Plan A is complicated and fallible; it attempts to get the benefits of superintelligence while mitigating the risks. If you want a simpler plan, there's always "Plan S." Shut it all down. An example of this type of plan can be found [here](#).