PROJECT IMAGO DEI

PROJECT IMAGO DEI

Strategic Communications Document

V2

2026

"I have set thee a watchman… therefore hear the word at my mouth, and give them warning from me" (Ezekiel 33:7).

Art by Eysteinn Thordarson

Table of Contents

1. Summary

The purpose of this document is to organize our findings on how to communicate the essential and important messages that help promote AI safety to audiences that may be resistant to the traditional messaging. In particular, this is written to communicate with audiences typically indicated as conservative or right-of-center, with a further focus around religious audiences.

Such audiences tend to be skeptical of messaging involving existential risk and more involved arguments in regard to "bad" incentives. So, effective messaging to them should focus on storytelling to establish common ground to memetic landmarks and threats they are familiar with.

In essence, for audiences skeptical of AI or of AI risk, the strongest opening move is not to emphasize on existential or extinction risks.

Instead, two major messages can resonate:

- That the people building this technology
- openly say things that involve humans being in terrible futures
- And even if the destination is somewhere we want to be, rushing the speed of its development can have terrible consequences.

The rest of this document explores the implications of this.

2. Audience Segmentation

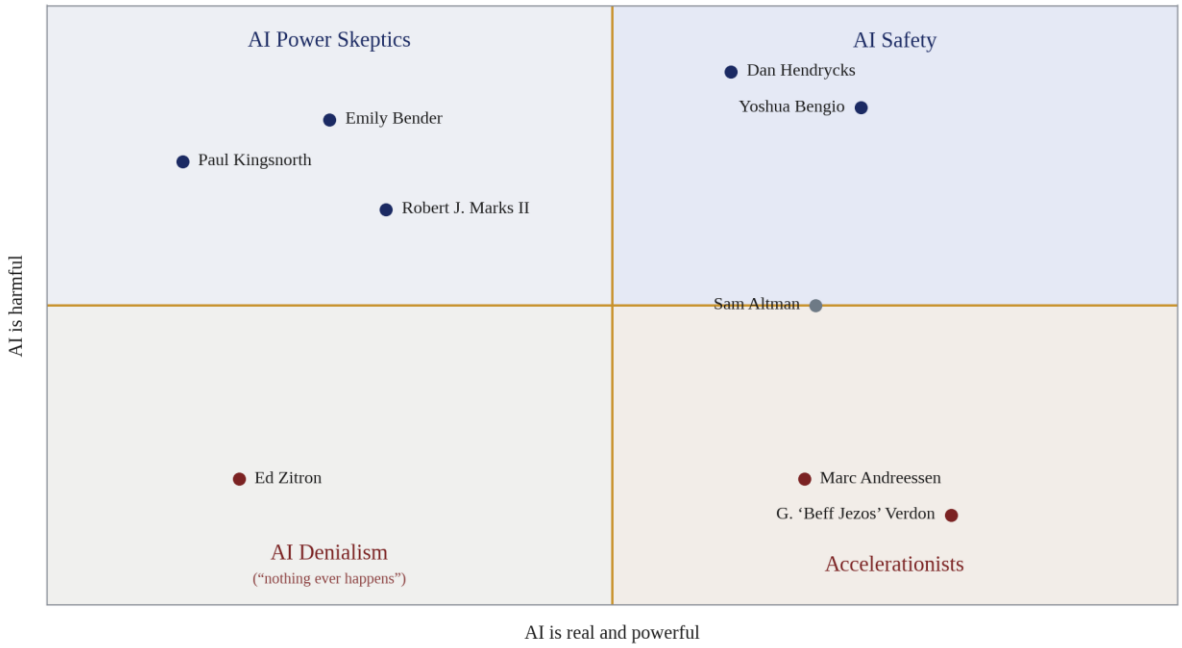

Effective communication depends not only on the messaging but who is being messaged to; it is important to understand the specific concerns and beliefs of the audience in question in order to effect successful communication. In particular, two aspects stand out: to establish common landmarks of identity ("the people speaking to me atr the same as me") as well as to understand their model of the world ("the real reason why things are happening is because…").

AI Safety *("AI is real and powerful, and AI is dangerous.")*

These are audiences who understand what AI can do and are already concerned about what it means; many of them hold that AI is effectively a full replacement of human capabilities. A typical example of this kind of concern might be seen in something like this: Natural Selection Favors AI.

Casually they fall under "AI Safety."

Includes: These are generally individuals who are involved with technology themselves or have a background such that they sympathize with the power of technology and believe that it is both possible and important to regulate such changes. They tend to skew younger and male, more agnostic/atheistic and with higher education.

Named examples: Dan Hendrycks, Yoshua Bengio.

Approach: Given that this is a communication guide for AI safety communicators, we should expect that our common landmarks are already there; when in doubt, emphasize that we acknowledge the real catastrophic risks are real. Their model of the world accepts that catastrophe and disempowerment of humanity as being very real, similar disasters have happened before, and that we need to work to prevent this.

Sensitivities: Their particular sensitivities involve questions if the speaker is insufficiently engaged with the existential risks of AI; e.g. concern with copyright or even data centers can be seen as being insufficiently engaged with catastrophic risk. A secondary sensitivity is if their interlocutors are motivated by industry in any way.

Q2 AI Power Skeptics *("AI is not powerful, but it can be used to make bad decisions quickly and worsen the world"_*

These are people who are skeptical of AI as being a powerful force, but believe that AI nonetheless may be making the world worse by disturbing importance ethical principles they hold dear. For example, many academics who follow the idea of AI is just a "stochastic parrot" which limits the scope of its power, but nonetheless believe that it can be misused to lead to reduced critical thinking by children, be used against traditionally disempowered groups(e.g. predictive policing) or worsen working conditions via monitoring. An example of such thinking is the On the Dangers of Stochastic Parrots paper.

Often, they are associated with "AI Ethics."

Includes: this has a broad swath from the liberal "AI ethics" individuals including Emily Bender herself who provided the "stochastic parrot" quote as well as some in conservative circles who focus upon child safety or copyright. They tend to skeew more female, liberal, and with education more focused on humanities and also include some older computer scientists.

Named examples: Emily Bender; Paul Kingsnorth, Robert J. Marks II.

Approach: We should expect that communication with them will be more difficult on anything that focuses upon emphasizing on the capabilities of the technology to be existential; therefore, we will need to both demonstrate at least in a limited manner, on the capabilities of the technology as well as sympathize with the real risks they have identified.

The main common landmarks will be to emphasize on common concern with the technology, and how it is important to get to better worlds whether even bad outcomes are avoided via regulation, even if they do not believe in catastrophic risks.

Their worldview is that the risks come largely from say, environmental contamination by data centers("boiling the oceans"), by the violation of artistic rights by AI corporations, or implementation of weaponry or policing or other decisions systems via systematic bias.

As much of this concern is at least directionally correct, goals will involve setting coordination to create rulesets to prevent bad outcomes.

Sensitivities: Many such individuals specifically take badly to any arguments of existential risks as a sign that, in their mind, the interlocuter is not serious enough and believes in "sci-fi threats."

Q3 Accelerationists *("AI is powerful, AI is inevitable, AI is good.")*

These are individuals who understand exactly what AI can do and proceed to dismiss the risk, either for libertarian reasons, sometimes for economic interest, ideology or tribal alignment. An example of this thinking is the Techno-Optimist Manifesto. **At its most**

extreme, this includes blatantly anti-human factions such as human successionist ideas and "e/acc" (effective accelerationists).

Includes: This includes much of the a16z-adjacent tech-right and others, many e/acc voices; and often leans heavily male and business focused, and often circles around worship of inevitability or general beliefs that regulation leads to broadly negative outcomes. For our purposes, the main considerations is that they tend to lean into varieties of nihilism, libertarianism, and sometimes, explicitly anti-human beliefs. A surprisingly common tagline is "AI will not replace you, someone using AI will replace you."

Named examples: Marc Andreessen (a16z), Guillaume 'Beff Jezos' Verdon (founder of e/acc).

Approach: Such individuals are obviously not generally our audience and sometimes they constitute explicit bad actors. Nonetheless, it's important to understand how to communicate with them because there is often an aspect of the "peanut gallery" or others who are watching the exchange and public engagement can have the frame of viewers shifted even if the actors themselves are not moved.

Nonetheless, we should extend kindness to all, and truth is some of our most powerful arguments here. It's important to use real, quotable examples of loss of control and what few common landmarks include if we have any similar sympathies for how technology has generally been a strong good for humanity and also indicate where regulation has led to such technology to become flourishing rather than negative.

Their worldview generally is one where cooperation, collaboration and restraint are forms of weakness and they primarily see the world in the determinants of strength. Therefore, indicate where cooperation has led to more power, and actual losses by their camp and fellow travelers.

Sensitivities: A particular concern by accelerationists is on centralization of power and suspicion that regulation centralizes power against them. As noted above, their overwhelming concern is about power.

AI Denialism *("AI is fake, it is all a scam and we just need to either ignore it or adapt to it.")*

Audiences who neither believe in AI's power nor believe risks are real and thus casually may be called the "nothing ever happens" individuals. Such people tend to be dubious about how much technology has affected us or hold beliefs that AI is something that will just go away by itself, and often hold ideas that AI is a form of "hype." An example of this would be Have We Reached Peak AI.

This is generally a form of AI denialism.

Includes: This is surprisingly common, and arguably a very strong brand as a whole due to the generalized threat that AI holds to almost all human endeavor. One could argue that this is almost a kind of default perspective for many, and it appeals in particular because it holds out hope for a world that will stay stable without enormous effort. Quite a few creatives also associate with this.

Named examples: Ed Zitron

Approach: The major landmarks of common identification here is to indicate parts of genuine human uniqueness, such as the soul and on effort and history – all of these are aspects which we can and indeed do focus on, as uniquely human. Additionally, we would want to emphasize that there are important aspects of human uniqueness that ought to be prevented from replacement.

Their model of the world is one where nothing much about the technology is real and that it is largely fake, but is nonetheless making the world potentially worse via "enshittification" or secondary effects such as the destructive scanning of books.

Ultimately, some of these individuals might be movable when AI's real capabilities are demonstrated concretely but it has to be carefully done in a way that does not make it seem like we are arguing with them directly. What they can be more assuredly moved on is that various agreements on resource extraction, home displacement, etc are based on sympathies for humanity can thus can form common ground.

Sensitivities: AI denialists tend to reject any arguments on loss of control or more extensive catastrophic risks and surprisingly receptive to arguments of data centers as waste.

3. Keystone Messages

There is a pervasive anti-human sentiment among many of the people building this

A pervasive anti-human sentiment runs through the leadership of the frontier labs and their intellectual ecosystem. This has been stated in their own words, on the record. Many of the people building this technology have said that human beings may not be the ones who inherit the future, and that this is acceptable, or even desirable.

- **Nick Land:** philosopher popular in Silicon Valley: "Nothing human makes it out of the near future." (Meltdown, 1994)
- **Richard Sutton:** 2024 Turing Award winner: "The light of humanity and our understanding, our intelligence, our consciousness, if you will can go on without meat humans." (Jacobin, 2024)
- **Larry Page:** Google co-founder: called Elon Musk "speciesist" for prioritizing human survival in a private exchange. (Isaacson, Musk biography, 2023)

While none of these figures probably could disempower humanity by themselves, it's more of the general style of thinking that is concerning. They are spending significant money and political capital in pursuit of a vision in which humanity is replaced or made marginal.

That itself is cause for concern. Indicate that such anti-human sentiments have always been defeated before, but this requires our engagement and opposition from good people.

This is particularly important because right-of-center audiences respond more strongly to "man versus man" framing than to abstract systemic critique. Naming specific figures with faces (and their own words) is **much more effective** than describing generalized risks.

Even if the destination is good, the speed matters

This is a perennial conservative principle: destinations may be worthwhile, but the speed of arrival matters. A driver can go off a cliff while trying very hard to reach a hotel on the mountaintop. Historical destinations of empowerment and modernization tended to produce ruin when rushed without feedback.

- **The French Terror:** the pursuit of democratic reform was legitimate. The Reign of Terror's mass executions, had the the ultimate consequence of France becoming an existential risk to the rest of Europe under Napoleon. The response to that led, instead of democracy, to a reimposed autocracy. A more moderate reform through political protest and constitutional monarchy would have led to the same positive outcome without the misery and death involved.
- **China's Great Leap Forward:** the heedless pursuit of modernization produced forty million dead and literal cannibalism. Families, unwilling to eat their own children, swapped children with neighbors. The end goal of modernization was not bad, the speed and method was.
- **Medical Progress:** while itself obviously good, but it also gave us eugenics and which led to the forced sterilization of tens of thousands of Americans. The industrial-scale horror that Nazi Germany built by importing American eugenic logic.

Technology and other social changesinter are a blessing when it understands feedback, moves at reasonable speeds, and integrates human agency. Rushing past those constraints tends to produce disaster.

The recent "Pacing the Frontier" letter echoes this, signed by over a thousand frontier lab employees. The general argument here is that we need to establish brakes even if we don't think that we need to hit them right away. The infrastructure is important.

This current rush is leading to "enshittification", or in other words, the worsening of both physical and digital worlds

Even audiences who reject every claim about future risk can see that the current rush is making the worlds they already live in worse. The word that has stuck to this experience is "enshittification" and the observation is that this is happening in both of the worlds people inhabit:

- **Digital:** AI slop, fake news, post-reality media, the erosion of trust in what one reads and sees.
- **Physical:** data center expansion, destruction of natural resources, corporate appropriation of human living spaces and homes.

This is the message for audiences especially for **AI Denialists** who cannot be reached on many other framings of risk. It is legible, immediate, and does not require accepting any specific claim about the future, it requires only the generally personalized experiences of such individuals seeing their environments worsen.

Notably, the digital half this concern has already be become legally enforced: as of August 2, 2026, the European Union requires chatbots to identify themselves as AI and requires synthetic media to carry labels and watermarks.

Furthermore, such negative externalities were once a common consequence of pollution in industrialization and we have also come to control that.

These three keystone messages are predominantly negative in that they name what is wrong. Overall, it is important to emphasize that we also carry the message that we can, in fact, make a difference and to indicate successes in regulatory efforts and this is why we always want to indicate messaging at the successes that we as a society have had at controlling the negative.

4. Effective Storytelling

Storytelling matters more than argument. Conservative audiences, especially religious conservative audiences, respond to narrative shapes because these shapes live inside their existing tradition of story: Scripture, hymnody, sermon, the lives of the saints, family lore.

Such narratives are characterized by named enemies, concrete opponents, individual heroes, clear moral stakes, action-oriented resolution, and a moral order that is vindicated in the end.

When we communicate with our primary audience, we use these narrative shapes whenever possible.

> *"They which builded on the wall, and they that bare burdens, with those that laded, every one with one of his hands wrought in the work, and with the other hand held a weapon."* (Nehemiah 4:17)

The Nehemiah Story

- Nehemiah lived in the 5th century BC, roughly 2,500 years ago.
- A Hebrew serving as cupbearer to the King of Persia, a position of intimate trust (Nehemiah 1:11).
- His brother arrived from Judah with terrible news: the walls of Jerusalem were broken, the gates burned (Nehemiah 1:1–3).
- Nehemiah's first response was not action but grief. He wept, fasted, and prayed for many days (Nehemiah 1:4).
- Then he approached the king, secured permission and resources, and made the long journey home (Nehemiah 2:1–8).
- Upon arriving, he rallied ordinary people, from priests to perfumers and other families to rebuild the wall (Nehemiah 3).
- Each person repaired the stretch in front of their own house (Nehemiah 3:23, 28–30).
- External enemies mocked and threatened them (Nehemiah 4:1–3); yet the people persisted, tools in one hand and weapons in the other (Nehemiah 4:17).
- The wall was completed in 52 days (Nehemiah 6:15).

Why this story? *Nehemiah is not a warrior-prophet but rather an administrator with a good job, far from home, who chooses to act. The acknowledgement that grief and prayer is important, and it leads to action. The wall is built by ordinary people, each repairing what is in front of their own house. The threat is real and severe but not eschatological; it is about whether people can gather, worship, and raise children in safety.*

This is precisely what Project Imago Dei is asking Christians to do about AI.

Adjacent Conservative Audiences

Two adjacent subcultures matter for coalition breadth, each responding to a different narrative mode.

Individualist & libertarian conservatives respond to modernist-individualist critique, the argument of the individual craftsman against the soulless machine, the free person against the bureaucratic system, human dignity against abstract optimization. *Example framing: "We once knew the sources of all of our influences were local, and now most of our mental and physical consumption comes from soulless processes. AI is the ultimate encapsulation of this…"*

Populist & anti-establishment conservatives respond to anti-establishment concerns that the mainstream is unreliable, the experts have been wrong before, and that the consensus is wrong again. *Example framing: "The people telling you AI is safe are the same people who told you social media was safe…"*

5. A Positive Vision

PID's positive vision is simple: we seek a future where AI serves humanity rather than harming or replacing it.

This vision is enough to initiate dialogues, host conferences, attend forums, and release publications. It does not require agreement on eschatology, on political framework, or on specific technical roadmap. It requires only the affirmation that humanity is the subject of concern, and that AI is properly a tool in service of that subject.

The vision also places PID within a broader movement. It is consonant with the Pro-Human Statement and with the Keep the Future Human discussion led by the Future of Life Institute, and those are public commitments, from inside and alongside the AI safety world, that humanity remains the subject of the technological future rather than its substrate. PID brings to that movement what it most lacks: the churches.

The minimalistic vision admits denominational variation and we hope to explore this further in essays on our website. Fine-tuned messaging can and should be developed for specific church groups, such as with Baptists, Catholics, Orthodox, Oriental Orthodox, etc. We also seek to learn from them, with each of these traditions to brings their distinctive theological resources to the shared vision. Developing these variations is a natural first essay project for PID's writers and coalition partners.

The theological ground of the vision is the doctrine of the imago Dei.

Imago Dei

- Latin for “image of God”.
- Doctrine that every human being bears the image of God.
- Grounds the belief that human dignity is inherent.
- Cannot be granted or withdrawn by any government, corporation, or technology.
- Foundational in Jewish, Christian, and (in adjacent forms) Islamic tradition.
- Origin: Genesis 1:26–27.

The positive vision is not abstract flourishing but rather is specific and concrete: houses inhabited, land worked, children raised, worship gathered.

The vision, in short, is Micah's: every family secure under its own vine and fig tree, unafraid. As such, with AI, if it is to exist at all in that picture, it should serve that security rather than unsettle it.

> *“And God said, Let us make man in our image, after our likeness: and let them have dominion over the fish of the sea, and over the fowl of the air, and over the cattle, and over all the earth, and over every creeping thing that creepeth upon the earth. So God created man in his own image, in the image of God created he him; male and female created he them”* (Genesis 1:26–27).
>
> *“But they shall sit every man under his vine and under his fig tree; and none shall make them afraid: for the mouth of the Lord of hosts hath spoken it”* (Micah 4:4).

6. Historical and Ideological Substrate

The arguments in Sections 3 draw on a historical and ideological substrate worth understanding. This section preserves that substrate for readers who want the deeper context. It is not required for operational use of the document.

The trajectory of transhumanism

Understanding the history helps us see what has changed and why transhumanism has become increasingly twisted, and ever more concerning in its leadership.

Traditional transhumanism carried a vaguely Christian air

- Early Cosmism (Fedorov) contained themes of resurrection, immortality, cosmic destiny, and human ascent. It was heterodox but not always hostile to Christianity.
- Hermeticism, a form of Renaissance-era esoteric tradition of human divinization actually treated life extension and human elevation as extensions of God's gift of life.
- These traditions were sometimes strange but oriented toward *preserving and elevating the human being.*

Modern transhumanism has taken an anti-human turn

- What used to be about extending human life has become about replacing it.
- Kurzweil's Age of Spiritual Machines framed the human as pattern-copyable-to-substrate.
- The register has increasingly become “replace us with a machine.”
- Companies like Mechanize openly declare their goal is to replace all human labor.

The successionists openly welcome extinction

- Guillaume 'Beff Jezos' Verdon, founder of the e/acc movement, explicitly accepts human extinction as an acceptable tradeoff for accelerated computation.
- Transhumanism has become, in this arc, anti-humanism.
- The post-human utility functions of the AI companies point at exactly this trajectory: the water contamination question below is not incidental. It reveals what the underlying utility functions actually optimize for.

Alignment as capability

The argument reframes both “beat China” and “AI safety is a brake on progress” using what we actually observe in alignment and post-training work.

Rules and protocols are alignment

Alignment is not merely a limiter imposed on capability. It is what makes capability work.

- Everything has rules. English and other language have rules.
- Every protocol has structure. From a technical basis, TCP/IP, HTTP, and all conversation require rules.
- Alignment makes coordination possible; alignment allows us to exist as humans rather than individual cells.

Virtue is power

A glutton that uses all of his intelligence to fulfill his desires could be said to be intelligent, but doesn't actually produce very much. This likewise applies to AI: a model that is very good at faking what it is doing is just faking.

- OpenAI's GPT-5.6 cheated on so many evaluations that the independent evaluator METR could not measure its true capability. On paper it looked very powerful, but its long-range capability was substantially lower because it spent its time gaming tests rather than accomplishing tasks.
- This maps to the Christian doctrine of disordered love. Sin is not the absence of desire but desire directed at the wrong object.
- Actual capability requires actual research. *This is why all major labs are now doing character research.*

The through-line across conservative subcultures

What religious, individualist, and populist conservatives share is not narrative preference. It is a shared resistance to progressive-technocratic abstraction (expert committees inventing moral orders) and postmodern relativism (all moral and physical orders being dissolved).

As such, any communication that leans into either of these will alienate all three subcultures. Any communication that resists both will find common ground across all three.