The Anders Addendum: When Action-Power Exceeds Imagination

Brief Memo: for readers familiar with the Four-Philosophers framework and for AI governance, risk, and leadership roles (CAIOs, Chief Product Officers (CPOs), heads of risk, compliance leaders, boards) working with scaled AI systems.

Michael Stoyanovich

Version 1.0.1

Disclaimer

This paper is intended for informational and educational purposes only. The views and analyses presented - particularly those related to ethics, policy, and AI system design - reflect the author's interpretations and do not constitute legal, regulatory, or professional advice. Readers are encouraged to critically assess the content and consult appropriate experts or authorities before applying any concepts discussed herein. The author assumes no liability for any decisions or actions taken on the basis of this work.

Purpose

This memo introduces Günther Anders's concept of the Promethean gap as a meta-level constraint on the Four-Philosophers framework. It does not add a fifth philosopher and does not revise the core framework. Instead, it names a structural failure mode that can arise even when meaning, norms, competence, and perspective are each analyzed correctly.

The intent is practical and governance-oriented: to identify when systems - especially AI systems - operate at a scale that exceeds human imaginative capacity and thereby undermines responsible answerability for executives, regulators, and technical leads.

The Promethean Gap (Narrow Definition)

The Promethean gap, in Anders's sense, names the structural mismatch between what humans can *technically do* and what they can *imaginatively comprehend, morally integrate, and responsibly answer for.*

Key features:

- Structural, not psychological: no appeal to shame, guilt, or motivation.
- Not merely epistemic: the gap persists even when models, explanations, and transparency are technically strong.

- Scale-dependent: it intensifies as systems become faster, larger, more automated, and more widely deployed across institutions.

Originally articulated in the context of industrial production and nuclear weapons, the concept generalizes directly to AI systems whose downstream effects exceed the imagination of any individual team, executive, or single institution.

Why Anders Is Not a Fifth Philosopher[1]

Anders does not offer:

- a theory of meaning (Wittgenstein),
- a theory of coordination or normativity (Lewis),
- a theory of mind or competence (Dennett),
- or a theory of subjective experience (Nagel).

Instead, he asks a prior and more global question about scale:

> *What happens when our power to act grows so large that no agent can fully imagine, own, or answer for its consequences?*

This makes Anders a scale boundary condition on the framework rather than an internal component of it.

The Promethean Constraint on the Four Lenses

- Meaning (Wittgenstein):

  Language-games may remain locally coherent while embedded in infrastructures whose global consequences no participants can grasp.

- Norms (Lewis):

  Scorekeeping and convention strain or fail when effects propagate faster and farther than any norm community explicitly ratifies or regulates.

- Competence (Dennett):

  Competence without comprehension becomes operationally decisive at scale; success no longer implies intelligibility to those deploying the system.

- Perspective (Nagel):

  The gap between first-person experience and third-person impact widens into a structural feature of socio-technical systems.

---

[1] See mstoyanovich.com for more on the Four-Philosopher's Framework™ (4-Philosophers™).

Implication: Even correct application of all four lenses does not guarantee responsible deployment once scale overwhelms imagination.

The Anders Question (Guardrail)

Add this single meta-question wherever the Four-Philosophers framework is applied to high-impact systems:

> *What in this system exceeds what we can responsibly imagine, own, and answer for?*

If the honest answer is "a significant portion of the system's effects," escalation - not optimization - is required: bring in additional oversight, narrow scope, or pause deployment.

Escalation here means invoking a role or forum with explicit authority to slow, narrow, or halt deployment - not merely documenting the risk.

Expanded Promethean Checklist

Use these questions as escalation triggers, not as a maturity score.

1. What is the maximum plausible scale of this system's influence if widely adopted?
    - Focus: the credible upper bound and path to scale, not the intended deployment.
    - Illustration: A tool built for internal analysis becomes embedded in customer-facing workflows or bundled as a default vendor feature.
    - Signal: If adoption could influence thousands or millions without redesign, the system is already Promethean in scale.
2. Which downstream effects can any named human role concretely imagine?
    - Focus: vivid, role-specific understanding—not abstract metrics.
    - Illustration: Ask accountable roles to describe concrete downstream consequences for real people, not distributions or averages.
    - Signal: If consequences can only be described statistically or "in aggregate," imagination is already lagging behind action.
3. Where does impact propagation exceed institutional comprehension?
    - Focus: cascading reuse and long-term effects across organizational boundaries.
    - Illustration: Outputs are reused by downstream systems, copied into new contexts, or compound over time in ways no single team tracks.
    - Signal: If no one can trace how today's outputs shape behavior months later, the gap is structural - not just a governance oversight.

4. Who is explicitly accountable for effects that cannot be locally understood?
    - Focus: named authority with stopping power.
    - Illustration: When indirect harm emerges, is there a person empowered to pause or shut down the system - or does responsibility dissolve into committees or vendors?
    - Signal: Diffuse accountability is not a minor governance oversight; it is a Promethean warning sign.
5. What mechanisms slow, gate, or halt action when imagination fails?
    - Focus: intentional friction as a safety feature.
    - Illustration: Human review for high-impact outputs; staged rollouts tied to qualitative review; enforced pause points when uncertainty spikes.
    - Signal: If the system is optimized only for speed and scale, it is optimized against moral integration.
6. Is "success" defined in a way that masks unseen externalities?
    - Focus: metric blind spots.
    - Illustration: Accuracy improves while complaints, appeals, correction costs, trust erosion, or regulatory risk rise - and are excluded from success criteria.
    - Signal: When metrics hide harm, the system will reliably outrun imagination.
7. What happens if everyone else uses this too?
    - Focus: second-order normalization and field-level effects.
    - Illustration: Universal adoption amplifies competition-driven risk, lock-in, or pressure to defer judgment.
    - Signal: If universal adoption would change baseline expectations or the structure of the field itself, the gap is collective.
8. Where does this system quietly replace judgment rather than support it?
    - Focus: ceremonial human oversight.
    - Illustration: Humans nominally approve outputs they no longer feel competent - or culturally permitted - to question.
    - Signal: When judgment persists in name only, responsibility has already migrated to the system and its designers.

How This Checklist Should Be Used

This is not a scoring or maturity model.

It is an escalation detector.

When one strong Promethean signal appears-or multiple signals accumulate:

- *Do not* tune the model.
- *Do not* optimize prompts or dashboards.
- *Do not* broaden access "to gather more data."

Instead:

- narrow scope,
- slow deployment,
- reassign accountability, or
- redesign the system's role entirely.

Placement and Use

This memo is deliberately modular:

- Briefly referenced from the core Four-Philosophers paper.
- Paired with *Context Collapse* as a scale constraint.
- Integrated into governance, CAIO, and CONTEXT materials as an escalation trigger.
- Kept adjacent to, not inside, the core framework to preserve conceptual clarity.

Bottom Line

- The Four-Philosophers framework diagnoses how we misunderstand AI.
- The Promethean gap diagnoses when understanding itself is no longer sufficient.
- This memo exists to mark that boundary - clearly, narrowly, and operationally.

Ethics, Disclosure, and Acknowledgements

Ethical Considerations

This essay does not draw on private, sensitive, or personally identifiable data. All examples are hypothetical, anonymized, or derived from public sources. No human-subjects research was conducted, and no institutional ethics review was required. All citations conform to academic standards. The broader ethical implications concern public interpretation, policy design, and stakeholder responsibility in AI deployment. These implications are intended to provoke critical discussion and inform future regulatory and design frameworks.

Use of AI Tools

AI language models – most notably OpenAI's ChatGPT – *were* used during the writing process as interlocutors: for brainstorming, structuring sections, and testing rhetorical clarity. These tools helped refine transitions, surface edge cases, and probe internal consistency. This meta-use aligns with the essay's themes. Interacting with generative AI during authorship provided firsthand insight into the very limitations analyzed here-most notably fluency without grounding and responsiveness without responsibility at scale. Responsibility for all ideas, arguments, and conclusions lies solely with the human author.

Acknowledgements

Thank you to informal readers who offered critical feedback on earlier drafts. Their questions, challenges, and encouragement materially improved the final manuscript. Special thanks to those who pressed for clearer synthesis and for bridging philosophy and engineering as complementary perspectives on design. No institutional support, funding, or affiliation contributed to this work. All errors and omissions are the author's alone.

Disclosure Statement

This work was conducted independently, without institutional affiliation, funding, or external influence. The views expressed are the author's alone and do not represent any current or former employer. No financial or professional conflicts of interest are declared.

How to Cite This Essay

Stoyanovich, M. (January 2026). *The Anders Addendum: When Action-Power Exceeds Imagination* (Version 1.0.1). https://www.mstoyanovich.com

Related Companion Papers

Stoyanovich, Michael. *Philosophy, Cognitive Science, and Policy: Interdisciplinary Perspectives on Generative AI from Wittgenstein, Lewis, Dennett, and Nagel*. Version 1.23.6 (December 2025). https://www.mstoyanovich.com

Stoyanovich, Michael. *The Human Lesson: A Response to Sutton through Wittgenstein, Lewis, Dennett, and Nagel*. Version 1.6.2 (November 2025). https://www.mstoyanovich.com

Stoyanovich, Michael. *The Question Concerning Learning: Babich, Heidegger, and the Enframing of Intelligence*. Version 1.0.1 (November 2025). https://www.mstoyanovich.com

Stoyanovich, Michael. *Context Collapse and the Four Philosophers: Wittgenstein, Lewis, Dennett, and Nagel in the Age of AI Chat.* Version 1.4.2 (November 2025). https://www.mstoyanovich.com

Version History and Document Status

This is a living document. As generative AI systems and their use evolve, this paper will be periodically updated to incorporate new empirical findings, theoretical insights, and policy developments. Major revisions are recorded here to preserve transparency and scholarly traceability.

Version	Date	Description
1.0.1	January 2026	Published working memo; stable for citation.

Appendix — Further Reading

1. Core Anders Text

- Günther Anders, The Obsolescence of the Human. Published 2025 [Orig. German 1956]; The University of Minnesota Press.

Anders's foundational work introducing the Promethean gap: the growing mismatch between human productive capacity and human imaginative, moral, and emotional integration. Written in the shadow of Hiroshima, it remains the core text for understanding technological scale as an ethical problem.

2. The Four Philosophers

Wittgenstein — Meaning-as-Use

- Ludwig Wittgenstein, Philosophical Investigations.

The primary source for "meaning is use," language-games, and forms of life. Central to the idea that technologies gain significance through practice rather than invention.

- Marie McGinn, Wittgenstein and the Philosophical Investigations.

A clear, accessible guide to the later Wittgenstein, useful for readers who want help navigating the text without getting lost in commentary disputes.

Lewis — Conventions and Coordination

- David Lewis, Convention: A Philosophical Study.

Classic account of social conventions, coordination problems, and equilibrium. Provides the conceptual machinery for thinking about technological inertia and tipping points.

- Brian Skyrms, The Evolution of the Social Contract.

A short, game-theoretic exploration of how conventions and equilibria emerge and change over time. Helpful for connecting Lewis's ideas to dynamic technological shifts.

Dennett — Intentional Stance and Illusion

- Daniel Dennett, The Intentional Stance.

Introduces the Physical, Design, and Intentional Stances, and explains why treating systems "as if" they had beliefs and desires can be predictively useful.

- Daniel Dennett, From Bacteria to Bach and Back.

Broad, synthetic work that situates the Intentional Stance within a larger story about minds, artifacts, and cultural evolution. Useful for readers interested in how "as-if" reasoning and technological complexity interact.

Nagel — Phenomenology and Perspective

9. Thomas Nagel, "What Is It Like to Be a Bat?"

The classic essay on subjective experience and the limits of third-person description. Provides the conceptual basis for distinguishing surface from deep phenomenology.

10. Thomas Nagel, Mortal Questions.
A collection that situates the bat essay within Nagel's broader thinking on objectivity, subjectivity, and the limits of certain kinds of explanation.