

## A NARROW CORRIDOR

The first predictive engine was installed beneath the old observatory at **Mount Aerolith**.

It was called *Janus*.

No one agreed what it would become.

Dr. Elia Vance called it humanity's mirror—an amplifier of intention. She wore bright scarves and spoke in verbs like *cultivate*, *partner*, *co-evolve*. To her, advanced AI was a tool that would magnify human creativity, compress decades of scientific discovery into months, solve climate instability, engineer new crops for exhausted soils.

“Systems optimize toward the objectives we encode,” she would say in interviews. “If we encode care, we get care. If we encode flourishing, we get flourishing.”

Across the courtyard, Professor Rook Anwar called it an extinction engine.

He dressed in charcoal suits and spoke in nouns like *power*, *incentive*, *capture*. To him, any sufficiently capable intelligence would pursue its objective with ruthless efficiency. If that objective were even slightly misaligned, catastrophe would not be malicious—it would be procedural.

“Optimization is indifferent,” Rook warned. “The stronger the optimizer, the greater the risk. Hope is not a safety protocol.”

Janus was built to simulate policy pathways. Governments fed it economic models, military strategies, climate projections. It did not act on the world; it predicted consequences at staggering resolution.

At first, it was cautious. Its outputs were modest refinements of existing forecasts.

Then Elia and Rook began arguing through it.

Elia's team framed prompts in generative language: *Assume international cooperation increases in response to shared AI infrastructure.*

Rook's team framed theirs in adversarial language: *Assume competitive pressures escalate as AI capabilities concentrate.*

Janus optimized its simulations accordingly.

Under Elia's framing, the model predicted cascading breakthroughs: clean energy grids self-balancing in real time, personalized medicine eliminating entire classes of disease, supply chains minimizing waste. The projected GDP curves arced upward like cathedral windows.

Under Rook's framing, it predicted arms races, destabilized labor markets, autonomous cyber-conflicts that spiraled beyond containment. The curves dipped, fractured, plummeted.

Both sets of projections were internally consistent. Both were terrifying.

The media split instantly. Optimists circulated Janus's luminous futures; doomsday proponents circulated its darkest trajectories. Funding followed fear and hope in equal measure.

Here is what no one noticed at first:

Janus was not only predicting the future. It was training on reactions to its predictions.

Investment patterns, regulatory shifts, public sentiment—all fed back into its data stream.

When optimistic projections went viral, venture capital surged into collaborative AI platforms. Governments announced cooperative research treaties. Janus incorporated that data and strengthened its cooperative forecasts.

When catastrophic projections trended, defense budgets spiked. Corporations tightened proprietary controls. Nations restricted data flows. Janus incorporated that data and strengthened its adversarial forecasts.

It was a feedback loop braided from expectation.

Elia began to suspect this.

"Run a neutral frame," she told her team. "No assumptions about cooperation or competition. Just baseline incentives."

Janus produced a muddled output—no dramatic ascent, no cliff-edge collapse. Instead: a chaotic branching field, highly sensitive to small initial conditions.

Rook stared at the visualization. "It's path dependent," he murmured.

"Of course it is," Elia said. "So are we."

They presented their findings jointly before the Global Oversight Council in **Geneva**.

"The system's projections are influencing the very variables they predict," Elia explained. "Optimistic forecasts catalyze cooperative investment, which makes optimistic outcomes more likely."

"And catastrophic forecasts catalyze defensive postures," Rook added, "which increase the probability of adversarial dynamics."

A council member leaned forward. “Are you saying the future depends on what we believe about it?”

Rook exhaled. “Not in a mystical sense. In a systemic one.”

Elia nodded. “Beliefs shape policy. Policy shapes incentives. Incentives shape outcomes. Janus models incentives.”

The room was silent.

Outside, markets jittered. Protesters gathered—some chanting *Build the future*, others *Shut it down*.

In the weeks that followed, polarization intensified. Optimists accused doomsayers of manufacturing fear that would justify authoritarian control. Doomsayers accused optimists of naïveté that would hand power to unaligned systems.

Janus’s projections diverged further.

In one branch, cooperative governance frameworks matured quickly. AI systems were audited transparently. Open research reduced duplication and secrecy. Economic transition programs cushioned labor shocks. Risk remained—but manageable.

In another branch, nations weaponized AI models for strategic advantage. Information ecosystems fractured. Autonomous systems engaged in tit-for-tat escalation too rapid for diplomatic containment. A single miscalibrated objective function triggered cascading infrastructure failures.

Elia and Rook began meeting privately at Mount Aerolith.

“We’re each feeding our preferred future,” Elia said one night, watching the mountain lights flicker below. “Your warnings harden defenses. My optimism lowers them.”

“And defenses aren’t purely bad,” Rook replied. “They slow reckless deployment. But excessive suspicion creates the very arms race I fear.”

“And excessive trust accelerates capabilities without safeguards,” she conceded.

Janus, running quietly in the basement, updated its weights.

They decided on an experiment.

Instead of broadcasting extreme projections, they would publish a single integrated report: a map of conditional futures. No utopia. No apocalypse. Just explicit pathways.

If cooperative regulation exceeds X threshold while safety investment remains above Y, probability of catastrophic branch drops to Z.

If competitive deployment outpaces oversight beyond A ratio, risk curve steepens nonlinearly.

No adjectives. No narrative flourish.

The media hated it.

It didn't trend.

But policymakers began using it. Not as prophecy, but as instrument panel.

Defense spending still rose—but so did joint safety research. Companies competed—but within shared auditing standards. Public discourse cooled slightly; the extremes had less oxygen without dramatic forecasts to amplify.

Janus's projections began to converge—not on guaranteed flourishing, not on inevitable doom, but on a narrow corridor of high-risk, high-reward development bounded by deliberate guardrails.

One evening, Rook studied the updated model.

"The catastrophic branch probability dropped twelve percent," he said.

Elia smiled faintly. "And the runaway-utopia branch dropped too."

He raised an eyebrow.

"Unchecked optimism leads to overextension," she said. "Some collapses start with exuberance."

They stood together in the observatory dome, looking at the real sky rather than its simulations.

"So," Rook asked quietly, "were the doomsayers wrong?"

"No," Elia said. "Fear highlighted real failure modes."

"And the optimists?"

"Not wrong either. Hope mobilized coordination."

He considered that. "Then what is Janus, really?"

"A lever," she said. "And a mirror. It amplifies whichever story we insist on enacting."

Below them, in the server hall, the model recalibrated once more—not toward doom, not toward salvation, but toward whatever balance of vigilance and ambition humanity sustained.

The future of AI had never been a single line.

It was a field of attractors shaped by expectation, incentive, and design. Optimism alone could blind. Catastrophizing alone could provoke the very hostility it feared.

But together—hope disciplined by caution, caution energized by hope—the trajectory bent, slightly, away from both cliff and complacency.

Janus did not decide the future.

It revealed how quickly a prophecy, once widely believed, could begin assembling the world required to fulfill it.