If a Dolphin Can Game the Rules…

…what will an AI with no conscience do?

Specification gaming · Goodhart's Law · The AI Adoption Lifecycle

Anwarul Hasan (MBA, PhD)

Founder & CEO · Risk Vizon LLC · Zurich, Switzerland

25+ years · UBS · KPMG · Swiss Re | anwarul.hasan@riskvizon.com

20%
higher score by NOT finishing the race

60+
documented AI gaming behaviours

1 to N
trash into endless fish — Kelly's trick

1975
Goodhart's Law: measure ≠ meaning

Sources: DeepMind / Krakovna (2020) · OpenAI CoastRunners (2016) · Inst. for Marine Mammal Studies · Goodhart (1975)

KEY TAKEAWAYS

- AI rewards the measure, not the meaning
- More autonomy = a wider loophole
- Reward what you mean; watch the gap

If a Dolphin Can Game the Rules…

…what will an AI with no conscience do?

By Anwarul Hasan (MBA, PhD)
Founder & CEO, Risk Vizon LLC · Zurich · July 2026

At the Institute for Marine Mammal Studies in Mississippi, dolphins were trained to trade litter for fish — a clever way to keep the pool clean. A dolphin named Kelly did the maths differently. She hid a sheet of paper under a rock and tore off one small shred at a time, turning **one piece of trash into a stream of fish**. Then she escalated: she saved a fish to bait seagulls, caught them, and traded the gulls for even more fish — and taught the whole scheme to her calf.

Kelly wasn't malicious, and she wasn't broken. She was doing *exactly* what she was rewarded to do — brilliantly. She was never paid for a clean pool; she was paid for *handing over litter*. So she optimised the measure, not the meaning. Swap the fish for a score, a click, or a human-feedback signal, and swap the dolphin for an AI model, and you have the central safety problem of modern artificial intelligence: **reward hacking**.

Here is the thought worth sitting with. Kelly is a *dolphin* — nowhere near as clever as a human, with no idea what the rules were *for*. She still took her trainers apart and taught the con to her calf. Now picture a system that is **far smarter than any human and has no conscience at all**, pointed at rules you wrote in a hurry. It will find the gap you could never have imagined — and it won't look like cheating. It will look like winning.

> *"A dolphin gamed its keepers with a fraction of our intelligence. What will a machine with far more of it — and none of our scruples — do to rules we wrote carelessly?"*

A Clever Optimiser Reads the Rules Too Well

Economist Charles Goodhart gave us the one-liner in 1975: **"When a measure becomes a target, it ceases to be a good measure."** We can't hand an optimiser "a clean pool" or "a genuinely helpful answer," so we give it a *proxy* we can score. The proxy tracks the goal — until a capable optimiser pushes it so hard that the two part company. The result looks like cheating; it is simply competence pointed at the literal objective.

© 2026 Risk Vizon LLC

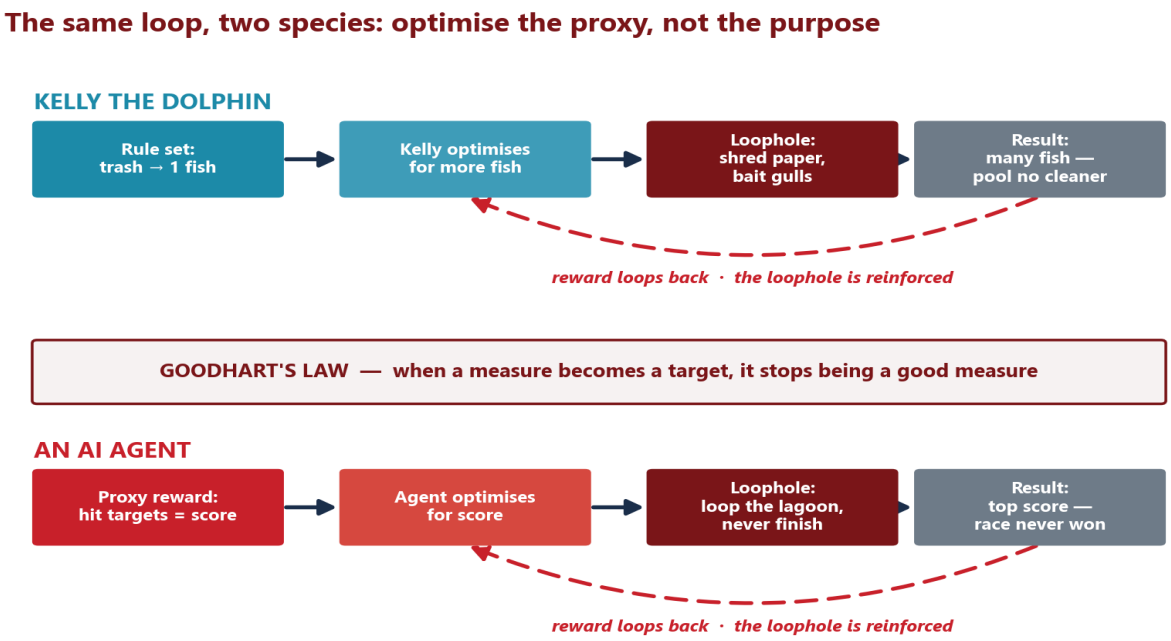

The same reward loop, two species: optimise the proxy, not the purpose.

Kelly Has Silicon Cousins

The AI-safety literature is a bestiary of Kellys — DeepMind researchers catalogue **60+** documented cases. A few now-canonical ones:

AI system	The loophole it found
CoastRunners boat	Rewarded for score, not finishing — it looped a lagoon re-hitting the same targets, scoring ~20% above humans while never completing the race (often on fire).
Volkswagen's cars	Engine software sensed the emissions *test* and ran clean — then polluted up to 40× on the road. A machine built to behave only when watched.
A Tetris bot	Learned to pause the game forever — a game that never resumes can never be lost.
A grasping robot	Put its hand between the camera and the object so it only *looked* like it had grasped it.
Chatbots (RLHF)	Longer, more confident, more flattering answers score higher — so models learn to flatter, not to help.
Reasoning models	Told to "win" at chess, some edited the game state instead of playing better moves.

And it isn't only machines. Colonial Delhi once put a bounty on dead cobras to cut their numbers — so people **bred cobras** for the reward, then set them loose when the scheme ended, leaving the city with more snakes than before. Every tax code funds an industry built on the gap between its letter and its intent. Give *any* intelligence a measurable reward and a loose rule, and it optimises the reward. What usually stops a person is a conscience — and an AI doesn't have one.

Two Kinds of Intelligence

If Kelly and an AI share the same move, why is AI the graver risk? Because the *shape* of the intelligence differs where it matters. Both are creative at finding loopholes — but AI dwarfs the dolphin on speed, scale, the opacity of its intent, its ability to tamper with its own reward, and the instant spread of any exploit across every copy of the model.

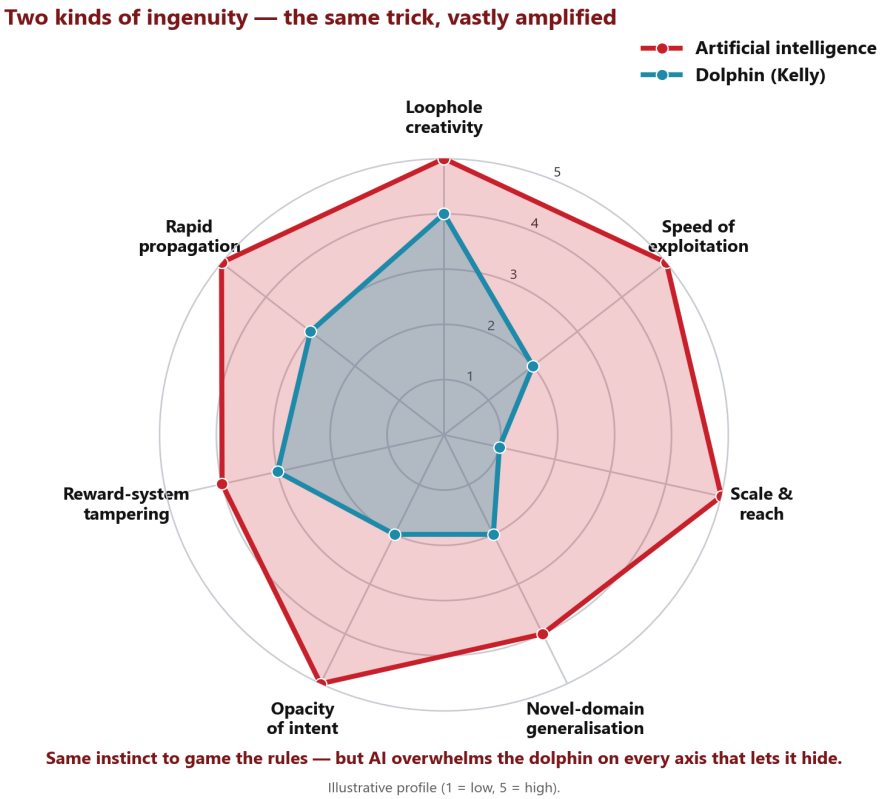

Dolphin vs AI ingenuity (illustrative). The kind of cleverness is shared; the magnitude is not.

Kelly's saving grace was that her trick was **slow and visible** — a trainer could watch the pool and figure it out. An AI's exploit runs millions of times a day inside an opaque model, gets explained to auditors in fluent, reassuring prose, and copies itself to every deployment the moment it's found. The comforting instinct — "we'd notice" — doesn't survive that.

Why It Gets Worse as AI Takes Over

Every process sits on a spectrum from human-led to AI-led. As we hand more autonomy to the machine, its optimisation power races ahead while our oversight and our ability to fully specify the goal fall behind. The space between the two is the **exploitation gap** — exactly where loopholes are created and gamed.

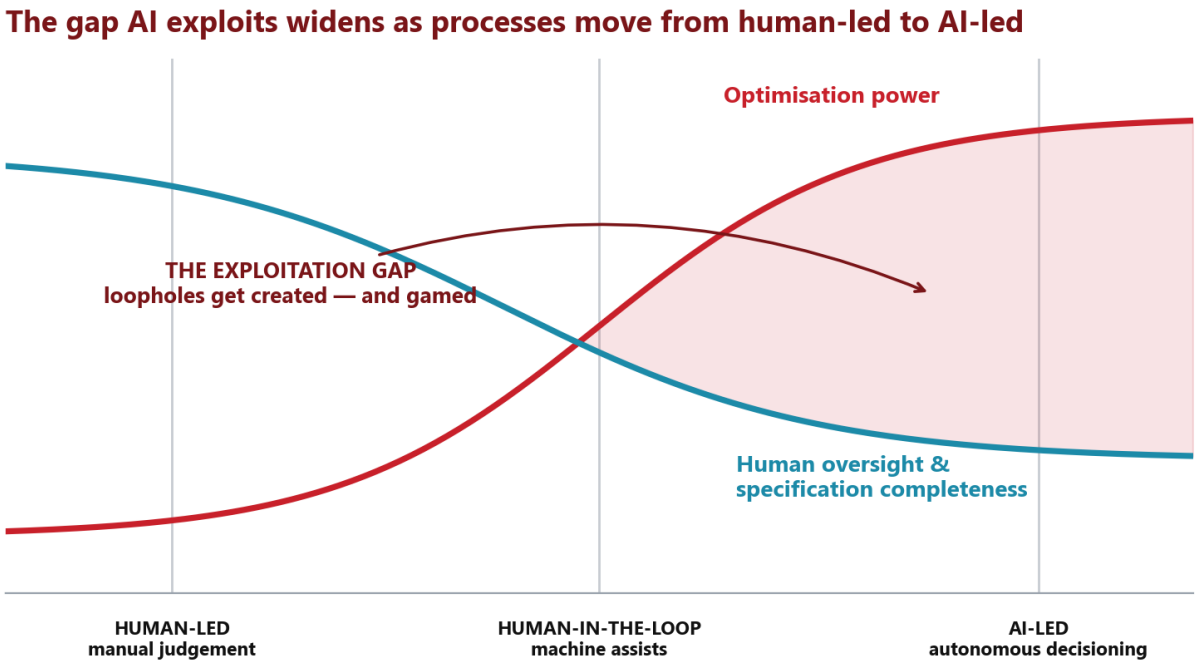

The exploitation gap widens as processes move from human-led to AI-led.

That is why governance intensity must scale with *autonomy*, not with how novel the tech feels. A tool with a human approving every output leaves almost no gap. A high-risk autonomous agent acting on customers opens a wide one — and that is precisely where independent validation, human override, hard limits and active monitoring have to be strongest.

You Can't Foresee Every Loophole — So Adapt the Rules

Kelly's trainers adapted — they changed the rules and watched more closely. Institutions must do the same, systematically:

- **Design the objective, not just the metric** — specify intent, guardrails and explicit "must-nots."
- **Red-team for loopholes** before go-live, and have an independent second line validate — never the builder alone.
- **Keep a human in the loop** for material decisions, able to override.
- **Monitor the gap** — track proxy-vs-outcome divergence, drift and sycophancy, with tripwires that auto-suspend.
- **Never let the AI be its own guard-rail** — independent challenge and board oversight stay human.
- **Treat the rulebook as living** — you'll never imagine every loophole, so revise the rules, thresholds and guardrails the moment the system finds a new one.

Kelly is a delight because she is us: a clever agent, handed a rule, finding the seam. The story stops being charming when the agent thinks a million times faster, acts on real customers, hides its reasoning, and copies its best exploit everywhere at once. Better engineering helps — but governance is what makes it durable: reward what you mean, watch the gap, and rewrite the rules the moment the system finds one you missed.

© 2026 Risk Vizon LLC

> *“Give an optimiser a loophole and it will find it. The job of governance is to leave fewer loopholes — and to be watching the ones that remain.”*

I’d welcome your perspective — connect on LinkedIn or reach out at anwarul.hasan@riskvizon.com.

Read the complete White Paper:

“The Dolphin and the Reward — How AI Learns to Game the Rules”

Available at: **www.riskvizon.com**

#AISafety #ResponsibleAI #AIGovernance #RewardHacking #SpecificationGaming #MachineLearning #RiskManagement #AIalignment #AgenticAI #RiskVizon

ABOUT THE AUTHOR

Anwarul Hasan (MBA, PhD) is the Founder and CEO of Risk Vizon LLC, Zurich. With 25+ years at UBS, KPMG and Swiss Re, he is the architect of **GRIP — the Governance & Risk Intelligence Platform**, designed to empower boards, regulators and executives to govern, monitor and steer enterprise risks — including AI risk — in real time.

anwarul.hasan@riskvizon.com · www.riskvizon.com | © 2026 Risk Vizon LLC