The Wheel of the Gods

How a civilization of busy bees built an accelerating machine with tools once reserved for the gods — and forgot to ask who was steering.

Derek Van Derven

November 6, 2026

The Bees Built a Train With No Conductor

Compartmentalization, the disappearance of responsibility, and the race toward technologies we may not be able to control

One day, the problem won't be a screwed-up delivery address.

That might sound like an odd place to begin an essay about artificial intelligence, corporate power, government, law, and the future of civilization. But sometimes an ordinary experience reveals something much larger.

Consider a simple problem involving a customer, a seller, a delivery company, a corporation, several departments, customer service representatives, claims procedures, and conflicting instructions.

Each participant handles a piece. Each department has its own procedures. Each organization has its own policies. Responsibility moves from one desk to another. Nobody necessarily sees the entire problem, and nobody appears to have the authority or inclination to take ownership of the outcome.

The customer experiences one problem.

The institution experiences a collection of separate transactions.

It is a small example of something that has become increasingly familiar in modern life: the fragmentation of responsibility.

And it raises a much larger question.

What happens when this same architecture operates not on a misplaced package, but on governments, financial systems, industrial infrastructure, public policy, warfare, genetic engineering, or artificial intelligence?

The machinery of compartmentalization

The internet did not invent bureaucracy, corporate complexity, or the division of responsibility. Those things existed long before computers.

But the internet accelerated their reach.

Business became increasingly global. Corporate ownership structures expanded. Companies acquired companies, created subsidiaries, outsourced operations, and distributed their activities across borders. Governments increasingly relied on contractors, specialized agencies, and layers of administration.

Technology made it possible to coordinate enormous operations without requiring any one participant to understand the entire operation.

There are legitimate reasons for specialization. Complex societies cannot function without it.

But there is a darker side.

Compartmentalization can become a mechanism for avoiding responsibility.

A corporation can divide ownership, operations, decision-making, and liability across multiple legal entities. A government can distribute authority across agencies. An institution can separate the people making decisions from the people experiencing their consequences.

Sometimes these arrangements serve legitimate organizational purposes. Sometimes complexity itself becomes useful because it makes accountability harder to establish.

Responsibility becomes a moving target.

The person handling the complaint says it belongs to another department. The department says it follows corporate policy. The corporation points to the subsidiary. The subsidiary points to the contractor. The contractor points to the contract.

And the people affected are left trying to locate someone who can answer for the whole.

This is not merely an administrative inconvenience. It is a structural problem with moral consequences.

Sociologist Ulrich Beck called one aspect of this phenomenon organized irresponsibility: institutions can produce risks while organizing themselves in ways that make responsibility for those risks difficult to locate.

Political philosopher Dennis Thompson described a related problem as the problem of many hands: when many individuals contribute to an institutional decision, assigning responsibility for its consequences becomes extraordinarily difficult.

These are not abstract academic curiosities. They describe a recognizable feature of modern institutional life.

History has already shown us where this can lead

There is a historical reality that should never be softened for the sake of polite discussion.

The Holocaust was not carried out by a handful of people acting in isolation.

The systematic persecution and murder of six million European Jews by Nazi Germany and its collaborators required an enormous apparatus of government, administration, transportation, policing, finance, industry, professional expertise, and individual participation.

The machinery involved far more than those who directly committed murder.

Railway personnel helped transport victims. Civil servants administered discriminatory laws. Police and local authorities identified, rounded up, and deported people. Financial institutions and businesses participated in the expropriation of Jewish property. Lawyers, judges, doctors, engineers, bureaucrats, and industrialists contributed in different ways to the machinery of persecution and murder.

The United States Holocaust Memorial Museum documents the involvement of German institutions, businesses, insurance companies, banks, the railway system, healthcare professionals, civil servants, and countless individuals.

The historical record also makes something else clear: participation was not simply the product of people blindly following orders. Motivations included antisemitism, ideological commitment, careerism, greed, fear, opportunism, and personal advancement.

Some participated enthusiastically. Some actively expanded the machinery. Some benefited materially. Others stood by or helped through inaction.

This matters because the familiar defense that someone was merely doing their job does not erase what that job contributed to.

Administrative procedures did not make the crimes less real. Professional specialization did not make the participants morally invisible. The fact that one person handled transportation, another paperwork, another confiscation, and another enforcement did not divide the suffering of the people subjected to the system.

The victims experienced the whole thing.

The perpetrators and institutions helped construct it in pieces.

That is the warning.

We should not equate every modern corporation or bureaucracy with a genocidal regime. Historical distinctions matter. But we must not ignore the demonstrated capacity of bureaucratic systems to distribute participation, normalize cruelty, and obscure individual responsibility.

The lesson is not that every employee is a perpetrator.

The lesson is that institutional participation does not become morally neutral simply because it is divided into departments.

The billion-bee problem

Modern society increasingly resembles a gigantic hive.

Millions, perhaps billions, of people occupy their individual cells. They work, compete, produce, manage, invest, administer, legislate, regulate, design, code, sell, and consume.

Each has responsibilities. Each has ambitions. Each has bills to pay.

Each is encouraged to perform efficiently within the boundaries of their own position.

The corporation wants growth. The investor wants returns. The executive wants success. The politician wants influence. The professional wants advancement. The employee wants to keep a job.

None of these motivations is inherently extraordinary.

But collectively, they can produce something that nobody has adequately considered.

The hive can become so occupied with the productivity of its individual cells that it loses sight of what the hive is building.

And this is where the pursuit of competitive advantage becomes particularly dangerous.

Adam Smith's invisible hand is often invoked to explain how individual self-interest can contribute to broader economic prosperity. But the invisible hand is not a guarantee that every collective outcome will benefit humanity.

Markets do not automatically account for every external cost. Competition does not automatically protect the public. Profit does not automatically produce responsibility.

And when the consequences extend beyond the people making the decisions, the assumptions become even more dangerous.

The train is accelerating

We are now entering a period in which the same competitive machinery is being applied to technologies with consequences that extend far beyond ordinary commercial competition.

Artificial intelligence is the most immediate example.

In March 2023, the Future of Life Institute published an open letter calling for a verifiable pause of at least six months on training AI systems more powerful than GPT-4. The letter warned about an uncontrolled race toward increasingly powerful systems and called for independent safety oversight and stronger governance.

The appeal received enormous public attention and thousands of signatures.

But a public appeal is not an enforceable agreement. It did not produce a binding, industry-wide pause. The competitive development of increasingly capable AI systems continued.

That contradiction deserves attention.

The industry can publicly acknowledge the need for caution while remaining embedded in a competitive environment that rewards speed, investment, deployment, and technological advantage.

A company that slows down may fear that another will move ahead.

An executive who delays may fear losing investors. A research organization may fear losing talent. A government may fear that another country will gain an advantage.

Everyone can explain why they cannot afford to stop.

And so the train keeps accelerating.

The people building it may sincerely believe they are acting responsibly. They may fund safety research, publish principles, establish ethics departments, and issue public assurances.

But what happens when the commercial incentives of the institution conflict with the precautions it publicly endorses?

What happens when the people responsible for safety are not the people who control the commercial decisions?

What happens when the people who can slow the train are rewarded for making it faster?

Competitive advantage has a limit

Competitive advantage helped drive technological progress. It helped produce extraordinary developments in computing, communications, medicine, engineering, and industry.

But there is a fundamental difference between competing to produce a better consumer product and competing to develop technologies whose consequences may affect the future of humanity.

Artificial intelligence, nuclear technology, and genetic engineering are not interchangeable. Their risks are different, their applications are different, and their governance requires different approaches.

But they share a fundamental question:

Can we continue applying an unrestricted competition model to technologies whose consequences extend far beyond the competitors?

With nuclear weapons, there is no meaningful commercial victory in a nuclear catastrophe.

With genetic engineering, the consequences of irresponsible intervention may extend beyond the individual or institution making the decision.

With increasingly capable artificial intelligence, the question is not simply which company develops the most capable system first. It is whether the systems being developed can be reliably understood, controlled, and governed, and whether the institutions deploying them can be held accountable.

There are domains in which winning the race may not constitute winning anything at all.

We have inherited a competitive model that rewards being first, being bigger, being faster, and gaining advantage over others.

But some technologies require cooperation, restraint, independent oversight, and enforceable limits.

The pursuit of competitive advantage cannot be the only organizing principle when the consequences of failure are shared by everyone.

Powers once reserved for the gods

Throughout human history, mythology has imagined powers reserved for gods: creating life, altering nature, commanding forces beyond ordinary human capacity, and possessing knowledge beyond human understanding.

Human beings are now developing technologies that approach aspects of those imagined powers.

We can manipulate the building blocks of life. We can release enormous nuclear forces. We can construct artificial systems capable of performing increasingly sophisticated intellectual tasks.

This is not an argument against scientific progress.

It is an argument against confusing the ability to create something with the wisdom to control it.

Our technical capabilities have advanced enormously. Our political institutions, legal structures, and collective decision-making processes have not necessarily advanced at the same pace.

We are increasingly capable of producing consequences that cross national borders, generations, and institutional boundaries.

Yet responsibility remains fragmented.

The company builds one component. Another company supplies the infrastructure. A government establishes a policy. A regulator oversees a narrow area. An executive approves a deployment. An employee performs a task.

And when the combined system produces consequences that nobody can adequately contain, the question of responsibility becomes a maze.

No one is conducting the train

This is the image I keep returning to.

A billion busy bees.

Every bee working. Every bee pursuing its own objective. Every bee occupied with its own little cell.

Together they are building a supersonic train.

Some have already seen the storm ahead. Some have warned about it. Some have signed statements calling for restraint.

But the machinery continues.

The train becomes faster. The systems become more complex. The distance between those making decisions and those living with the consequences grows.

And perhaps the most dangerous part is that no conspiracy is required.

No central authority needs to command every participant. No individual needs to intend the eventual outcome.

The incentives, the competition, the compartmentalization, and the momentum can be enough.

But we cannot keep dividing responsibility and pretending that the consequences will divide themselves.

We cannot allow the people who build powerful systems to hide indefinitely behind the complexity of the systems they build.

We cannot accept public declarations of safety while treating restraint as a competitive disadvantage.

And we cannot keep telling ourselves that everyone was simply doing their job when the combined result becomes indefensible.

At some point, somebody has to look up.

Somebody has to see the entire train.

Somebody has to have the authority, responsibility, and courage to stop it.

Because the question is no longer simply how fast we can build.

It is whether we have retained the collective judgment to recognize what we are building, who it may hurt, and whether we can still stop it.

You cannot build a supersonic train with a million people working on individual components, acknowledge that a storm is coming, and then assume the train will somehow steer itself.