

AI AND THE PHYSICAL ECONOMY

How Resource Constraints Shape Technological Velocity

KRISH VASUDEVAN

Founder, TERAXYS

TERAXYS • SEPTEMBER 2026

© 2026 TERAXYS HOLDINGS LLC. All rights reserved.

I. The Physical Constraint

The development of artificial intelligence faces a dilemma similar to interstellar travel. At relativistic velocity, each additional increment of speed requires progressively more energy, with the requirement rising toward infinity as velocity approaches the speed of light. Artificial intelligence may not face an equivalent fixed physical constant; however, it faces a more complex set of physical constraints that can be both dynamic and interactive. Power, grid capacity, semiconductors, cooling, water, land, materials and industrial capacity expand on different timetables, and pressure on one can alter the severity of another constraint. The result is not a single immutable constant, but an evolving physical drag on AI's technological velocity.

Technological velocity

This difference is significant as technological velocity cannot be measured solely by improvements in model capability. It must also reflect the rate at which the next superior capability can be built, realized and deployed at economically significant scale. The resources supporting that process differ materially in their elasticity, lead times and capacity for substitution. Some can expand relatively quickly in response to investment; others require years of planning, construction, and industrial development. A constraint that appears manageable at one level of deployment can become binding at another, while relieving one bottleneck can expose another. The relevant limit is therefore not a fixed quantity of any single resource, but the changing capacity of the broader system of resources that support the next incremental unit of technological velocity.

At sufficient scale, artificial intelligence does not merely draw on productive capacity; it begins to alter the resource environment on which its further development depends. Rising demand can tighten particular constraints, while higher prices, investment, substitution and efficiency can induce new capacity in some places and relax others. The binding constraint can therefore migrate as the system responds, and pressure on one resource can change the importance of another. Scarcity by itself is not the central problem. The harder problem is that technological velocity and the physical system supporting it evolve together, while the interactions between them are not adequately captured by treating resources as fixed external inputs. The relevant question becomes which resource, or combination of interacting resources, is exerting the greatest marginal drag on technological velocity at a given point in time—and where relieving that drag creates the greatest room for further acceleration.

Drag, not scarcity

That drag can arise when the physical system stops expanding in step with technological velocity. Additional capacity in one part of the system may have limited economic value if a necessary complement in another part cannot keep pace. A shortage of one input can therefore become a bottleneck that delays, raises the cost of, or prevents the next increment of deployable capability. Drag, in this sense, is the reduction in attainable technological velocity produced by those physical constraints. It is not synonymous with scarcity: a resource can be scarce without materially slowing AI development, while a relatively narrow bottleneck can exert disproportionate drag if it constrains a much larger system or several dependent systems.

II. Endogenous Resource Drag

The severity of that drag also changes through time because of the varying responsiveness of resources supporting AI. The coefficient of drag, b , represents the severity of these constraints under prevailing conditions and can rise or fall as capacity, scarcity, substitution and efficiency change. A constraint that is relatively unimportant at one level of deployment can become increasingly consequential at another, while interactions among constraints can magnify their combined effect. Physical drag may therefore become nonlinear rather than rising in a smooth or proportionate manner.

A dynamic formulation

Given the dynamic movement in technological velocity and the corresponding changes in physical constraints, the endogenous relationship can be stated differentially as:

$$\frac{dV}{dt} = \underbrace{\sum_i A_i(\cdot)}_{\text{forces increasing } V} - \underbrace{b(Q, t)V^q}_{\text{physical drag}} \quad \text{Here, } V \text{ represents}$$

technological velocity, while Q represents the prevailing state of the physical-resource system. The first term captures the combined forces that increase it, including capital, technological improvement and the physical capacity available to support further development. The second captures the resistance imposed by that same physical system as its resources become constrained. The coefficient $b(Q, t)$ reflects the severity of prevailing resource conditions and can rise or fall as capacity tightens, expands or shifts between bottlenecks. The exponent q describes how that drag scales with technological velocity: more slowly than velocity, proportionately, or more rapidly. The critical feature is that the two terms are not independent. Capital and physical capacity can accelerate technological development, while the resulting acceleration increases demand for the same resources and can strengthen the drag confronting the next unit of growth.

Forms of drag

Physical drag can take several forms. Singular drag arises when one resource becomes the binding constraint on further technological velocity; additional capacity elsewhere has limited effect until that bottleneck is relieved. Contributory drag occurs when several constraints independently impose limits without any one of them determining the outcome. Interactive drag is more complex: pressure on one resource changes the severity of another, so the combined effect cannot be understood by examining each constraint individually. A shortage of grid capacity, for example, may alter where infrastructure can be built, which in turn changes requirements for land, cooling, water or transmission. A contributory constraint can become singular as one bottleneck tightens, while several initially separate constraints can become interactive as expansion proceeds. The composition of physical drag can therefore change even when the underlying technological objective remains the same.

Capacity responds - and is consumed

Physical drag is not fixed. Higher prices and tighter constraints can induce additional capacity, substitution and efficiency, while artificial intelligence can itself improve how resources are produced, allocated and used. This can reduce b by increasing the amount of technological velocity that the physical system can support. Yet the mechanism is endogenous in both

directions. The AI capability used to improve resource efficiency itself creates demand on the resources of the physical system whose efficiency it is attempting to improve. The harder question is therefore whether resource capacity and efficiency can improve faster than the demand on the resources. If technological velocity rises faster than supporting capacity can be added, the absolute quantity of resources may continue increasing while becoming progressively less sufficient relative to demand. Rising scarcity can reinforce that effect through price: each unit of capital purchases fewer units of the more expensive constrained resource, slowing the rate at which physical capacity itself can expand. Efficiency and induced capacity may therefore reduce drag without eliminating it; the resulting headroom can enable further technological expansion, which consumes part of the capacity just released and can recreate the constraint at a higher level.

required to support it. The analytical capability required to allocate those resources carries its own capital and resource cost. The question then shifts from where resources should go to what kind of organization can economically sustain and exploit such a system.

III. Geography and the Moving Bottleneck

The coefficient of drag should therefore vary substantially across geography, depending on factors beyond the absolute quantity of endowed resources. What matters is the ability of a jurisdiction to expand the physical substrate as a system. Very few economies combine abundant energy, grid capacity, water, land, advanced fabrication, materials, construction capability, industrial supply chains and the institutional ability to bring new capacity online at speed. Some possess sophisticated technology and capital but face tight physical or infrastructural constraints; others can expand infrastructure rapidly while remaining dependent on critical technologies or inputs produced elsewhere. Strength in one resource can therefore coexist with severe drag elsewhere in the value chain. The most favorable environment for technological velocity is therefore the one capable of relieving successive constraints as they emerge. Geography consequently affects not only the level of b , but potentially the way drag itself scales as technological velocity increases.

The point of maximum pressure

The economic significance of physical drag lies not in scarcity as a general condition, but in identifying which resource is imposing the greatest constraint, where that constraint is emerging and when it is likely to bind. Those pressure points are not fixed. The sequence in which those constraints appear also matters. Investment committed to relieve one bottleneck can alter the value chain elsewhere before it becomes productive, meaning that similar starting conditions can lead to very different outcomes depending on which constraint emerges first and how quickly the system responds. Treating the physical substrate as a single aggregate therefore risks obscuring the point of maximum pressure within the resource system. The allocation problem is consequently dynamic: capital and productive capacity must be directed toward constraints whose importance is constantly changing. The critical question is whether those inflection points can be identified early enough for the relevant physical capacity to arrive before technological acceleration exposes the next debilitating constraint.

IV. Allocation Under Moving Constraints

If those pressure points move faster than conventional periodic analysis can track, the allocation problem increasingly requires an AI-assisted system capable of continuously identifying, comparing and ranking them by their effect on technological velocity. The cost of sustaining that velocity therefore extends beyond the original investment in artificial intelligence to the continuing expansion of the physical system