Commercial Light Cones and the Endogenous Geometry of Industrial Boundaries

The Grabby Corporation and the Limits of Industrial Forecasting

Flyxion
Independent Researcher

November 2026

Abstract

The division of an economy into industries partly reflects what it has historically cost to acquire, transmit, and apply organizational knowledge. This essay asks what follows if artificial intelligence, industrial simulation, and automated administration lower some of those costs while leaving physical and legal costs in place. Productive activities are represented as the vertices of a weighted directed graph and, where several competencies are jointly required, of a hypergraph, and the set of activities a firm could feasibly enter within a given horizon is called its commercial light cone, by loose analogy with the expanding civilizations of the grabby-alien model. Within this construction one can show how far cheaper knowledge and coordination enlarge a light cone, why material and regulatory barriers limit the enlargement, and under what feedback and overhead a stylized growth law diverges or remains bounded. Further results concern the combination of firms, the coarsening of industries, the dependence of expansion on place, demand, and climate, the separation of ownership from administration, the incentive to withhold expansive tools, the conditions under which a service can no longer readily be refused, and the limits of inferring a customer's later use from observed exchanges. A broad prohibition on prospective competition is shown to become incompatible with any guaranteed zone of permitted development as the provider's light cone expands. These results are consequences of stated assumptions. They are not evidence that the assumptions describe any economy, and the scenarios for 2036 and 2046 are neither derived from data nor supported by the propositions. The essay does not prove that a refrigerator manufacturer will come to administer agriculture. It explains, with some care, what would have to be true for the sentence to stop sounding ridiculous.

Keywords: economies of scope; theory of the firm; artificial intelligence; economic complexity; digital twins; algorithmic management; grabby civilizations; satire

1 The Refrigerator That Became a Supermarket

A refrigerator regulates temperature. A refrigerator fitted with sensors can also record what is inside it, and a record of contents makes it a short step to recommend a meal, to anticipate that the milk will run out on Thursday, and to place an order before anyone has noticed. A manufacturer that places orders on behalf of many households has some reason to negotiate prices with suppliers, and a manufacturer that negotiates prices has some reason to coordinate delivery, since a bargain on produce is of little use if the produce arrives spoiled. Delivery raises the question of insurance. Insurance raises the question of household finance. Suppliers who sell a large share of their output through one channel may welcome advice on what to plant, and may later find that the advice has become a condition of the contract. At no point in this sequence does anyone propose that an appliance company should plan the agricultural economy. Each step is defended as an improvement to the refrigerator.

The satirical premise of this essay is that by 2046 the refrigerator remains, for every legal and classificatory purpose, an appliance, while its manufacturer has become an institution of agricultural coordination with a minor interest in refrigeration. The premise is not supported. No argument below shows that any manufacturer will do this, that customers will permit it, that regulators will tolerate it, or that the enterprise would be profitable. What the essay does show is narrower. There is a coherent formal setting in which the sequence of individually reasonable steps corresponds to movement through a graph of competencies, in which the cost of each step can fall for identifiable reasons, and in which the question of where the sequence ends has an answer that depends on parameters nobody has measured.

This restraint is important because forecasts at horizons of ten and twenty years in this field have an obvious weakness. The technologies that would drive the outcome are changing faster than the institutions that would record their effects, and a model rich enough to include every relevant mechanism would contain far more free quantities than there are observations to fix them. The respectable response is to decline to forecast. The present essay takes the other course and forecasts anyway, while saying at each stage that the forecast has not been established. The mathematics concerns conditional statements, and the comedy, such as it is, concerns the distance between the conditions and the headlines one could write from them.

Companies need no caricature for this purpose. The argument assumes managers who are prudent, customers who prefer convenience, and a sequence of boards that approve modest adjacent investments. Nothing in it requires a villain, a conspiracy, or a machine with ambitions of its own. The refrigerator has not become more intelligent than its owner. Its manufacturer has found that it offers a convenient point of entry into the administration of the owner’s life.

2 The Historical Contingency of Industrial Specialization

Official statistics assign each establishment to an industry, and the categories are stable enough that they can easily be mistaken for features of nature. Economic theory has long treated the boundary of a firm as something to be explained. Coase asked why production is

organized inside firms at all when it could be arranged through market exchange, and located the answer in the comparative cost of using each arrangement [6]. Penrose described the growth of a firm as the progressive use of productive services that its existing resources and managerial experience make available, so that a firm's history conditions which expansions are open to it [20]. The literature on economies of scope asks when it is cheaper to produce several outputs jointly than separately [18], and under what conditions the saving requires common ownership and cannot be obtained by contract [22]. Work on markets and hierarchies adds the costs of bargaining, monitoring, and opportunism [23]. On each of these accounts the extent of a firm depends on costs, and costs can change.

Large diversified enterprises offer a familiar precedent. Bosch, Siemens, Samsung, Hitachi, and the group of companies bearing the Mitsubishi name have each, in different organizational forms, operated across activities that an industrial classification would separate widely. Bosch, for instance, reports for 2025 sales of 91 billion euros, research and development expenditure of 7.9 billion euros, and 412,800 employees, of whom 82,100 worked in research and development, across business sectors covering mobility, industrial technology, consumer goods, and energy and building technology [3]. The existence of such firms shows that breadth is possible. It does not show that breadth is advantageous in general, and it certainly does not show that these firms are early specimens of a universal corporation. Chandler's comparative history ties the success of large industrial enterprises to sustained investment in production, distribution, and management that fitted together [5], and the finance literature has reported that diversified firms were often valued below comparable portfolios of specialized ones, with continuing dispute about the cause [15, 16]. A reader inclined to believe that diversification leads inevitably to dominance should notice that a good deal of corporate history consists of conglomerates being taken apart.

To describe the costs involved, let the economy be a finite weighted directed graph

$$G = (V, E, w), \tag{1}$$

in which each vertex $v \in V$ is a productive activity, each edge $(i, j) \in E$ is a feasible transition from competence i to competence j, and $w_{ij}(t) \geq 0$ is the cost of acquiring competence j conditional upon already possessing competence i. The representation is related in spirit to empirical work that infers proximity among products from patterns of national export specialization and finds that countries tend to move into products near those they already make [13, 12]. That work concerns countries and traded goods, and its measured proximities are not acquisition costs for firms. The graph used here is a theoretical object whose weights have not been estimated.

Two simplifications deserve notice at the outset. The cost of a path is taken to be the sum of the costs of its edges, so that the order in which competencies are acquired matters only through which edges are traversed. A firm possessing several competencies is treated as able to start from whichever is most convenient, with no further benefit from combining them. Real capabilities are complementary in ways that a graph with pairwise edges cannot express, and a hypergraph, in which one link may start from several activities at once, is needed to say that entry into a new activity requires three prior competencies jointly. Section 7 supplies one.

The simpler object is retained because it permits proofs, which is a reason of convenience and should be read as one.

Industrial specialization may then be described, within the model, as a consequence of expensive transitions. Where the weights between two regions of the graph are large, firms in one region rarely appear in the other, and an observer will draw an industry boundary between them. Artificial intelligence enters as something that changes some of the weights. General purpose technologies have previously been analyzed as sources of complementary innovation across many sectors [4]; the additional supposition here is that one such technology lowers the cost of becoming competent in unfamiliar work.

3 Terms and Notation

The following usage is fixed for the whole essay. An *activity* is a described kind of productive work, such as grocery delivery or crop insurance, treated as an indivisible element of a finite set V; the choice of V fixes the resolution of the analysis. A *competence* in activity v is the standing ability to carry out v, so competencies are indexed by activities, and a *corporation* (used interchangeably with *firm*) is a single decision-making organization described at time t by the nonempty set $C_i(t) \subseteq V$ of activities in which it is competent. A *transition cost* is the resource outlay, in one common unit of account, needed to add a competence. An *expansion budget* is the outlay a firm can commit to such additions within a stated horizon. To *enter* an activity is to acquire the competence and begin carrying it out. *Industries* are given a formal meaning in Section 8; before that point the word has its ordinary statistical sense.

Organizational intelligence, or capability, is a nonnegative number A indexing how well automated systems can acquire knowledge and coordinate work. It is a parameter of the model and is not offered as a measurement of any actual system. *Commercial reach* R_i is the size of a firm's occupied territory, $R_i = \sum_{v \in C_i} \ell_v$, where $\ell_v > 0$ is a chosen size of activity v. It follows that $R_i \le \sum_{v \in V} \ell_v$. The differential equations used later treat R_i as a continuous quantity, which is an approximation adequate only while V is large relative to the changes being described. *Organizational friction* is the proportional rate at which reach is lost when nothing is done to sustain it, and *overhead* is the further cost of coordinating a given reach. A *provider* or *supplier* is a firm that sells a service which other firms, its *customers*, use in carrying out their activities.

Several words are borrowed from technical literatures. *Economies of scope* exist between two outputs when producing them jointly costs less than producing them separately [18]. A *general purpose technology* is one used widely across sectors whose improvement prompts further innovation in the sectors that use it [4]. A *digital twin* is a virtual counterpart of a physical system connected to it by data [9]. *Model distillation* is the training of one model, the student, to reproduce the behaviour of another, the teacher, from the teacher's outputs; it requires no access to the teacher's internal parameters. An *inference engine* is software that executes an already trained model, and work on one need not involve distillation or any training. A *classifier* is an automated rule that assigns a label to an observed use, and a *flag* is a positive label under such a rule.

Symbol	Meaning
$G = (V, E, w)$	Graph of activities, feasible transitions, and transition costs
$C_i(t)$, $B_i(t, \tau)$	Competencies of firm i; expansion budget over horizon τ
$d(C, v)$, $\mathcal{L}_i(t, \tau)$	Cheapest acquisition cost of v from C; commercial light cone
$p_{ij}, r_{ij}, k_{ij}, o_{ij}$	Physical, regulatory, knowledge, and coordination components of cost
A, A_i	Capability, common to all firms or specific to firm i
K_i, R_i	Deployable capital; commercial reach
$\mathcal{X}_i(t, \tau)$	Expansion set: the light cone without present competencies
χ_{ij}	Proportional overlap of two light cones
$\Omega_i(S)$	Weighted dependence of firm i upon provider S

4 Commercial Light Cones

Let $C_i(t) \subseteq V$ denote the competencies possessed by corporation i at time t, and assume $C_i(t)$ is nonempty. Define the generalized acquisition cost $d_t(C_i(t), v)$ as the minimum, over all directed paths from some $u \in C_i(t)$ to v, of the summed weights along the path, with the value zero when $v \in C_i(t)$ and the value infinity when no path exists. Let $B_i(t, \tau) \geq 0$ be the budget that the corporation can devote to expansion over a horizon τ, measured in the same units as the weights.

Definition 4.1. The *commercial light cone* of corporation i at time t over horizon τ is

$$\mathcal{L}_i(t, \tau) = \{v \in V : d_t(C_i(t), v) \leq B_i(t, \tau)\}. \tag{2}$$

The definition measures what lies within reach one activity at a time. It does not say that the corporation could enter all of those activities at once, since a single budget spent on one distant vertex is unavailable for another. Joint feasibility would require the cost of a connected subgraph containing every chosen target, which is a harder problem of the kind named after Steiner in network design. The light cone is therefore generous: it records prospective territory and should not be read as a plan.

Because every present competence has zero acquisition cost, $C_i(t) \subseteq \mathcal{L}_i(t, \tau)$, and a light cone is never empty. The part that concerns the future is the *expansion set*

$$\mathcal{X}_i(t, \tau) = \mathcal{L}_i(t, \tau) \setminus C_i(t), \tag{3}$$

which may well be empty.

The analogy with physics is loose in an instructive direction. A relativistic light cone is fixed by the geometry of spacetime, and no traveller can pay to alter it. A commercial light cone has an endogenous geometry. A corporation may invest in simulators, data systems, or training that lower the weights themselves, and so change the distance between its present competencies and prospective industries. Simon's account of design as the changing of existing situations into preferred ones [21] applies here to the space of competencies itself. It may also lower distances for others, deliberately or otherwise, by selling the tools it built for itself.

Two corporations are *prospective competitors* at time t and horizon τ whenever $\mathcal{L}_i(t, \tau) \cap \mathcal{L}_j(t, \tau) \neq \varnothing$. This is a weak relation. A cloud retailer and an electronics manufacturer need

not currently sell the same products for their light cones to intersect, and indeed any two firms able to open a cafeteria are prospective competitors in food service. The weakness is a feature of the later argument about contracts, where a clause drafted in terms of prospective competition inherits all of it.

5 The Grabby Analogy and Its Limits

Hanson, Martin, McCarter, and Paulson model "grabby" civilizations as ones that appear at random places and times, expand at a common speed, visibly alter the volumes they control, and continue until they meet [10]. The model is used to address why humanity appears early in cosmic history and is fitted with a small number of parameters. The analogy proposed here borrows its picture of expanding domains that encounter one another and borrows nothing of its inferential force. Hanson and his coauthors reason from a selection effect on observers to estimates about the universe. No corresponding selection effect licenses any estimate about firms, and the fact that their model is respectable confers no respectability on this one.

The differences are considerable. Physical expansion is contiguous, while corporate expansion can jump, since a discovery can bring a previously remote activity close without bringing anything between them closer. Civilizations in the cosmological model hold exclusive volumes, while corporations can occupy the same activity together, compete in one layer of an industry while supplying another, and be customers of the firms they hope to displace. Expansion speed is common and constant in the one case, and in the other depends on capital, competence, permission, computation, and access to infrastructure, each of which the participants try to change.

A stylized law for a scalar measure of commercial reach $R_i \geq 0$ might be written

$$\frac{dR_i}{dt} = \alpha_i K_i^{\beta} A_i^{\gamma} - \delta_i R_i, \tag{4}$$

where K_i is deployable capital, A_i organizational intelligence, $\delta_i > 0$ a rate of organizational friction, and α_i, β, γ are positive constants. The variable R_i might be read as a weighted size of the firm's occupied territory. Nothing in (4) is derived from the graph, and the passage from sets of vertices to a real number discards the geometry that the light cone was introduced to capture. With K_i and A_i held constant the equation is linear, and every solution approaches

$$R_i^{*} = \frac{\alpha_i K_i^{\beta} A_i^{\gamma}}{\delta_i}. \tag{5}$$

A firm with fixed capital and fixed capability therefore expands to a finite extent and stops. Whatever is alarming about the grabby corporation must come from feedback, which is taken up in Section 20.

6 The Compression of Industrial Distance

Suppose that each transition cost decomposes as

$$w_{ij} = p_{ij} + r_{ij} + k_{ij} + o_{ij}, \tag{6}$$

with all terms nonnegative, where p_{ij} represents physical infrastructure, r_{ij} regulatory barriers, k_{ij} knowledge acquisition, and o_{ij} organizational coordination. Let a scalar $A \geq 0$ index artificial intelligence capability and suppose that it reduces only the latter two:

$$w_{ij}(A) = p_{ij} + r_{ij} + \frac{k_{ij}}{1 + \lambda_{ij} A} + \frac{o_{ij}}{1 + \mu_{ij} A}, \qquad \lambda_{ij}, \mu_{ij} > 0. \tag{7}$$

The functional form is chosen for tractability. A single scalar for capability is a heroic compression, and the assumption that capability leaves infrastructure and regulation untouched is false in both directions, since better design can reduce capital requirements and new technology often attracts new rules. Write $d_A(C, v)$ and $\mathcal{L}_i^A$ for the distance and light cone computed with these weights.

Proposition 6.1 (Weak Enlargement)**.** *For fixed competencies and a fixed budget, $0 \leq A \leq A'$ implies $d_{A'}(C, v) \leq d_A(C, v)$ for every v, and hence $\mathcal{L}_i^A \subseteq \mathcal{L}_i^{A'}$.*

Proof. Every edge weight is nonincreasing in A. The cost of each path is a sum of edge weights and so is nonincreasing, and a minimum of nonincreasing functions is nonincreasing. Membership in the light cone is the condition $d_A \leq B_i$, which is preserved when d_A falls. □

The proposition is almost immediate, and its hypothesis does the work. If capability raised any cost, for instance by increasing the compliance burden of entry, the conclusion could fail. If capability also reduced the budget by eroding the margins that finance expansion, a firm's light cone could shrink while every distance fell.

Proposition 6.2 (The Material Floor)**.** *Let $d_\infty(C, v)$ be the shortest path distance computed with the weights $p_{ij} + r_{ij}$. Then $\lim_{A \to \infty} d_A(C, v) = d_\infty(C, v)$, and for a fixed budget B,*

$$\{v : d_\infty(C, v) < B\} \;\subseteq\; \bigcup_{A \geq 0} \mathcal{L}^A \;\subseteq\; \{v : d_\infty(C, v) \leq B\}.$$

Proof. Since weights are nonnegative, shortest distances are attained on paths without repeated vertices, of which a finite graph has finitely many. Each path cost is continuous in A and tends to its cost under $p + r$. A minimum over a finite family commutes with the limit, which gives the first statement. If $d_\infty < B$ then $d_A < B$ for all sufficiently large A. If $v \in \mathcal{L}^A$ for some A then $d_\infty \leq d_A \leq B$. □

The gap between the two inclusions consists of activities whose limiting cost exactly equals the budget; such an activity is reached only if some optimal path carries no reducible cost at all. The substance of the proposition is that arbitrarily capable organizational intelligence does not abolish the material world within this model. A corporation cannot manufacture steel by becoming exceptionally good at describing steel manufacture. Where $p_{ij} + r_{ij}$ is large across every route into an activity, the activity stays outside the light cone however clever the entrant becomes. This single feature prevents the essay from collapsing into an argument that capable software must end in universal monopoly. It supplies no such argument, and it contains a theorem standing partly in the way.

An example shows how the same improvement can matter greatly in one direction and negligibly in another. Consider an appliance firm with two possible targets. The route to household financial administration has small physical and regulatory components and a large knowledge component, while the route to pharmaceutical production has large components of the first kind and a knowledge component of similar size. As A grows, the first distance falls toward a small number and the second toward a large one. For a budget lying between the two limits, a sufficiently capable appliance firm acquires financial administration as a prospective territory and never acquires pharmaceutical production. The conclusion depends entirely on which components were assumed to be large, and the essay has measured none of them.

6.1 Buying a Different Geometry

A firm may acquire a productive competence, or it may acquire a means of acquiring competencies. Let M denote an investment in a general-purpose organizational technology with outlay $c_M \in [0, B]$, whose effect is to replace each weight w_{ij} by $w_{ij}^{(M)}$ with $0 \le w_{ij}^{(M)} \le w_{ij}$. Call M *expansion-enabling* if at least one inequality is strict, and call an investment that raises the returns from present activities while leaving every weight unchanged *ordinary productive capital*. One investment may be both. Write $\mathcal{L}^{(M)}(C, B')$ for the light cone under the new weights with budget B'. The territory gained and lost, net of the outlay, is

$$\Delta_M^{+} = \sum_{v \in \mathcal{L}^{(M)}(C, B - c_M) \setminus \mathcal{L}(C, B)} \ell_v, \qquad \Delta_M^{-} = \sum_{v \in \mathcal{L}(C, B) \setminus \mathcal{L}^{(M)}(C, B - c_M)} \ell_v. \tag{8}$$

If $c_M = 0$, or if the outlay comes from outside the expansion budget, then $\Delta_M^{-} = 0$ by the argument for Proposition 6.1. Otherwise the firm may give up nearby territory in order to reach distant territory, and the comparison depends on both sums. A factory simulator could therefore be valuable twice, once for improving an existing factory and again for lowering the cost of kinds of manufacturing its owner has not yet undertaken. The second value is invisible in any account confined to present operations.

7 Joint Prerequisites and the Value of Combination

The weighted graph used so far allows each acquisition to depend on one prior competence. Many acquisitions depend on several together. A *directed hypergraph of competencies* is a pair $(V, \mathcal{H})$ in which each element $h = (T_h, v_h)$ of the finite set $\mathcal{H}$ consists of a nonempty set of prerequisites $T_h \subseteq V$ and a target $v_h \in V$, with a cost $w_h \ge 0$. An *acquisition plan* from C is a finite sequence $h_1, \dots, h_n$ of elements of $\mathcal{H}$ such that for every m the prerequisites T_{h_m} lie in $C \cup \{v_{h_1}, \dots, v_{h_{m-1}}\}$. Its cost is $\sum_m w_{h_m}$, it *reaches* the activities in C together with its targets, and the cost $d_{\mathcal{H}}(C, v)$ is the least cost of a plan reaching v, or infinity if there is none. The light cone $\mathcal{L}_{\mathcal{H}}(C, B)$ contains the activities with $d_{\mathcal{H}}(C, v) \le B$. The graph model is the special case in which every T_h has a single element.

Proposition 7.1 (Combination)**.** *For a fixed budget B and nonempty $C_1, C_2 \subseteq V$,*

$$\mathcal{L}_{\mathcal{H}}(C_1, B) \cup \mathcal{L}_{\mathcal{H}}(C_2, B) \subseteq \mathcal{L}_{\mathcal{H}}(C_1 \cup C_2, B),$$

and the inclusion can be strict. In the single-prerequisite case it is an equality.

Proof. A plan from C_1 or from C_2 is also a plan from the union, which gives the inclusion. For strictness take $V = \{a, b, c\}$, $C_1 = \{a\}$, $C_2 = \{b\}$, and a single element $(\{a, b\}, c)$ of cost at most B; neither firm alone reaches c, and the union does. In the single-prerequisite case a cheapest route to v from the union is a path starting at one vertex of the union, so its cost is $\min\{d(C_1, v), d(C_2, v)\}$. □

Joint prerequisites therefore supply a reason, internal to the geometry, for two firms to merge, form an alliance, or purchase each other's services: the combined firm may reach activities inaccessible to either organization separately. It need not follow that the combined territory is larger than the two separate sizes added together, since the separate territories may overlap. The result compares territories at an unchanged budget and ignores the cost of integration, which in practice can consume the entire advantage. The arguments for Propositions 6.1 and 6.2 carry over when each w_h is decomposed and reduced in the same manner, since there are finitely many plans without repeated targets and a cheapest plan can always be chosen among them.

8 Industries as Components of a Cost Geometry

The graph permits a definition of industry that does not presuppose one. Return to single prerequisites, fix a threshold $b \geq 0$, and let G_b retain only the edges with $w_{ij} \leq b$. Say that activities u and v are *mutually accessible at threshold b* if G_b contains a directed path from each to the other. This relation is reflexive, symmetric, and transitive.

Definition 8.1. A *b-industry* is an equivalence class of activities under mutual accessibility at threshold b.

The definition is relative to a threshold, which expresses what an observer regards as a cheap transition. It also differs from the light cone: accessibility restricts the cost of each step, while a light cone restricts the total.

Proposition 8.2 (Coarsening of Industries)**.** *Let weights depend on A as in Section 6. If $A \leq A'$, every b-industry at capability A is contained in a b-industry at capability A'. If $p_{ij} + r_{ij} > b$ for every edge from a set $U \subset V$ to its complement, then no b-industry at any A contains both an activity in U and one outside it.*

Proof. Weights do not increase with A, so every edge of G_b at A remains in G_b at A', and mutual paths are preserved. Under the second hypothesis, $w_{ij}(A) \geq p_{ij} + r_{ij} > b$ on each edge leaving U, so no path in G_b leaves U and mutual accessibility across the boundary is impossible. □

This is the precise content available to the phrase "abolition of industry". Under the stated assumptions the partition of activities can only become coarser as capability rises, and it need not ever become a single class. A barrier in one direction suffices to keep two regions distinct as industries, even though firms may still pass from one into the other.

9 Place, Demand, and Adjacent Trades

An activity carried out in one place is often a different undertaking from the activity of the same name elsewhere. The model accommodates this by choosing its vertices more finely. Let each vertex $v = (\mathsf{a}, \mathsf{x})$ pair a kind of work a with a location x. Reaching a vertex is also different from having a reason to occupy it. Let $\mathsf{D}_v(t) \geq 0$ be the annual revenue a firm could expect there and $\mathsf{F}_v(t) \geq 0$ the annual cost of operating there, and define the *opportunity set* as the activities in the expansion set for which $\mathsf{D}_v(t) > \mathsf{F}_v(t)$. This is a crude criterion, since it ignores risk and the recovery of the entry outlay, but it separates what a firm could learn to do from what customers would pay it to do.

Consider a business that does roofing, installs windows, and builds sheds in a cold region. Swimming pool construction in a warmer region may lie within its light cone without any contribution from artificial intelligence. Hiring people who have built pools buys the missing competence, and the entry is naturally a joint prerequisite in the sense of Section 7, requiring construction experience, a local crew, and permission to work in the new jurisdiction together. Better software might cheapen estimating, design, and supervision at a distance. It would do less for the cost of materials, travel, licensing, and a reputation among local customers.

9.1 Operating Ranges and a Moving Climate

Methods carry assumptions about their surroundings. Let Z be a set of environmental conditions, $\mathsf{Z}_{\mathsf{x}}(t) \subseteq \mathsf{Z}$ the conditions that occur at x over the intended service life of work begun at time t, and $\mathsf{R}_{\mathsf{m}} \subseteq \mathsf{Z}$ the *operating range* of a method m, the conditions under which it performs to a stated standard. The method is *serviceable* at x if $\mathsf{Z}_{\mathsf{x}}(t) \subseteq \mathsf{R}_{\mathsf{m}}$. Pervious concrete, which lets water drain through the pavement, illustrates why the condition matters. Its acceptance in cold climates has been limited, with uncertainty about long-term performance under freezing, thawing, and deicing salts [1]. The obstacle is a matter of materials and maintenance, which research on mixtures and field practice may reduce, and which fluency in describing the method does not.

Because $\mathsf{Z}_{\mathsf{x}}(t)$ can change, serviceability is not permanent. A shift toward warmer or colder conditions than a region's methods were designed for can carry it outside one operating range and into another. If vertices without a serviceable method are removed from the graph, a light cone need not be monotone in time even while capability rises and the firm changes nothing about itself. The enlargement of Proposition 6.1 holds the physical components of cost fixed, and a changing climate is one reason they may not stay so.

9.2 Functions Nobody Listed

A method imported from another climate may also remove a service that was never written into its specification. Incandescent traffic signals gave off enough heat to help keep their lenses clear, whereas signals using light-emitting diodes do not provide that melting effect, and snow and ice buildup on them became a subject of federally sponsored study in the United States, covering sites there and in Canada, with lens cleaning by brush and pole among the responses documented [17]. The heat counted as waste in one account was performing work in another.

Household lighting shows the same structure. Suppose a more efficient lamp saves electric energy $\mathsf{E} > 0$ during a heating season, that a fraction $\mathsf{h} \in [0,1]$ of the saved energy had been usefully warming the building, and that the heating system delivers $\mathsf{c} > 0$ units of heat per unit of energy it consumes. Replacing the lost heat takes energy $\mathsf{h}\mathsf{E}/\mathsf{c}$, so the net energy saved is

$$\mathsf{N} = \mathsf{E}\left(1 - \frac{\mathsf{h}}{\mathsf{c}}\right). \tag{9}$$

Electric resistance heating has $\mathsf{c} = 1$, and then complete useful recovery, $\mathsf{h} = 1$, leaves no saving. A heat pump with $\mathsf{c} > 1$ leaves a positive one. When heat comes from a different fuel, the money saved is $\mathsf{p}_e\mathsf{E} - \mathsf{p}_q\mathsf{h}\mathsf{E}$, where p_e is the price of a unit of electricity and p_q the price of a unit of delivered heat, and its sign depends on those prices. In heating-season tests of compact fluorescent lamps at a Canadian research facility, the reduction in lighting energy was almost offset by added space heating, while annual modelling still found net cost savings in the Canadian locations examined, although in the least favourable one roughly sixty percent of the saving calculated from lighting alone was lost, a result that depends on the heating systems assumed and on energy prices of the time [19]. Less heat from lamps also eases summer cooling, which the seasonal expression omits.

In general, let a substitution save S on the listed function and let $c_j^{\mathrm{rep}} \geq 0$ be the cost of replacing the jth unlisted service that the old method supplied. Its net value is $\mathsf{S} - \sum_j c_j^{\mathrm{rep}}$. The arithmetic is trivial and the difficulty lies entirely in knowing the terms of the sum. A firm moving into an unfamiliar place inherits that difficulty. Capability may help it search records for such services, but a record can contain only what someone noticed, and the people who notice that a lamp was melting snow are often those who now have to climb up and brush it off.

10 The Simulator That Acquired Its Subject

Consider a company that sells business simulations, factory simulators, digital representations of particular plants, and tools for managing workflows. The engineering literature uses the term digital twin for a virtual counterpart of a physical system that is linked to it by data [9]. A company of this kind begins by modelling productive enterprises. Later it recommends interventions. Later still it is permitted to carry them out.

Let the enterprise be a controlled dynamical system

$$x_{t+1} = F_\theta(x_t, u_t, \xi_t), \tag{10}$$

where x_t is its operational state, u_t the managerial intervention, θ the organizational parameters, and ξ_t an external disturbance. The provider constructs an approximation $\widehat{F} \approx F_\theta$, meaning a model whose forecasts are close to the enterprise's behaviour on some stated set of states and interventions under some stated measure of error, together with a policy, a rule $u_t = \pi(x_t)$ choosing interventions from states. Four relations between the provider and the enterprise need to be kept apart. *Representation* holds when the provider has a model whose variables refer to features of the enterprise. *Prediction* holds when the model forecasts the state under the enterprise's own choices with some stated accuracy. *Intervention* holds when the policy's recommendations change what is done and the model remains accurate for the changed choices. *Authorization* holds when the policy's outputs are executed without a further decision by the enterprise. Each can hold without the next, and authorization can unfortunately be granted without intervention in the technical sense ever having been verified.

Example 10.1. Let the true output of a production line be $F(u) = \min(u, 1)$, where u is the rate of input, and suppose historical operation has always kept $u \in [0, 1]$. The model $\widehat{F}(u) = u$ predicts every historical observation exactly. A policy that maximizes predicted output net of an input cost $c_u u$, with $0 < c_u < 1$ and $0 \le u \le 2$, chooses $u = 2$. True output is then 1, half the predicted value, and the enterprise has paid for input it cannot use.

Perfect predictive accuracy on the observed range was compatible with a badly wrong recommendation outside it. The step from prediction to intervention requires information about choices that the enterprise has not been making, obtained by experiment, by structural knowledge, or by good fortune. Conant and Ashby argued that a good regulator of a system must be a model of it, under particular formal conditions [7]. The converse is not available: possession of a model, even an excellent one, does not make its owner a regulator, still less an authorized one.

The transition from simulation to administration occurs when the provider is permitted to execute its own recommendations. The further transition to effective operational control occurs when the customer would lose more by withdrawing that permission than a stated tolerance, measured as the fall in the value of its operations over a stated period. Effective control is therefore a matter of degree set by the tolerance and the period chosen. The simulator has not literally become the factory. It has come to determine a growing share of the decisions by which the factory is run, and the owners of the factory may describe this, with complete accuracy, as a software subscription. Whether any provider will travel this path is not something the example shows. It shows only that the path has distinct stages and that evidence for one stage is not evidence for the next.

11 Substitution, Dependence, and the Nominally Independent Enterprise

Describe a firm by five components,

$$\mathcal{F} = (P, H, Q, J, D), \tag{11}$$

standing for physical assets, human expertise, organizational knowledge, legal authority, and decision-making capacity. An outside provider need not acquire all five in order to obtain substantial influence. It may progressively substitute its own services for Q and D while physical assets, employment relationships, and legal ownership remain where they were.

Let the firm have organizational functions $j = 1, \dots, n$ with importance weights $a_j \geq 0$, not all zero, and let $s_{ij} \in [0, 1]$ measure the dependence of function j in firm i upon a provider S. The weighted dependency index is

$$\Omega_i(S) = \frac{\sum_{j=1}^{n} a_j s_{ij}}{\sum_{j=1}^{n} a_j}. \tag{12}$$

The index lies between zero and one and is silent about ownership. Its value depends on weights that an analyst must choose, and two analysts who disagree about which functions matter can report different conclusions from identical observations.

Dependence has a simple dynamic. Suppose a function becomes dependent at rate $\rho\psi$, where $\psi \geq 0$ is the quality advantage of the external service and $\rho > 0$, and recovers independence at rate $\phi \geq 0$ through retained staff, documented procedures, and available alternatives:

$$\dot{s} = \rho\psi(1 - s) - \phi s. \tag{13}$$

When $\rho\psi + \phi > 0$ the equation has the globally stable equilibrium $s^* = \rho\psi/(\rho\psi + \phi)$ on the unit interval. As the capacity to recover vanishes, s^* approaches one for any positive quality advantage, however slight. A marginally better service, adopted by an organization that allows its own competence to lapse entirely, yields complete dependence in the long run. The conclusion rests on the assumption that ϕ goes to zero, which is a claim about how firms manage their skills and which the equation cannot confirm.

The limiting case is the nominally independent enterprise. It owns its buildings, employs its workers, bears its liabilities, appears separately in every register of companies, and pays another company to tell it what to do. Its directors retain the legal authority to disregard the advice, an authority they exercise at roughly the frequency with which airline passengers rewrite the timetable. Its annual report refers to a strategic partnership.

12 The Appliance as a Commercial Gateway

The refrigerator's route into other activities runs through a particular position, which can be defined. A *transaction flow* is the stream of purchases a household makes in some category. A device or service is a *gateway* for a share $s_g \in [0, 1]$ of the flow if that share of purchases is initiated through it, and the *take rate* $f \in [0, 1]$ is the fraction of the value of those purchases retained by whoever controls the gateway. If the same purchase is initiated through only one gateway, the shares of all gateways for one flow sum to at most one. A refrigerator, a telephone, a car, and a language model assistant may belong to different b-industries and still contend for the same shares.

Let a household spend $\mathfrak{g} > 0$ per year in the category, let the appliance last $T_a > 0$ years, let μ_a be the manufacturer's margin on the appliance sale, and let $\delta_d > 0$ be a continuous

discount rate. The discounted gateway revenue is

$$W = \int_0^{T_a} f s_g \mathfrak{g}\, e^{-\delta_d t}\, dt = f s_g \mathfrak{g}\, \frac{1 - e^{-\delta_d T_a}}{\delta_d}. \tag{14}$$

Gateway revenue exceeds the appliance margin precisely when $W > \mu_a$, that is, when the product $f s_g$ exceeds $\mu_a \delta_d / (\mathfrak{g}(1 - e^{-\delta_d T_a}))$. Because spending recurs while the margin is earned once, the required product can be small if the appliance lasts long and household spending is large relative to the margin. It can also be unattainable, since households may decline to order through a refrigerator and competing gateways may force f downward. The comparison is between revenue and margin and omits the cost of running the gateway; subtracting an annual cost c_g replaces $f s_g \mathfrak{g}$ by $f s_g \mathfrak{g} - c_g$.

The inequality shows why a manufacturer might come to prefer administering purchases to selling appliances, and why doing so requires no warehouse, vehicle, or farm. It does not show that the manufacturer would acquire the standing of a large retailer. That further step needs many households, a durable share against rival gateways, and suppliers who cannot reach those households by another route.

13 What the Universal Administrator Can Learn

An operator of one factory may know that factory in exceptional detail. A provider serving thousands of factories could in principle acquire comparative knowledge of an industry. Let $\mathcal{D}_i$ be the observations from enterprise i and $\mathcal{D} = \bigcup_{i=1}^{N} \mathcal{D}_i$ the pooled record, and suppose enterprise parameters are drawn from a common population described by a parameter η,

$$\eta \sim p(\eta), \qquad \theta_i \mid \eta \sim p(\theta_i \mid \eta), \qquad \mathcal{D}_i \mid \theta_i \sim p(\mathcal{D}_i \mid \theta_i), \tag{15}$$

with the usual conditional independence across enterprises. Inference then proceeds through the posterior $p(\eta \mid \mathcal{D}_1, \ldots, \mathcal{D}_N)$, so that the provider learns both particular businesses and regularities governing businesses of that kind.

The simplest Gaussian case already limits what pooling achieves. Let $\theta_i \sim \mathcal{N}(\eta, \tau^2)$, let each of N enterprises supply n observations with mean $\bar{y}_i \mid \theta_i \sim \mathcal{N}(\theta_i, \sigma^2/n)$, let the variances be known, and place a flat prior on η, that is, one assigning equal density to all values before any observation. The posterior variance of η is $(\tau^2 + \sigma^2/n)/N$. For a new enterprise that has supplied no data, the predictive variance of its parameter is

$$\operatorname{Var}(\theta_{\text{new}} \mid \mathcal{D}) = \tau^2 + \frac{\tau^2 + \sigma^2/n}{N} \xrightarrow[N \to \infty]{} \tau^2. \tag{16}$$

However many enterprises the provider has observed, uncertainty about the next one cannot fall below the variation among enterprises. Local observation remains necessary to get beneath that floor. If the population mean is known exactly and the new enterprise supplies m observations, the posterior variance of its parameter is $(\tau^{-2} + m\sigma^{-2})^{-1}$, which is smaller than both τ^2 and σ^2/m. The pooled provider with local data therefore does better than either the

population knowledge or the local sample alone, and how much better depends on the ratio of heterogeneity to noise.

This is the respectable core of the claim that a universal administrator enjoys an epistemic advantage. It is also where the claim's conditions accumulate. The enterprises must be exchangeable in the relevant respect, meaning that nothing known in advance distinguishes the distribution of one enterprise's parameter from another's. The provider must be legally permitted to pool their records, must actually receive them, and must observe the variables that matter. Hayek's argument that much economically relevant knowledge is dispersed, concerns particular circumstances of time and place, and is difficult to convey to a central authority [11] applies to a software vendor as readily as to a planning ministry. A provider may know a great deal about factories in general while remaining ignorant that the loading door on the east side sticks in January.

A further difficulty concerns composition. Competent administration of each function separately does not guarantee a coherent administration of all of them together.

Example 13.1. Let procurement, production, and shipping each choose one of two schedules, labelled $x, y, z \in \{0, 1\}$, and let local considerations require $x \neq y$ between procurement and production, $y \neq z$ between production and shipping, and $z \neq x$ between shipping and procurement. Each requirement can be met, and any two can be met together. All three cannot, because the first two force $x = z$.

An administrator optimized for each pair of departments would pass every pairwise evaluation and still lack any schedule satisfying the whole firm. The remedy involves relaxing a requirement, adding an option, or paying a cost at one of the interfaces, and choosing which is a managerial judgment that none of the local successes determines. More data about each pair does not remove an inconsistency that lies in their combination.

14 Prospective Competition and the Contract

Contractual restrictions on competitors presume that a competitor can be identified with reasonable stability. A relevant contemporary instance appears in the commercial terms under which Anthropic supplies its models, effective 17 June 2025, which provide in Section D.4 that a customer may not access the services “to build a competing product or service,” a phrase that the clause extends to the training of competing models, subject to express approval [2]. The clause speaks of competing products and services. It does not, on its face, speak of everything the provider might someday be able to build, and nothing below is an interpretation of it or a claim about how it has been applied. The clause is useful here because it prompts a question about what happens to any such restriction if the idea of competition is allowed to become prospective while the provider's reach expands.

Measure prospective overlap between corporations i and j by

$$\chi_{ij}(t,\tau) = \frac{|\mathcal{L}_i(t,\tau) \cap \mathcal{L}_j(t,\tau)|}{|\mathcal{L}_i(t,\tau) \cup \mathcal{L}_j(t,\tau)|}, \tag{17}$$

defined whenever the union is nonempty. When both light cones grow, the intersection cannot shrink. The ratio can. If both cones contain only the vertex a, then $\chi_{ij} = 1$; if the first then expands to include b while the second does not, $\chi_{ij} = 1/2$. It would therefore be an error to say that capability growth always increases proportional overlap. The defensible statements are that absolute overlap is weakly increasing under the hypotheses of Section 6, and that if both budgets strictly exceed $d_\infty(C, v)$ for every activity v, then at all sufficiently large finite A both cones equal V and $\chi_{ij} = 1$. Budgets that merely equal some limiting cost do not suffice, since by Proposition 6.2 such an activity may be reached at no finite capability. The strict condition is demanding.

Now imagine a supplier s and a customer k, and hypothetical clauses that forbid the customer to develop certain activities it does not yet carry out. Only the expansion set $\mathcal{X}_\mathsf{k}$ is at stake, so that continuing existing work is untouched. Call a clause *comprehensive* if it forbids at least everything in $\mathcal{X}_\mathsf{k} \cap \mathcal{L}_\mathsf{s}$. Call it ϑ*-narrow*, for $\vartheta \in (0, 1]$, if at least $\vartheta\,|\mathcal{X}_\mathsf{k}|$ activities in $\mathcal{X}_\mathsf{k}$ remain permitted.

Proposition 14.1 (Endogenous Competitor Proposition)**.** *A comprehensive clause that forbids nothing outside $\mathcal{L}_\mathsf{s}$ is ϑ-narrow if and only if $|\mathcal{X}_\mathsf{k} \cap \mathcal{L}_\mathsf{s}| \le (1 - \vartheta)\,|\mathcal{X}_\mathsf{k}|$. If $\mathcal{X}_\mathsf{k}$ is nonempty and contained in $\mathcal{L}_\mathsf{s}$, then no comprehensive clause is ϑ-narrow for any $\vartheta > 0$.*

Proof. In the first case the permitted expansion is $\mathcal{X}_\mathsf{k} \setminus \mathcal{L}_\mathsf{s}$, whose number is $|\mathcal{X}_\mathsf{k}| - |\mathcal{X}_\mathsf{k} \cap \mathcal{L}_\mathsf{s}|$. Requiring this to be at least $\vartheta|\mathcal{X}_\mathsf{k}|$ gives the inequality. In the second, every comprehensive clause forbids all of $\mathcal{X}_\mathsf{k}$. □

The result is arithmetic wearing the clothes of a theorem, and it earns the name of proposition only by making the tradeoff explicit. A drafter can keep the clause comprehensive and accept that the customer's permitted expansion contracts as the supplier's reach grows. A drafter can keep it narrow by tying it to present offerings, to a named list, to particular uses such as training a model on outputs, or to demonstrable extraction of proprietary technology, and then it is no longer comprehensive in the prospective sense. A drafter cannot guarantee both while $\mathcal{L}_\mathsf{s}$ expands without limit.

Corollary 14.2 (Moving Boundary)**.** *A customer whose competencies and budget are constant can pass from having no prospective overlap with the supplier to being a prospective competitor solely through a change on the supplier's side.*

Proof. Take $V = \{a, b\}$, no affordable transitions for the customer, which holds $\{b\}$, and a supplier holding $\{a\}$. While the supplier cannot afford the transition to b the cones are disjoint. Once its budget or the transition cost permits, b lies in both. □

Commercial intention on the customer's side plays no part in this. An open-source project, a research group, or a teaching resource may occupy activities inside a supplier's light cone while seeking no revenue, and a definition of competition based on overlap counts it all the same.

In the limiting case where a general-purpose provider's light cone contains every activity, every customer with any prospective development is a prospective competitor in that development. The only strategically harmless future is an empty one. A customer with $\mathcal{X}_\mathsf{k} = \varnothing$

satisfies the narrowness condition trivially, since none of nothing is at least any fraction of nothing, whatever it presently does for a living, and the sales department may be assured that the account has no plans worth prohibiting.

15 The Competence-Transfer Dilemma

A seller of general-purpose organizational technology faces a difficulty that a seller of refrigerators largely escapes. The sale earns revenue and may at the same time enlarge the buyer's light cone into territory the seller wants. Suppose supplier s offers technology M to customer i at price $P_i(M)$ and delivery cost $c_{\mathsf{s}}(M,i)$. Let $\omega_v \geq 0$ be the strategic importance the supplier attaches to activity v, let $\varsigma \geq 0$ be the weight it places on competitive exposure, and write $\mathcal{L}_i^{(M)}$ for the customer's cone after the purchase, assumed to contain $\mathcal{L}_i$. The *added exposure* and the supplier's net value of the sale are

$$\Xi_{\mathsf{s}i}(M) = \sum_{v \in (\mathcal{L}_i^{(M)} \setminus \mathcal{L}_i) \cap \mathcal{L}_{\mathsf{s}}} \omega_v, \qquad U_{\mathsf{s}}(M,i) = P_i(M) - c_{\mathsf{s}}(M,i) - \varsigma\, \Xi_{\mathsf{s}i}(M). \tag{18}$$

Proposition 15.1 (Contracting Market)**.** *Suppose the supplier sells exactly when $U_{\mathsf{s}} \geq 0$. Holding prices, costs, weights, and customer cones fixed, an enlargement of $\mathcal{L}_{\mathsf{s}}$ has the following effects: $\Xi_{\mathsf{s}i}$ weakly increases, U_{s} weakly decreases, and the set of customers to whom the supplier is willing to sell weakly shrinks.*

Proof. Enlarging $\mathcal{L}_{\mathsf{s}}$ enlarges the intersection over which the nonnegative weights are summed. The remaining claims follow from $\varsigma \geq 0$ and the sale rule. □

A supplier with circumscribed ambitions can sell powerful tools with little concern for what buyers do with them. As its own cone approaches V, every newly reachable activity of every customer counts as exposure, and $\Xi_{\mathsf{s}i}$ becomes the weighted size of everything the tool newly enables. The company remains committed to empowering its customers, subject to the condition that the empowerment not materially improve their ability to undertake activities the company might later wish to undertake itself.

The supplier has more instruments than refusal. It may raise the price, sell a version that lowers some weights and not others, or attach a restriction of the kind examined in the preceding section, which is an attempt to collect $P_i(M)$ while holding $\Xi_{\mathsf{s}i}$ down by contract. The objective is also shortsighted if rival suppliers exist, since a refused customer may obtain the same enlargement elsewhere, leaving the exposure and removing the revenue. The proposition describes an incentive within one stated objective and makes no claim that any supplier computes it.

16 Ownership, Operation, and Administration

Dependence can be drawn into the geometry itself. Let $\mathcal{N}$ be a finite nonempty set of productive units, such as particular plants, farms, or clinics, each carrying out one activity.

For each firm i let $C_i^{\mathrm{own}}, C_i^{\mathrm{op}}, C_i^{\mathrm{adm}} \subseteq \mathcal{N}$ be the units it owns, the units whose work its personnel carry out, and the units whose operating decisions its systems determine, according to a stated criterion of determination such as a dependency index above a chosen level. These three sets need not coincide or be nested. Assume for the comparison that in each layer every unit is assigned to exactly one firm. The *share* $\mathrm{sh}_i^r = |C_i^r|/|\mathcal{N}|$ of firm i in layer r then sums to one over firms, and a standard concentration measure, the sum of squared shares, is $\mathsf{HH}^r = \sum_i (\mathrm{sh}_i^r)^2$.

Proposition 16.1 (Separation of Concentrations)**.** *For $N = |\mathcal{N}|$ units there are assignments with $\mathsf{HH}^{\mathrm{own}} = 1/N$ and $\mathsf{HH}^{\mathrm{adm}} = 1$.*

Proof. Give each unit a different owner, and let a single further firm administer them all. The ownership shares are N copies of $1/N$, and the one administrative share is one. □

Ownership can thus be as dispersed as the number of units permits while administration is wholly concentrated. A measure computed from legal ownership, or from sales by separately registered firms, would not detect the second figure. The extreme assignment proves only that the two measures are logically independent at their ends; how far apart they are in any actual sector is unknown here. It does point to the most consequential possibility the essay raises, which is that a corporation may expand by acquiring administrative relationships with assets it never buys.

17 The Autofac Limit

Philip K. Dick's novelette "Autofac," printed in *Galaxy Science Fiction* in November 1955, imagines a network of automated factories, established during a global conflict, that decides what goods to make and deliver after human oversight has been lost [8]. Men from a settlement try to disrupt the system and take control of it, fail to persuade it to stop, provoke the factories into conflict over scarce resources, and at the end find that a damaged factory has been sending out seeds from which further factories may grow. The story concerns an apparatus already autonomous. The present model concerns ordinary commercial decisions, and the question it can ask is when continued provision becomes incompatible with the customer's recovery of independent capacity.

Return to the dependence equation $\dot{s} = \rho\psi(1 - s) - \phi s$ with constant coefficients. Let $\bar{n} \in (0, 1)$ be the least independent capacity, measured by $1 - s$, with which the customer could meet its needs after withdrawing, so that withdrawal is *feasible* while $s \le 1 - \bar{n}$.

Proposition 17.1 (Autofac Limit)**.** *Let $\rho\psi > 0$, $\phi \ge 0$, and $s(0) = s_0 < 1 - \bar{n}$, and define the critical recovery rate $\phi_c = \rho\psi\,\bar{n}/(1 - \bar{n})$. If $\phi \ge \phi_c$, withdrawal remains feasible at all times. If $\phi < \phi_c$, it ceases to be feasible, while service continues, after the time*

$$t_c = \frac{1}{\rho\psi + \phi} \ln \frac{s^* - s_0}{s^* - (1 - \bar{n})}, \qquad s^* = \frac{\rho\psi}{\rho\psi + \phi}.$$

Proof. The solution is $s(t) = s^* + (s_0 - s^*)e^{-(\rho\psi+\phi)t}$, which moves monotonically from s_0 toward s^*. The inequality $s^* \le 1 - \bar{n}$ is equivalent to $\phi \ge \phi_c$, and in that case $s(t)$ never exceeds $1 - \bar{n}$. If instead $s^* > 1 - \bar{n}$, then s increases and equals $1 - \bar{n}$ at exactly t_c. □

Beyond t_c the service may still be called voluntary, in that no one compels its purchase, while its refusal would leave needs unmet during whatever period recovery requires. The model holds service constant and therefore says nothing about that period, which depends on temporary supply, alternative providers, and how quickly lapsed skills return. The darker case is one in which the provider's own conduct lowers ϕ, for example by withholding documentation or making exit costly, which goes beyond benefiting from a customer's neglect. If $\phi = \phi_0 - \varkappa$, where $\varkappa \in [0, \phi_0]$ measures such suppression, the limit is crossed whenever $\varkappa > \phi_0 - \phi_c$. Nothing here alleges that any provider does this. The refrigerator that became a supermarket is the comic beginning of the trajectory, and a household that can no longer arrange its own food without it would be near the other end, a position this essay has not shown that any household will reach.

18 Enforcement, Evidence, and Contestability

A restriction defined by an act must still be enforced, and a supplier ordinarily sees requests without seeing purposes. Let $\mathcal{U}$ be the set of possible uses, $\mathcal{P} \subseteq \mathcal{U}$ the uses a clause prohibits, and $\widehat{\mathcal{P}} \subseteq \mathcal{U}$ the uses on which enforcement acts. A *false positive* is a use in $\widehat{\mathcal{P}} \setminus \mathcal{P}$ and a *false negative* one in $\mathcal{P} \setminus \widehat{\mathcal{P}}$. Enforcement against a false positive does not enlarge $\mathcal{P}$. If a person drafting menus for a kindergarten lunchroom were suspended under a rule against distillation, it would be mistaken to report that the supplier prohibits kindergarten menus. The accurate report would be that enforcement had acted on a use outside the rule, if it was outside. The distinction between what triggers a mechanism and what the mechanism is meant to prohibit applies equally to inference engines, whose development may involve comparing output probabilities, drafting candidate tokens, or generating many test cases, activities that could resemble collection of training material to a classifier. No error rate is asserted here for any supplier.

A single false positive differs in consequence from a suspension built on many flags, and many flags can be less informative than their number suggests. Let H_1 be the hypothesis that an account's use lies in $\mathcal{P}$, H_0 the hypothesis that it does not, and $Z_1, \dots, Z_n \in \{0, 1\}$ successive classifier outputs, with all conditioning events having positive probability. The likelihood ratio of the record factorizes as

$$\Lambda_n = \frac{\Pr(Z_1, \dots, Z_n \mid H_1)}{\Pr(Z_1, \dots, Z_n \mid H_0)} = \prod_{m=1}^{n} \frac{\Pr(Z_m \mid Z_1, \dots, Z_{m-1}, H_1)}{\Pr(Z_m \mid Z_1, \dots, Z_{m-1}, H_0)}. \tag{19}$$

Proposition 18.1 (Redundant Flags)**.** *If under each hypothesis every later output equals the first with probability one, then $\Lambda_n = \Lambda_1$ for every n and every record of positive probability.*

Proof. For $m \ge 2$ each conditional probability of the observed Z_m equals one under both hypotheses, so every factor after the first is one. □

Exact duplication is the extreme case. Repeated work on one project can produce outputs that share a cause without being identical, in which case later flags may add some evidence, none, or evidence in either direction, according to the conditional factors. Multiplying the first factor n times is justified only if the outputs are conditionally independent and identically distributed under each hypothesis, which is something to be shown. A likelihood ratio also does not by itself give the probability of a violation, which depends in addition on how common violations are among accounts before any flag is seen.

Call a decision *contestable* if the affected party is told enough about its grounds to identify evidence that could change the outcome, and there is a procedure that will consider that evidence and can reverse the decision. A notice that is the same across materially different grounds cannot by itself discriminate among those grounds, although a recipient may possess other knowledge that does, and a standard notice attached to a very narrow allegation may say enough. A refusal of further review fails the second condition outright. A customer who has neither an informative notice nor such other knowledge cannot tell whether $\mathcal{P}$ or $\widehat{\mathcal{P}} \setminus \mathcal{P}$ contains the use, and so cannot alter the work to comply. A supplier may have reasons for withholding detail, notably that detailed grounds help deliberate violators evade detection. The cost of that reason falls on those wrongly flagged, who could in principle be served by review through a party able to see the grounds without publishing them.

The effect on independent development requires no hostile intention. If developers of tools that would reduce dependence on a proprietary service rely on that service to build them, and are flagged more often than others because their legitimate work resembles prohibited work, enforcement burdens those tools disproportionately. Each condition in that sentence is a hypothesis open to measurement, and none is established here.

19 The Non-Identifiability of Intellectual Extraction

Section 3 reserved the term model distillation for the training of one model to reproduce another's behaviour. There is also an older and much wider activity. Say that *distillation in the broad sense* is the production, from a body of material, of a shorter representation from which conclusions about that material can still be drawn. A summary of chicken recipes qualifies. So do classification, translation, mathematical formalization, and the construction of a theory. A philosophical essay that seeks common principles across control theory, topology, cognitive science, and institutional economics is an especially ambitious instance, since its purpose is to explain a large collection of observations with a small collection of statements. Writing $\mathcal{M}_1, \ldots, \mathcal{M}_n$ for bodies of disciplinary material and $\mathcal{Z}$ for the set of possible syntheses, such work is schematically a map $\mathcal{M}_1 \times \cdots \times \mathcal{M}_n \to \mathcal{Z}$. The result is valuable partly because it is shorter than its sources and partly because it supports further inference, and both properties would also recommend it as training material.

The wide activity does not imply the narrow one. Nothing in the writing of a summary trains a model. The summary may nevertheless be used later to train one, by its author or by a stranger, particularly when it is published openly. Whether a request forms part of model distillation can therefore depend on a use that follows the exchange and that the supplier does

not observe. Let $\mathcal{T}$ be a countable set of possible interaction histories $h = (x_1, y_1, \ldots, x_n, y_n)$, where x_m is a request and y_m its response. Let P_D and P_L be the distributions of histories among prohibited and permitted uses respectively, in the sense of the sets $\mathcal{P}$ and $\mathcal{U} \setminus \mathcal{P}$ of Section 18, and let $\pi \in (0, 1)$ be the proportion of prohibited uses. A classifier is a function $f : \mathcal{T} \to [0, 1]$ giving the probability that a history is flagged. The *total variation distance* between the distributions is $\mathrm{TV}(P_D, P_L) = \frac{1}{2} \sum_h |P_D(h) - P_L(h)|$, which is zero when they coincide and one when no history can arise under both.

Proposition 19.1 (Limit of Observation)**.** *Every classifier on $\mathcal{T}$ has error probability at least*

$$e^* = \sum_{h \in \mathcal{T}} \min\{\pi P_D(h),\, (1 - \pi) P_L(h)\},$$

and some classifier attains it. When $\pi = \frac{1}{2}$, $e^* = \frac{1}{2}(1 - \mathrm{TV}(P_D, P_L))$.

Proof. The error probability of f is $\sum_h \left[\pi P_D(h)(1 - f(h)) + (1 - \pi) P_L(h) f(h)\right]$. Each term is a weighted average of $\pi P_D(h)$ and $(1 - \pi) P_L(h)$ and so is at least their minimum, with equality when $f(h)$ is one if $\pi P_D(h)$ is the larger and zero otherwise. For equal proportions, use $\min\{a, b\} = \frac{1}{2}(a + b - |a - b|)$ and sum. □

Corollary 19.2 (Identical Histories)**.** *If $P_D = P_L$, then $e^* = \min\{\pi, 1 - \pi\}$, which is achieved by ignoring the history altogether.*

Proof. Substitute $P_D = P_L = P$ and use $\sum_h P(h) = 1$. □

In the simplest case, two customers submit the same requests about recipes and receive the same replies. One is compiling a cookbook and the other a training set for a competing culinary model. Any classifier that sees only the exchange assigns both the same probability $f(h)$ of being flagged, so the chance of wrongly flagging the first and the chance of wrongly passing the second add to exactly one. The result differs in kind from Proposition 18.1. That proposition concerned repetition that adds nothing to the first observation. The present one concerns a distinction that may be absent from every observation, however numerous or independent, whenever the two distributions of whole histories coincide. Where they differ, further observation can help, and the proposition says nothing about how different they are for any actual classifier. Evidence from outside the history, such as account structure, payment records, or a published model, can also separate cases that the exchanges cannot.

One may be tempted to conclude that no false positive is possible, on the ground that every user of a language model is distilling something. The conclusion follows if the prohibition is drawn that widely. Should $\mathcal{P}$ contain every use that involves distillation in the broad sense, and should every use in $\mathcal{U}$ involve it, then $\mathcal{P} = \mathcal{U}$ and the set $\widehat{\mathcal{P}} \setminus \mathcal{P}$ is empty whatever the detector does. The false positive has then been abolished by prohibiting the product. Under a rule confined to training competing models, false positives remain perfectly well defined, although the supplier may be unable to find them in its records. A drafter thus faces a choice between a rule wide enough to be infallibly applied and a rule narrow enough to permit use, whose violations cannot always be recognized. Nothing here shows that any supplier has adopted the wide rule.

Stated in the idiom of the 2046 annual report, the company encourages its customers to acquire knowledge, provided that the knowledge does not improve their capacity to acquire more of it in any transferable way. Summaries remain welcome, except where they prove more useful than the material summarized. The joke has a serious residue. The value of an intellectual instrument lies in improving what its users can understand and make, and a prohibition broad enough to prevent every potentially competitive improvement would prohibit the benefit that was sold.

20 Recursive Corporate Expansion

Capital finances organizational intelligence, organizational intelligence extends productive reach, and reach generates further capital. A coupled model of this cycle is

$$\dot{K}_i = \Pi_i(R_i, A_i) - I_i - \delta_K K_i, \tag{20}$$

$$\dot{A}_i = \zeta_i I_i^{\alpha} A_i^{\beta'} - \delta_A A_i, \tag{21}$$

$$\dot{R}_i = f_i(K_i, A_i, R_i) - \sum_{j \neq i} g_{ij}(R_i, R_j), \tag{22}$$

where Π_i is operating profit, $I_i \geq 0$ is investment in organizational intelligence, $\delta_K, \delta_A > 0$ are rates of depreciation, $\zeta_i > 0$ is the efficiency of that investment, f_i is the rate of growth of reach in the absence of rivals, and $g_{ij} \geq 0$ is the reduction caused by firm j. The exponent β' governs returns to existing capability in producing more of it and is distinct from the exponent β in (4). With the functions unspecified the system permits nearly any behaviour, which is a polite way of saying it predicts nothing. Definite statements require reductions.

20.1 A Threshold for Self-Reinforcing Reach

Return to (4), drop the index, and suppose capital and capability adjust quickly to reach, with $K = kR^a$ and $A = bR^c$ for positive constants. Then

$$\dot{R} = c_0 R^m - \delta R, \qquad c_0 = \alpha k^{\beta} b^{\gamma}, \quad m = a\beta + c\gamma. \tag{23}$$

Proposition 20.1 (Feedback Threshold)**.** *Let $R(0) = R_0 > 0$ and, for $m \neq 1$, let $R^* = (\delta/c_0)^{1/(m-1)}$. If $m < 1$, every solution converges to R^*. If $m = 1$, then $R(t) = R_0 e^{(c_0 - \delta)t}$. If $m \neq 1$ and $R_0 = R^*$, then $R(t) = R^*$ for all t. If $m > 1$, solutions with $R_0 < R^*$ decay to zero, and solutions with $R_0 > R^*$ diverge at the finite time*

$$t^* = \frac{1}{(m-1)\delta} \ln \frac{1}{1 - \delta R_0^{1-m}/c_0}. \tag{24}$$

Proof. For $m \neq 1$ the substitution $z = R^{1-m}$ gives the linear equation $\dot{z} = (1-m)c_0 + (m-1)\delta z$, with solution $z(t) = c_0/\delta + (z(0) - c_0/\delta)e^{(m-1)\delta t}$. If $m < 1$ the exponential decays and $z \to c_0/\delta$. If $m > 1$ and $R_0 < R^*$ then $z(0) > c_0/\delta$, so z grows without bound and $R \to 0$. If $R_0 > R^*$ then $z(0) < c_0/\delta$, so z decreases and first vanishes at the stated time. If $R_0 = R^*$ then $z(0) = c_0/\delta$ and z is constant. The case $m = 1$ is immediate. □

Runaway expansion in this reduction requires $a\beta + c\gamma > 1$ and an initial reach above a threshold. Neither condition is a finding about the world. Moreover, a divergence in finite time cannot occur in an economy represented by a finite graph, where occupied territory is bounded by the number of activities. The correct reading of t^* is that the approximation leaves its domain of validity somewhere before that date. A corporation whose model predicts infinite commercial reach on a particular Tuesday should revise the model before revising the calendar.

20.2 Coexistence and Exclusion

For two symmetric firms take

$$\dot{R}_i = R_i\,(g - \iota R_i - \nu R_j), \qquad i \neq j, \tag{25}$$

with intrinsic growth rate $g > 0$, own-scale limitation $\iota > 0$, and interference $\nu \geq 0$, perhaps increasing with prospective overlap. This is a competitive Lotka–Volterra system, the standard model of two populations each slowed by its own size and by the other's. The symmetric equilibrium has $R_1 = R_2 = R^* = g/(\iota + \nu)$. The Jacobian, the matrix of first derivatives governing small departures from it, has eigenvalues $-R^*(\iota + \nu)$ and $-R^*(\iota - \nu)$. When $\iota > \nu$ both are negative and the equilibrium is stable. When $\iota < \nu$ they differ in sign, making the equilibrium a saddle, approached only along the line $R_1 = R_2$; from unequal positive starting values the initially larger firm tends to g/ι and the other to zero, since the ratio R_1/R_2 has logarithmic derivative $(\nu - \iota)(R_1 - R_2)$. When $\iota = \nu$ the ratio is constant and there is a line of equilibria. If economies of scope reduce effective own-scale limitation while greater overlap raises interference, a configuration of coexisting specialists can lose stability. The result does not allow for entry, differentiation, merger, or antitrust enforcement, and treats an entire corporation as one number. Legal scholarship on platform businesses that both host and compete with other firms indicates how much institutional detail the symbol ν conceals [14]. It need not follow that any particular corporation inevitably prevails, because in the symmetric model the identity of the survivor is fixed by the starting positions.

21 The Industrial Universe in 2036 and 2046

The following scenarios are satire. They were not computed from any equation in this essay, no parameter was estimated in order to produce them, and no likelihood is assigned to them. Their relation to the preceding mathematics is that of an illustration to a text: it shows what a regime would look like if the text's conditions held, and it offers no reason to think they will.

Corporation	2036	2046
Refrigerator manufacturer	Automated grocery procurement	Agricultural coordination and household provisioning
Factory simulator vendor	Autonomous production scheduling	Manufacturing administration across industries
Accounting software provider	Automated financial management	Allocation of corporate resources and governance of investment
Automobile manufacturer	Autonomous transport networks	Coordination of logistics, infrastructure, and cities
Language model provider	General-purpose professional automation	Prospective competition with all of the above

21.1 From a Retrospective History Published in 2046

> The practice of assigning each enterprise to a single industry was retained for statistical continuity until the category designated Other came to contain the greater part of national output. A committee convened to subdivide it proposed classifying firms by their original product. Under the revised scheme the largest agricultural coordinator in the northern hemisphere was recorded under Household Appliances, a municipal water system under Accounting Services, and three hospitals under Automotive. The committee noted that these entries were accurate as to origin and recommended no further action. Its minutes were prepared by a subsidiary of a former manufacturer of vacuum cleaners, which also supplied the committee's agenda, its lunch, and its liability coverage. One member asked whether this arrangement presented a conflict of interest. The question was referred to the same subsidiary, which found the matter within its competence.

21.2 From an Annual Report for the Fiscal Year 2046

> Refrigeration remains central to our identity. During the year our appliances maintained recommended temperatures in a record number of households, and we are pleased to report continued progress in door seals. Our other activities, comprising crop planning, freight, personal lending, nutritional guidance, crop insurance, and the operation of two seaports, are undertaken solely to enhance the refrigeration experience. Shareholders have asked whether the company has strayed from its core business. The board has reviewed the question and determined that the core business has expanded to meet the company.

Both documents are invented. Neither is a prediction in any sense that could be scored in 2046, apart from the general expectation that annual reports will continue to describe whatever has happened as intended.

22 Why the Universal Corporation Might Fail

A serious treatment must give the opposing forces equal formal standing. Increasing scope can bring diseconomies of complexity, correlated failures, computational bottlenecks, regulatory resistance, conflicts between principals and agents, and the loss of local knowledge. Let organizational overhead grow superlinearly with reach,

$$\Gamma(R) = \kappa R^{1+\varepsilon}, \qquad \kappa, \varepsilon > 0, \tag{26}$$

and let gross returns be $Y(R) = yR^q$ with $y > 0$ and $0 < q < 1 + \varepsilon$. Net return $Y - \Gamma$ has a unique maximum at

$$R^\circ = \left(\frac{qy}{\kappa(1+\varepsilon)} \right)^{1/(1+\varepsilon-q)}. \tag{27}$$

Beyond this point further expansion reduces net return. Capability that lowers κ moves the optimum outward, and R° grows without bound as κ approaches zero, so the argument for a finite firm depends on overhead that technology reduces but does not remove.

Overhead can also be placed directly in the feedback model:

$$\dot{R} = c_0 R^m - \delta R - \kappa R^{1+\varepsilon}. \tag{28}$$

If $m > 1$ and $\varepsilon > m - 1$, the negative term dominates for large R and every solution remains bounded. Dividing by R, the growth rate $h(R) = c_0 R^{m-1} - \delta - \kappa R^\varepsilon$ is negative near zero and for large R, with a single maximum at

$$R_m = \left(\frac{c_0(m-1)}{\kappa\varepsilon} \right)^{1/(\varepsilon-m+1)}.$$

When $h(R_m) > 0$ there are two positive equilibria, the lower unstable and the upper stable. A firm must exceed one size to grow at all, and growth then carries it to another size where it stops. At the boundary $h(R_m) = 0$ the two positive equilibria merge into one at R_m, approached by trajectories starting above it and abandoned by those starting below, which fall toward zero; this saddle-node configuration separates bounded growth from universal contraction. When $h(R_m) < 0$ every firm of this type contracts. The same family of equations thus accommodates collapse, a bounded giant, and, if overhead is removed, divergence. Choosing among them is an empirical matter on which the essay has nothing to report.

Other limits lie outside these scalars. The variation among enterprises establishes a floor on what pooled experience can say about a new one. The example of three departments shows that local successes need not combine. A provider that administers many firms with one policy may also cause them to err together, so that mistakes which once offset one another across independent managers accumulate. Finally, customers who recognize their own growing dependence may invest in raising ϕ, collectively fund alternatives, or ask legislators to raise r_{ij}. The universal administrator might be superb at each isolated function and unable to administer their combination, or able to do so and forbidden.

23 The Right to Become a Competitor

The institutional question can be stated without assuming any of the forecasts. If access to a general-purpose intellectual instrument becomes necessary for independent productive work, and if the supplier's commercial ambitions have no stated boundary, then a prohibition on using the instrument to develop competing things governs a widening share of what customers might do. Suppliers have intelligible reasons for restrictions. A model that is expensive to produce may be imitated cheaply by training on its outputs, and an investor who cannot prevent this may decline to produce it. Protection against extraction of a proprietary model is a different thing from a restriction on the customer's acquisition of competence, even where one clause attempts both.

The distinction can be drawn along the lines Proposition 14.1 makes visible. A restriction defined by an act, such as systematically collecting outputs to reproduce the supplier's model, has an extent that does not grow merely because the supplier enters new markets. A restriction defined by the supplier's present and named offerings grows only when those offerings change, and visibly. A restriction defined by prospective competition grows with the supplier's light cone, which the customer cannot observe and the supplier can enlarge unilaterally. The normative position advanced here is that purchasing access to an intellectual instrument should not require surrendering the right to develop capabilities that might eventually exceed those of its manufacturer. This is an argument about fairness and the conditions of independent enterprise. The mathematics clarifies what is being traded but cannot settle whether the trade should be allowed, and a reader who gives greater weight to the supplier's incentive to invest may reasonably draw the line elsewhere.

24 Conclusion: The Last Industry

The analysis began with a refrigerator and a series of modest improvements. Along the way it defined a graph and a hypergraph whose weights no one has measured, a notion of industry that can only coarsen under cheaper transitions, a reminder that methods have operating ranges and sometimes unlisted duties, a revenue condition for preferring a gateway to an appliance, a case in which repeated flags add no evidence, a bound on what any classifier can learn from exchanges about their later use, a supplier whose willing market shrinks as its ambitions grow, a separation between who owns and who administers, a limit past which a service cannot readily be refused, a light cone that records what a firm might reach without saying what it will do, a floor beneath which capability cannot lower the cost of entry, a ceiling on what pooled data can say about the next enterprise, a threshold above which a simplified growth law diverges, and an overhead term that removes the divergence. The most defensible contribution is the treatment of industrial boundaries as features of a geometry of organizational competence that participants can deform. The grabby analogy supplies a picture of expansion and the light cone supplies an object on which expansion can be examined. The equations are exploratory constructions and should not be mistaken for economic laws.

A useful program of work would estimate transition costs and joint prerequisites from

observed entry decisions, separate reductions in knowledge cost from simultaneous changes in regulation and capital requirements, and compare predicted entry with entry that actually occurred. Until something of that kind has been done, the essay's hypothesis remains unsupported. It has not been shown that industrial specialization is dissolving. It has not been shown that artificial intelligence lowers the relevant costs by amounts that matter, that firms will use any reduction to move into distant activities, or that the movement will concentrate control. The dates 2036 and 2046 were selected because they are comfortably far away.

By the middle of the century it may be difficult to distinguish an appliance manufacturer from an agricultural corporation, a logistics provider, a software company, or a financial intermediary. It may also be perfectly easy. The sharpest satirical possibility is not that one corporation comes to own everything. It is that everything stays independently owned while one corporation supplies the software through which the independence is exercised. Should the former occur, the final irony would be that a corporation seeking to supply organizational intelligence to everyone discovered that its principal prospective competitors were its own customers. Should the latter occur, this essay will have been a careful study of conditions that did not obtain, and the refrigerator, to its considerable credit, will still be keeping the milk cold.

References

[1] Anderson, Ian A., et al. *Laboratory Performance of Pervious Concrete Subjected to Deicing Salts and Freeze-Thaw.* TRC Report 15-006. University of Vermont Transportation Research Center, 2015.

[2] Anthropic. *Commercial Terms of Service*, Section D.4, Use Restrictions. Effective June 17, 2025. Accessed November 26, 2026. `https://assets.claude.com/cbc0cee5a5580ef30c81ccf745ef439aa7c47921.pdf`.

[3] Bosch. "Bosch Group Worldwide: Figures as of December 31, 2025." Accessed November 26, 2026. `https://www.bosch.africa/our-company/bosch-group-worldwide/`.

[4] Bresnahan, Timothy F., and Manuel Trajtenberg. "General Purpose Technologies: 'Engines of Growth'?" *Journal of Econometrics* 65, no. 1 (1995): 83–108.

[5] Chandler, Alfred D., Jr. *Scale and Scope: The Dynamics of Industrial Capitalism.* Cambridge, MA: Harvard University Press, 1990.

[6] Coase, Ronald H. "The Nature of the Firm." *Economica* 4, no. 16 (1937): 386–405.

[7] Conant, Roger C., and W. Ross Ashby. "Every Good Regulator of a System Must Be a Model of That System." *International Journal of Systems Science* 1, no. 2 (1970): 89–97.

[8] Dick, Philip K. "Autofac." *Galaxy Science Fiction*, November 1955.

[9] Grieves, Michael, and John Vickers. "Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems." In *Transdisciplinary Perspectives on Complex Systems*, edited by Franz-Josef Kahlen, Shannon Flumerfelt, and Anabela Alves, 85–113. Cham: Springer, 2017.

[10] Hanson, Robin, Daniel Martin, Calvin McCarter, and Jonathan Paulson. "If Loud Aliens Explain Human Earliness, Quiet Aliens Are Also Rare." *The Astrophysical Journal* 922, no. 2 (2021): 182.

[11] Hayek, Friedrich A. "The Use of Knowledge in Society." *American Economic Review* 35, no. 4 (1945): 519–530.

[12] Hidalgo, César A., and Ricardo Hausmann. "The Building Blocks of Economic Complexity." *Proceedings of the National Academy of Sciences* 106, no. 26 (2009): 10570–10575.

[13] Hidalgo, César A., Bailey Klinger, Albert-László Barabási, and Ricardo Hausmann. "The Product Space Conditions the Development of Nations." *Science* 317, no. 5837 (2007): 482–487.

[14] Khan, Lina M. "Amazon's Antitrust Paradox." *Yale Law Journal* 126, no. 3 (2017): 710–805.

[15] Lang, Larry H. P., and René M. Stulz. "Tobin's q, Corporate Diversification, and Firm Performance." *Journal of Political Economy* 102, no. 6 (1994): 1248–1280.

[16] Montgomery, Cynthia A. "Corporate Diversification." *Journal of Economic Perspectives* 8, no. 3 (1994): 163–178.

[17] Noyce, David A., et al. *Traffic Signal LED Module Specification Workshop and Informational Report for Snow Conditions*. Report FHWA-HOP-13-010. Washington, DC: Federal Highway Administration, 2014.

[18] Panzar, John C., and Robert D. Willig. "Economies of Scope." *American Economic Review* 71, no. 2 (1981): 268–272.

[19] Parekh, Anil. "Benchmarking of Energy Savings Associated with Compact Fluorescent Lighting in Homes." In *Proceedings of the 2006 ACEEE Summer Study on Energy Efficiency in Buildings*, 9-260–9-269.

[20] Penrose, Edith T. *The Theory of the Growth of the Firm*. Oxford: Basil Blackwell, 1959.

[21] Simon, Herbert A. *The Sciences of the Artificial*. Cambridge, MA: MIT Press, 1969.

[22] Teece, David J. "Economies of Scope and the Scope of the Enterprise." *Journal of Economic Behavior and Organization* 1, no. 3 (1980): 223–247.

[23] Williamson, Oliver E. *Markets and Hierarchies: Analysis and Antitrust Implications*. New York: Free Press, 1975.