


BCG

THE BOSTON CONSULTING GROUP

BORGES' MAP

Navigating a World of Digital Disruption

by Philip Evans & Patrick Forth



. . . In that Empire, the Art of Cartography attained such Perfection that the map of a single Province occupied the entirety of a City, and the map of the Empire, the entirety of a Province. In time, those Unconscionable Maps no longer satisfied, and the Cartographers Guilds struck a Map of the Empire whose size was that of the Empire, and which coincided point for point with it.

**“ON EXACTITUDE IN SCIENCE”
JORGE LUIS BORGES**

DIGITAL DISRUPTION is not a new phenomenon. But the opportunities and risks it presents shift over time. Competitive advantage flows to the businesses that see and act on those shifts first. We are entering the third, and most consequential, wave of digital disruption. It has profound implications not only for strategy but also for the structures of companies and industries. Business leaders need a new map to guide them. This article explains the factors underlying these disruptive waves, outlines the new strategic issues they raise, and describes a portfolio of new strategic moves that business leaders need to master.

In the first wave of the commercial Internet, the dot-com era, falling transaction costs altered the traditional trade-off between richness and reach: rich information could suddenly be communicated broadly and cheaply, forever changing how products are made and sold. Strategists had to make hard choices about which pieces of their businesses to protect and which to abandon, and they learned that they could repurpose some assets to attack previously unrelated businesses. Incumbent value chains could be “deconstructed” by competitors focused on narrow slivers of added value. Traditional notions of who competes against whom

were upended—Microsoft gave away Encarta on CDs to promote sales of PCs and incidentally destroyed the business model of the venerable *Encyclopædia Britannica*.

In the second wave, Web 2.0, the important strategic insight was that economies of mass evaporated for many activities.¹ Small became beautiful. It was the era of the “long tail” and of collaborative production on a massive scale. Minuscule enterprises and self-organizing communities of autonomous individuals surprised us by performing certain tasks better and more cheaply than large corporations. Hence Linux, hence Wikipedia. Because these communities could grow and collaborate without geographic constraint, major work was done at significantly lower cost and often zero price.

Smart strategists adopted and adapted to these new business architectures. IBM embraced Open Source to challenge Microsoft’s position in server software; Apple and Google curated communities of app developers so that they could compete in mobile; SAP recruited thousands of app developers from among its users; Facebook transformed marketing by turning a billion “friends” into advertisers, merchandisers, and customers.

Now we are on the cusp of the third wave: hyperscaling. Big—*really* big—is becoming beautiful. At the extreme—where competitive mass is beyond the reach of the individual business unit or company—hyperscaling demands a bold, new architecture for businesses.

1. We use the term “economies of mass” to capture increasing returns to the volume of current activities (economies of scale), the breadth of current activities (economies of scope), and the cumulative volume of past activities (economies of experience, a proxy for organizational learning).

Up The Amazon

THESE WAVES of innovation have come one after another, but they have also overlapped and, in many cases, amplified each other. The exemplar of this is Amazon, whose successive innovations have been at the leading edge of each phase.

Jeff Bezos's initial idea was to exploit the Web to deconstruct traditional bookselling. With just a well-designed website that piggybacked on the inventory and the index of book wholesaler Ingram, Amazon offered a catalogue ten times larger than that of the largest Main Street superstore, at prices 10 to 15 percent cheaper.

But that was not a sustainable advantage: competitors such as BN.com would rapidly establish comparable selections and price points. Amazon went on to exploit the emerging economics of community. The Amazon Associates program allowed bloggers to post widgets endorsing books and to earn a commission on click-throughs. Amazon curated its reviewer community, encouraging the rating of reviews and awarding badges to the best-rated reviewers. It extracted insights from the behavior of its community of customers and became an early adopter of

collaborative filtering algorithms, goosing sales with messages that “people like you who bought X often buy Y.” On the selling side, the company launched Amazon Marketplace as a fixed-price rival to eBay: a platform hosting a community of small sellers that now numbers more than 2 million. All these strategies benefited from the network effect: the more participants, the more choices; the more reviews, the richer the experience.

Well ahead of others, Amazon also embraced what became the third wave of digital disruption, exploiting opportunities to hyperscale. It built a global network of 80 fulfillment centers and relentlessly broadened its product line to include almost any product that can be delivered by truck. It offered fulfillment services as an option for small merchants, which could thereby distribute almost as efficiently as Walmart. Amazon became the broad river of commerce suggested by its name. In parallel, and almost incidentally, it built impressive scale in its data centers and world-class skill in operating them. It then reconceptualized its own computing infrastructure as a product in its own right. The first step, in 2003, was to standardize the interfaces between data services and the rest of Amazon’s business. In 2006 (and in the teeth of criticism from Wall Street), Bezos opened Amazon Web Services (AWS)—cloud comput-

“Bezos saw business architecture as a strategic variable, not a given.”

ing—as a standalone service. This started as the simple rental of raw computing capacity but evolved into a complex stack of computing services. (Amazon even sells the service to competitors such as Netflix.) According to Gartner, in 2013, AWS had five times the capacity of the next 14 competitors put together.

From deconstruction to community curation to hyperscaling: at no point did Amazon sit back and wait for trends to emerge. Rather, it seized the strategic opportunities presented by

each successive wave of disruption, ruthlessly cannibalizing its own business where necessary. E-books were inevitable, so it launched the Kindle; customer information and scale in data processing are critical, so it sells cloud services to its own competitors. And at no point did Bezos restrict one business to protect another—Amazon is now run as four loosely coupled platforms, three of which are profit centers: a community host, supported by an online shop, supported by a logistics system, supported by data services.

Unlike many of his rivals, Bezos saw business architecture as a strategic variable, not a given. He did not harness technology to the imperatives of his business model; he adapted his business model to the possibilities—and the imperatives—of technology.

The Limit: Scale of 1:1

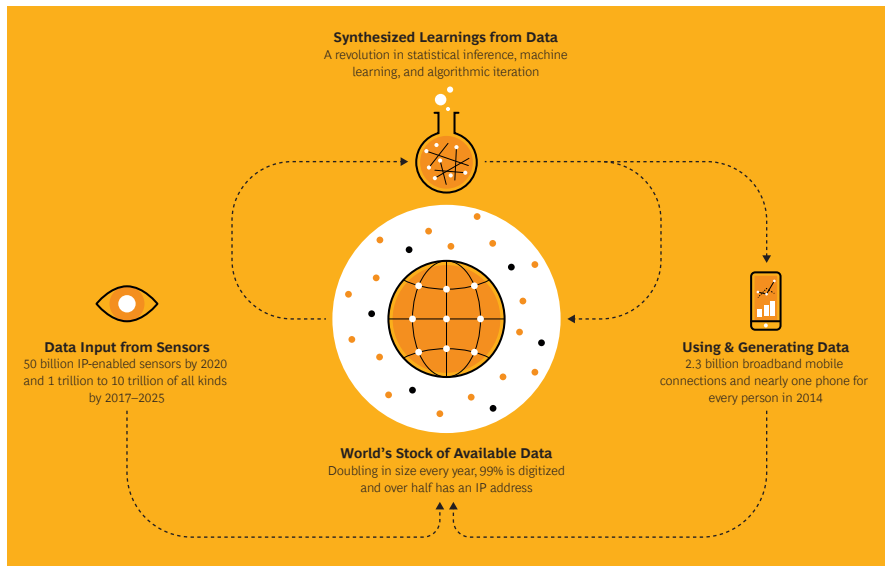
THE TECHNOLOGICAL IMPERATIVES described above are not unique to websites: they are universal. The underlying forces, of course, are the long-term falling costs of computing, communications, and storage. But in just the last six to eight years, these forces have begun to converge on an extraordinary pattern that begins to evoke Borges' imagined world. (See Exhibit 1.)

Ubiquitous Sensing. The number of Internet-connected devices hit 8.7 billion in 2012. IP-enabled sensors are projected to exceed 50 billion by 2020. The number of sensors of all types is variously projected at between 1 trillion and 10 trillion between 2017 and 2025. The lower estimate translates to 140 sensors for every man, woman, and child on the planet.

Ubiquitous Connectivity. Mobile broadband subscriptions reached 2.3 billion in 2014—five times the number in 2008. The smartphone is the fastest-adopted technology ever; the biggest absolute growth is in India and China. At the end of 2014 there were nearly 7 billion mobile-cellular subscriptions globally—nearly one per person on Earth.

Convergent Data. The world's production of data grew 2,000-fold between 2000 and 2012. Its stock of data is expected to double every two years; 99 percent of it is digitized and half has an IP address. This means that half of the world's data can now be put together, at near-zero cost, to reveal patterns previously invisible. Half of the world's data is already, technically, a single, universally accessible document.

Exhibit 1 Ubiquitous Sensing and Connectivity Drive an Emergent Digital Ecosystem



Source: BCG analysis.

All this data is linked by fixed and mobile communication networks and is managed by layers of modular, interoperable software. Software is replacing hardware, rapidly accelerating the speed of innovation: the life cycle of many products and services (previously defined by physical obsolescence) is shrinking from decades to just days between software updates. Information is comprehended and applied through fundamentally new methods of artificial intelligence that seek insights through algorithms using massive, noisy data sets. Since larger data sets yield better insights, big is beautiful. Data wants to be big, and businesses struggle to keep up.

The asymptote is where sensing, connectivity, and data merge into a single system. Every person and object of interest is connected to every other: the traffic readout on a mobile phone becomes the aggregation of all the data provided by all the mobile devices in the area reading the traffic. The world becomes self-describing and self-interpreting. At its outer limit, the digital map becomes the world itself. *The world and our picture of the world are becoming the same thing*: an immense, self-referential document. We are living in Borges' map.

A “perfect” map with a scale of 1:1 would encompass its world and describe its reality in complete detail. It would array the granular in the context of the universal. That is precisely the architecture toward which business (and human organization in general) is evolving: the arbitrarily large as a platform for arraying the arbitrarily small. And each arbitrarily small agent—whether a person, a thing, or a function—reads whatever parts of the map are needed to get to its goal.

Data in this world is *infrastructure*: a long-lived asset, general in purpose, capital intensive, and supporting multiple activities. Inference, by contrast, is short lived, real time, trivially cheap, specific to a problem or task, continuously adapting, and perpetually self-correcting. The organizational correlates of data and inference polarize in parallel. The minimum efficient scale for data systems and facilities is rising beyond the reach of individual business units within a company, and ultimately beyond that of many companies. Yet tens, thousands, maybe millions of devices or individuals or teams—sometimes sharing, sometimes competing—access that data to solve problems. Polarizing economics of mass are pushing the advantage simultaneously to the very big and the very small, and a new architecture is emerging for businesses of all sizes.

Hyperscale and Architectural Innovation

HEALTH CARE is a prime example of the transformational power of these new business architectures. This huge and dysfunctional industry is at the beginning of a transformation. The cost of sequencing a human genome in 2001 was \$100 million, and mapping just one (James Watson's) took nearly ten years. Today it costs less than \$1,000. In two or three years, it will cost \$100, and sequencing will take just 20 minutes. The number of sequences has grown as the cost has fallen: the Million Human Genomes Project is up and running—in Beijing. Gene mapping is shifting from an abstract research activity to a clinical one, in which a doctor customizes treatment to the patient's unique genomic makeup.

The pattern is clear: big-data techniques will be used to spot fine-grained correlations in a patient's genomic data, medical history, symptoms, protocols, and outcomes, as well as real-time data from body sensors. Medicine will advance by decoding *immense, linked, cheap, noisy* data sets instead of the *small, siloed, expensive, clean, and proprietary* data sets now generated independently by hospital records, clinical trials, and laboratory experiments. These databases will make it possible for practitioners and even groups

of patients to become researchers and for breakthroughs to be quickly shared around the world.

Of course, progress will be slower than the rush of early expectations. The real hurdle is a profound lack of cooperation. Medical records, even when digital, are kept in proprietary formats, and interoperable data standards are difficult to negotiate. But even after payers have cajoled providers into addressing that problem, how will all that data be melded when providers, insurers, device companies, pharma companies, Google, patients, and governments possess different pieces of the data elephant and view data as a source of competitive advantage? And even though pooled data makes clinical sense, how are privacy and patient rights going to be protected? The fundamental answer is *architecture*. Health care systems will need an infrastructure of trusted, secure, neutral data repositories.

“Economies of ‘mass’ are intensifying across the economy, driving new models of collaboration.”

This is already happening. Nonprofit organizations are becoming platforms for the curation of genomic databases, with an emphasis on data protection. Registries run by universities and medical associations are emerging as repositories for shared data specific to particular medical conditions. Security and encryption technologies are starting to reconcile the scientific imperative to share with the personal right to privacy. Drug research is becoming such a massive undertaking that competition is inefficient and prohibitively expensive, so even the big pharmaceutical companies are looking for ways to collaborate. Building a shared data infrastructure will be one of the strategic challenges of the next decade for the health care industry and for policy makers.

The health care industry is not an anomaly. Economies of “mass”—of scale, scope, and experience—are intensifying across the economy, driving new models of collaboration. The operating

system of a modern car has 100 million lines of code (and that's before Google displaces the driver), and since code is a fixed cost, the largest car makers have an advantage. So the smaller companies in that industry are adopting a shared, open-source model: the major component manufacturers launched the Automotive Grade Linux reference platform in April 2014.



The electricity generation industry is evolving toward its version of the Internet, the smart grid. Scale economies here are polarizing toward the very small, as domestic rooftop solar panels, electric-car batteries, and wind turbines become viable ways to feed power back into the grid. But they are also moving toward the very large, as the intermittent nature of these power sources requires new, shared, long-distance transmission networks large enough to arbitrage the regional vagaries of sun and wind. The shift to this ultrabroadband world will require a colossal investment in fiber infrastructure. This is a choke point in many markets, because competing service providers cannot rationally justify the high fixed cost of deployment. Municipal and national policy makers are increasingly recognizing that ultrabroadband will be fundamental to creating jobs and competing in the new economy, so they are stepping in to get those networks built as shared infrastructure: from Stockholm to Chattanooga, Singapore to Australia.

From Maps to Stacks

THE ANALOGY with the map imagined by Borges is not just a splash of literary fancy. The 750 MB digital map of an individual genome corresponds one-to-one with the 21 million base pairs of human DNA. Google aims to organize all the world's data. Look at something through the lens of Google Glass, and the object describes itself. Facebook wants to map the connections of everybody with everybody. Military planners aspire to “total battlespace awareness.” According to General Keith Alexander, former head of the U.S. National Security Agency, in order to find the needle, “you need the haystack”—the haystack being all messages, all conversations, all everything.

These maps describe reality with a granularity and comprehensiveness that is entirely new. But they also shape reality. Facebook has redefined what it means to be a friend. Waze maps the flow of traffic, and thereby equips its users to change it. Sensor technologies render a shopper browsing in a physical store as visible (and malleable) as he or she would be online. Consumers turn complementary technologies on retailers by using Yelp and Kayak to find alternative products and vendors. Google Search maps Web links; SEOs (search engine optimizers) map Google

Search; Google Search maps SEOs mapping... and so on. The map and the terrain, the sign and the signified, the virtual and the real become indistinguishable.

“A stacked ecosystem blows up the classic trade-off between efficiency and innovation.”

These technologies can be extended infinitely and are all converging on the instantaneous. They are built in layers, mixing real and virtual. A sensor is built into a parking space and another sensor is built into a car: one network enables the municipality to charge the driver for parking; another allows the driver to find an empty parking space. Richer networks will eventually enable an autonomous car to navigate and park itself. Still richer networks will enable cars to self-organize into “platoons,” lines of cars like a train that are given high-speed priority by smart traffic lights. Physical cars swarm on an infrastructure of roads, virtual agents swarm on an infrastructure of data: each is a layered system, but there is a one-to-one correspondence between them, and they continuously modify one another.

Modularity and layering, granularity and extensibility, the symbiosis of the very large with very small: these are the recurrent themes of the transformative technologies of our era. Borges’ map is a rich metaphor for an emerging architecture in business: the architecture of the stack.

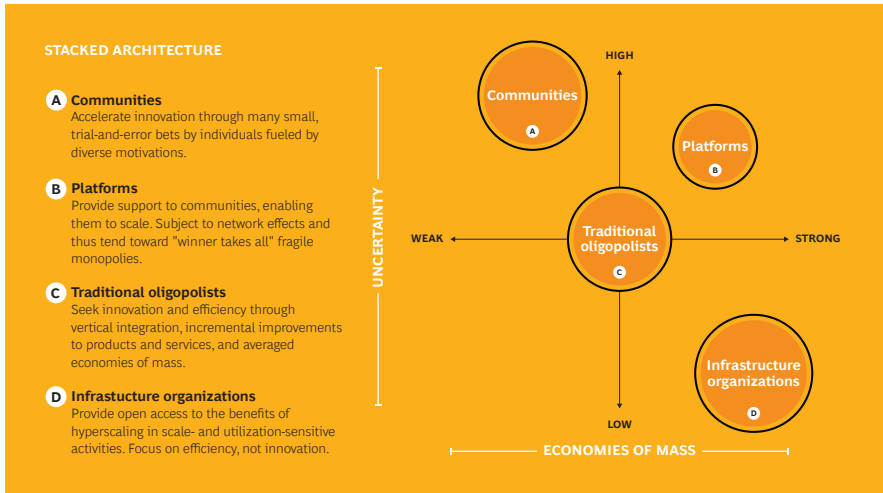
Stacks: The New Architecture of Business

BUSINESSES IN most industries have a classic oligopolistic structure, with a small number of companies competing on similar vertical value chains. In many cases, this will evolve into a much more diverse architecture of horizontal layers: shared infrastructure on the bottom, producing and consuming communities on the top, and traditional oligopolists competing in the middle. Borrowing a metaphor from technology, we call these industrial ecosystems “stacks.”

Stacks are a compelling model when the benefits of community-based innovation in higher layers and improved utilization in lower layers exceed the additional transaction costs incurred by breaking up value chains. Falling transaction costs make that trade-off progressively more favorable. In the right circumstances, a stacked ecosystem blows up the classic trade-off between efficiency and innovation.

Within a stack, different kinds of institutions coexist in a mutually sustaining structure, each focused on the activities where it has an advantage. While the pattern of layers varies, there are four broad types. (See Exhibit 2.)

Exhibit 2 Digital Disruption Has Enabled New Institutional Options with Distinct Economics



Source: BCG analysis.

Communities of users, professionals, and small entrepreneurs are typically found toward the top of the stack, receiving services from lower layers. They flourish when uncertainty is high but the economies of mass are weak and where innovation comes through many small, seat-of-the-pants, trial-and-error bets. Community members often innovate for their own use or amusement, and sharing or selling to peers is an afterthought. Uncoordinated, autonomous agents compete and collaborate. They chase a million dead ends, then flood the occasional success with a million tweaks. They vary in just how “communitarian” they are: at one extreme, small developers competing to write apps for the iPhone without much, if any, shared social capital; at the other, Linux hackers or Wikipediasts collaborating to build a massive body of shared intellectual property.

Infrastructure organizations are typically found at the bottom of the stack, since they provide services to other layers without receiving services in return. They are most useful when uncertainty is low and economies of mass (specifically scale) are overwhelming. Thus, long-lived monopolies, utilities, or coops. Their core competence is in long-term, numbers-driven capacity management. Their mission is to be efficient and maximize access. It is not, in general, their job to innovate, though they may deploy successive generations of technology innovations generated by others.

Curatorial platforms, narrowly defined as organizations that exist solely as hosts for communities, are a hybrid. In the stack, they lie immediately below the community they curate. Often they start as an innovation by some community member. But they can grow into something resembling infrastructure, sometimes with stunning speed. Where social capital is negligible, the platform becomes a marketplace or a tech standard. Where social capital matters, the platform is rule setter and cultural arbiter, legitimate by grace of the community it supports.

Traditional oligopolists occupy the broad middle of the stack.

They have the advantage when uncertainty is high but not incalculable, and economies of mass (scale, scope, and experience) are significant but not overwhelming. They exploit economies of scale and scope by placing big bets on technologies and facilities. They make incremental improvements in products and processes.

A company can participate in any of the four layers in a stack. Traditional oligopolists are companies by definition. Curatorial platforms may be nonprofits such as Wikipedia.org, but also corporations such as Facebook and InnoCentive. While some infrastructure organizations are owned by governments or municipalities, others are for-profit corporations, such as Amazon Web Services. Companies can even participate in communities as small ventures or venture capitalists, or indirectly by encouraging employees to contribute to projects such as Linux.

But what cannot be emphasized too much are the *differences* among these four types of activity. They require different skills and motives, present different financial profiles to investors, and need to be managed on different time horizons. A company can flourish in multiple layers—Amazon does it—but most organizations consistently underestimate the enormous challenges. Decades ago, in its evolution from mainframes to PCs, the computer industry moved from an oligopolistic to a stacked architecture. The Internet industry has had that architecture from the beginning, because the stacked architecture of the technologies served as a template for the stacked architecture of the institutions (corporate and noncorporate) that exploited them. The media industry is evolving painfully toward that structure. So is telecommunications. So is electrical power. So is transportation. So must health care. And every business

that impinges on these sectors, as supplier or customer, has a profound stake in this evolution.

Much of what is broken in today's economy stems from activities pursued with the wrong model.

Implications for Executives

IT IS FASHIONABLE (and correct) to assert that business leaders need to worry about disruption. But disruption takes very specific forms, and these forms are shifting. The disruptive impact of deconstruction—like that of low-cost technologies—is now widely understood, but the challenge of the very small, less so. And the challenge of the very large, hardly at all. Put them together and you pass from the familiar world of value chains to the world of platforms, ecosystems, and stacks. Extend that to the limits of ubiquity, and you enter the strange universe imagined by Borges.

So leaders need to focus on asymmetrical rivals and unlikely allies, on hackers and hobbyists, on rooftop solar panels and 3-D printers. They must also adapt their strategies to the possibility of shared infrastructure, to data that wants to be big, to the implacable embrace-and-extend bear hug of Google and Amazon and the National Security Agency. Conventional business models may be simultaneously too big and too small.

How should executives respond? Here are the four major drivers of the new industrial architecture and the key strategic imperatives for companies.

1. BIG DATA

Test your current analytics against the state of the art.


The field is moving so quickly that even well-versed companies can fall behind. There is currently a gold rush of new analytical methods: banks pricing mortgage collateral are adopting relational factor-graph techniques to predict the interdependence of adjacent property values; retailers focused on data-driven marketing are applying probabilistic graphical models to social-network data. Traditional spreadsheet methods are being applied at orders-of-magnitude larger scale, requiring new computer engineering even when the logic is unchanged. New data sources are becoming available. Families of problems trivial at small scale become noncomputable at large scale, so algorithms—successive guesses—substitute for closed-form solutions. Real-time computation replaces batch processing. Short cycles of experimentation and validation replace elaborate market tests. Organizations capable of all this will be ones in which business managers, programmers, and mathematicians talk each other's languages, where small teams iterate in fast cycles, where empirical validation counts for more than the judgments of hierarchies or senior executives.

Consolidate databases across the company.

Big data yields advantages from scope as well as scale, so siloed, business-unit-specific databases are quickly becoming antiquated. Data sets have value well beyond the silos within which they originate, but few companies can integrate their data across product lines or between online and offline channels. Tesco famously targeted promotions to members of its Clubcard loyalty program by developing an integrated understanding of buying patterns across households, time, and points of sale. Woolworths in Australia has used retail-shopping patterns to predict financial risk. It found that customers who drink lots of milk and eat lots of red meat are significantly better auto-insurance risks than customers who drink spirits, eat lots of pasta and rice, and fill their gas tanks at night.

Form partnerships to gain scale.


Given Tesco's head start, its archrival, Sainsbury's, faced long odds in trying to catch up by playing the same game. So it outflanked Tesco on scope. It formed Nectar: a loyalty card shared



with UK retailers such as BP, Homebase, and Argos—and operated by a third party called Loyalty Management Group. Consumers got the benefit of more points redeemable at more outlets, and retailers got the benefit of a wider set of behavioral data. The power of such aggregation lies in the million nonintuitive relationships between things like eating rice and driving safely. The value for sellers lies in more efficient promotion, and for buyers, in messaging that feels less like shrill coercion and more like helpful advice. Done with consideration for the consumer, this can be a win-win.

Manage data as a trustee.

Personal data collected by businesses cannot be treated as mere property, transferred once and irrevocably, like a used car, from data subject to data user. Data sharing will succeed only if the organizations involved earn the informed trust of their customers. Many such arrangements today are murky, furtive, undisclosed; many treat the data subject as a product to be resold, not a customer to be served. Those businesses risk a ferocious backlash, while their competitors are grabbing a competitive advantage by establishing trust and legitimacy with customers.



2. DECONSTRUCTION

Reorganize your business along its economic fault lines.

Define organizational units by their distinct competitive economics in their layer of the stack, and manage these units for stand-alone competitive advantage. Even if the strategy is to remain traditionally vertically integrated, this will give your managers a clear view of the threats they face and free them to compete as fiercely as any upstart. Never subordinate the competitiveness of one operation to the interests of another. Amazon functions at many different layers of a complex stack, but each part targets competitiveness on a standalone basis.

Look for opportunities to be the lateral aggressor.

Consider one example: the automotive and insurance industries are colliding. How and where a car is driven is the best predictor of the incidence and severity of accidents. For a few years now, innovative insurers such as Progressive have offered “black boxes” that track driving behavior and enable the company to undercut competitors in pricing policies for the best drivers. But cars are rapidly becoming computers on wheels for other reasons: the Ford Fusion contains 74 sensors, and each year’s model re-

ords and interprets more data: time and place, the identity and posture of the driver, seat belt usage, tire pressure, sharp braking, lane changes. All this data is uploaded to the mechanic and to services such as GM's OnStar. That means the OEMs will own the most detailed underwriting data, across all drivers (not just the self-selecting best), at zero incremental cost. The separate black box will disappear as the OEMs realize they can suck up all that data—and so much more—and use it to take the insurers out of the game. The OEMs have the opportunity to think “outside the black box” and become a lateral aggressor.

Identify where your value chain is most susceptible to lateral attack.

With their actuarial tables and even their black boxes rendered obsolete, how can traditional car insurers survive? The first (and hardest) step is to recognize the problem five years before it hits. First movers that acquire the lower-risk drivers will be able to hold onto them. In many countries, regulators will mandate that consumers have access to their own data, so insurers will not be out of the game, but rather competing on a level playing field. To win, they need to build advantage in other layers of the stack: the analytics that interpret the data, claims adjustment, cross-selling customer service. Even in countries where OEMs own the data, there will be major elements of the business in which they will have little interest. That suggests an ecosystem with alliances among insurer, network provider, and OEM. Players should begin to position themselves today.

The same challenges and strategies—for both aggressor and incumbent—apply for many businesses.

3. POLARIZATION OF ECONOMIES OF MASS

“Up-source” activities to a community.

Digital communities are able to perform many tasks cheaper and faster than companies can. Customers provide free reviews for Amazon and perform crowd-sourced technical support for Cisco and several telecommunications companies. They do it out of a mixture of altruism, ego, and self-advertisement. Innovation

contests with dollar prizes—such as GE’s Ecomagination Challenge, Netflix’s contest to improve its recommendation algorithm, or those posted on the InnoCentive platform—help companies accelerate the pace of innovation while decreasing the cost. The application programming interfaces provided by companies like Google and some telecommunications providers enable communities of entrepreneurs and programmers to create new applications quickly and cheaply by “mashing up” data streams. This drives users and metadata to the platform provider.

“Down-source” activities to shared infrastructure.

In mobile telecommunications, for example, there are significant scale economies at the bottom of the technology stack. In France, SFR and Bouygues Telecom have begun to share their infrastructure of towers and masts in lower-density service areas, allowing them to remove some 7,000 towers. Each company continues to compete with its own transponders. In the UK, carriers EE and Three share towers, masts, transponders, and backhaul, while larger rivals Vodafone and O2 have a passive sharing arrangement similar to the French plan. In Sweden, Telenor and Tele2 even share spectrum. In all these arrangements, competition is diminished in the lower layer of the stack, but the level playing field intensifies competition in the upper. There is additional complexity and some coordination costs in this kind of joint venture, but that is offset by the increased utilization of fixed assets. In the UK, these arrangements are forecast to save about £1 billion per year.

4. HOLISTIC, STACKED ARCHITECTURES

Curate a new industrial stack.

In light of evolving technologies, reevaluate your value added from first principles. To take just one example: imagine smart agriculture as a stack. Cheap, meshed sensors measure the temperature, humidity, and acidity of the soil; active repeaters embedded in agricultural machinery or in cell phone apps capture, aggregate, and relay the data; data services combine this local data with aggregate models of weather and crop prices; other services tap into their APIs to optimize planting, irrigation, fertilizing, and

harvesting. Farmers collect the data, share in the aggregation and pattern recognition, and follow prescriptions that give them a better yield on their crops. Such an ecosystem creates social and private value in both developed and developing economies. For large agribusinesses, this is a major opportunity that poses no challenge to the business model. But where farming is fragmented, these technologies scale beyond the reach of individual farmers. The opportunity is therefore wide open—to governments, NGOs, processors of fertilizer, and builders of agricultural machinery—to orchestrate a new industrial stack.

Many industries could be reconceptualized along these lines by participants with the necessary resources, strategic insight, and imagination.

Where you can't curate your own stack, seek advantaged roles in stacks curated by others.

Every company wants to be the master of its own fate, but not all have the scale and scope to be orchestrators. The “smart” home, for example, is a vision of how thermostats, motion detectors, lighting, home theater, door locks, appliances, phones, and tablets will act and interact intelligently. There are immense benefits in convenience, safety, and cost savings there, but adoption has been stymied by balkanized, overpriced systems that use different control pads and interfaces, run on different wired and wireless networks, and cannot talk to each other. Google, with its recent acquisition of Nest, and Apple, with its launch of HomeKit, are building stacked architectures for granular integration of the various subsystems of the smart home that will allow homeowners customization and increased efficiency. It is not clear how this battle will play out, but the implications for other players are evident and imminent: they must hedge their bets and focus on defensible niches.


For power utilities, to take one example, this is bitter medicine. Although energy savings is one of the biggest benefits of smart homes, the logic for vertical integration by the utility is weak because the price of power is just a number. Utilities have been conspicuously unsuccessful in their attempts to play orchestrator of the smart home. Their biggest advantage is in the field, in

installation, maintenance, and repair. They have a new opportunity to exploit their knowledge of grid behavior, neighborhood consumption patterns, and signals from smart-home devices to detect and anticipate mechanical failures in homes. This will broaden and deepen their relationship with customers, increase utilization of the field force, and ultimately reduce customer churn. They are better off fitting into a niche than trying to curate an ecosystem of their own.

Reshape regulation.

The logic of stacks has massive implications for the philosophy of regulation—and requires that both businesses and regulators think differently. Traditional metrics such as market power are insufficiently nuanced in an environment of polarizing economies of mass. Companies have a huge stake in how this thinking evolves, and they can and should influence policy in directions that favor efficiency at the bottom of the stack and open innovation at the top.

IN 1945, John von Neumann, one of the greatest mathematicians of the twentieth century, wrote a paper describing a “Turing machine” that made no distinction between its data and its instructions to process that data. This so-called von Neumann architecture became the design of the digital computer: treating data and code as one.



Just a few months later, Argentinian writer Jorge Luis Borges penned the one-paragraph story at the top of this article, recounting how a lost empire became its own map. Borges imagined reality becoming a description of itself. His map and reality, like von Neumann’s data and code, are indistinguishable.

How exquisite that these two extraordinary visions, one from a supreme scientist and the other from a supreme fabulist, were formulated almost simultaneously. Now, through three waves of digital disruption, technology is finally catching up with both. Executives in the next decade must chart their course through the labyrinth of Borges’ map.

About the Authors

Philip Evans is a senior partner and managing director in the Boston office of The Boston Consulting Group. You may contact him by e-mail at evans.philip@bcg.com.

Patrick Forth is a senior partner and managing director in the firm's Sydney office and the global leader of its Technology, Media, and Telecommunications practice. You may contact him by e-mail at forth.patrick@bcg.com.

