

GLOBAL MARKETS INSTITUTE

Goldman
Sachs

The Everything-as-a-Service Economy

The Goldman Sachs Group, Inc.

December 2018

Global Markets Institute
The Everything-as-a-Service Economy

Steve Strongin
Goldman Sachs & Co. LLC.

Amanda Hindlian
Goldman Sachs & Co. LLC.

Sandra Lawson
Goldman Sachs & Co. LLC.

Sonya Banerjee
Goldman Sachs & Co. LLC.

Deborah Mirabal
Goldman Sachs & Co. LLC.

Emma Campbell-Mohn
Goldman Sachs & Co. LLC.

The Goldman Sachs Group, Inc.
200 West Street
New York, NY 10282-2198

The Global Markets Institute is the public-policy and corporate advisory research unit of Goldman Sachs Global Investment Research. See the Disclosures at the end of this publication.

■ Preface

The global business landscape has changed dramatically over the past twenty years and continues to evolve, bringing the topic of *disruption* to the fore. Technological advancements have played a role in the accelerating pace of change we see across so many industries today – but while it is the most obvious driver, technology isn't the only catalyst for disruption.

In our view, there's a more complicated and more interesting dynamic at play. The ongoing quest for greater corporate efficiencies, enabled by the introduction of new technologies, has caused companies to adjust the way they operate, remodeling the front-office to the back-office and everything in between. In the process, companies are now leveraging more services provided by third-parties than they once did.

As a result, industries are reorganizing at a rapid rate, leaving companies – and their investors – with the challenge of navigating today's fast-moving business landscape. This is precisely the issue we aim to address.

This publication consists of two parts. The first part provides a theoretical framework to describe the way companies are evolving their operations. The second part addresses the business strategies we believe are likely to be successful in today's Everything-as-a-Service economy.

Our key takeaways are as follows:

First, the very nature of competition has changed as alterations to how businesses operate have eroded some traditional barriers to entry. In just about any market, a competitor – even a new entrant – can now scale quickly, with little to no capital and little to no staff. What's more, for a firm to become a disruptor it just needs to offer an incrementally better product or service than its competitors – and not even by that wide of a margin. Customers

can and will switch to new products and services to realize incremental benefits – and they'll do so for smaller gains than in the past as switching costs continue to decline. This is both an entrepreneur's dream and an incumbent's worst fear, since firms that fall behind risk rapid displacement.

The old competitive model typically pitted relatively evenly-matched giants against each other. For example, Macy's and Gimbels (a now defunct department store) once operated within blocks of each other. They competed for the same customers, using roughly the same operating structures and selling largely the same goods. As another example, Ford, General Motors and Chrysler were once centered in Detroit, with brands that paired off against each other in nearly every automobile market segment and in nearly every part of the country.

But today's competitive landscape looks quite different – it's not always obvious what a firm really does just by looking at what it sells. Furthermore, a customer in one area may be a competitor in another. For example, Walmart and Amazon compete in many retail categories, yet Amazon's edge in retail is largely driven by its e-commerce logistics platform. Apple, Samsung and Google, as another example, sell comparable mobile phones to similar customers, yet Samsung also supplies key components to Apple for the iPhone and leverages Google's Android operating system within its own devices.

Second, technological progress and increasing competition have driven businesses to reorganize and to re-engineer their operating structures in an ongoing process of disentanglement. Through this process each business function is refined and standardized to better leverage a company's capital and areas of expertise.

Third, we use the phrase the *Everything-as-a-Service economy* to describe how disentanglement has remade the business ecosystem. From technological infrastructure to manufacturing to delivery and after-market services, companies no longer need to

do it all. Instead, they can rely on third-parties for many of their needs. As a result, firms can better direct their resources to their areas of competitive advantage – concentrating their investments rather than diluting them by spreading them thinly everywhere. This is the focus of the first part of this publication, which is entitled “enabling disruption.”

The second part of this publication (which can be read independently from the first) is an examination of how companies can successfully navigate the Everything-as-a-Service economy by precisely identifying and investing in their sources of competitive advantage. To that end, we define four key drivers of economic advantage that we think companies can exploit to achieve long-run success: economies of scale, economies of scope, economies of fit and economies of learning.

While economies of scale and scope are well-understood concepts, the Everything-as-a-Service economy creates new possibilities both *to refine the focus of* and *to expand* the potential markets that businesses operating in these areas can leverage, and in ways that would not have previously been thought possible. In comparison, the concept of economies of fit (based on the economic concept of monopolistic competition) examines how to use today’s social and shopping platforms, which have created more flexible notions of “communities,” to find, to create and to supply new markets.

Lastly, we discuss learning companies, which leverage business models that are based on today’s proliferation of data. In particular, we develop a new framework that we refer to as “the learning curve” to describe when data can serve as a source of sustainable competitive advantage, in contrast to the more technical (but economically neutral) question of when new data-driven technologies can be used.

This new framework leads to a four-part test companies can take to determine whether data can serve as a source of competitive advantage. First, are there sufficient data to analyze? Second, are

The Everything-as-a-Service Economy

the insights gained from such data analysis novel enough to create significant value? Third, is the implementation of those insights complex enough to prevent competitors from simply copying the approach? And fourth, are the data scarce enough that a competitor cannot repeat the same analysis with relative ease?

If each of these questions elicits an affirmative response, building a sustainable competitive edge through data is possible. However, more often than not, this is unlikely to be the case, which means that data-based strategies tend to be a cost of entry rather than a source of sustainable competitive advantage and that robust second-mover strategies may be more cost effective than first-mover ones.

In the end, we find that the Everything-as-a-Service economy allows companies to better direct their efforts to their core areas of competitive advantage. Doing so well can pay enormous dividends, including creating stronger and more secure market positions, while failing to do so well can end in quick failures¹.

¹ The ideas contained within this publication benefited from extensive conversations with Goldman Sachs research analysts about disruption in their industries.

■ Table of Contents

Preface	5
Part 1: Enabling Disruption	11
The accelerating pace of disruption: the high-level theory	12
Five key observations about disentanglement	15
Disentanglement and disruption: a deeper dive	18
The decline of traditional competitive barriers	19
Case study 1: Ford	20
Case study 2: McDonald's	23
Case study 3: firmware (IBM)	26
Case study 4: user standards – Windows and iOS	28
Case study 5: ride-hailing services	30
Case study 6: ISO 9000	32
Case study 7: Netflix	33
Regulatory disentanglement	36
Competition policy	37
Part 2: New Rules for New Business Models	39
Focus on core areas of strength	39
New business models	40
Platform companies: exploiting economies of scale	44
Case study 8: Amazon's retail services business	47
Case study 9: Berkshire Hathaway	50
Case study 10: biopharma	51
Servicer companies: exploiting economies of scope	54
Case study 11: ADP in payrolls	56
Case study 12: payment processors	57

Case study 13: Qualcomm	59
Organizer companies: exploiting economies of fit	59
Organizers and the role of communities	61
Case study 14: lululemon athletica	64
Learning companies: the benefits of information	66
Learning companies and “the learning curve”	68
Data-smart companies	72
Data-asset companies	73
Data-feedback companies	76
Data-network companies	80
Concluding Thoughts	83
Bibliography	85
Disclosures	89

Part 1

Enabling Disruption

In this section, we discuss *how the structure of firms has evolved* and why today's operating environment, where nearly any business function can be sourced as a service, is a divergence from the past. We hone in on the ongoing process of *disentanglement* as the driving force behind disruption using a series of illustrative company-specific case studies. We also explain how this process has led to the modern business environment that we believe can best be described as the Everything-as-a-Service economy.

As the case studies will show, while firms' structural changes may seem sudden, they are reflective of a gradual progression over time. To a large extent, the push for ever-more corporate efficiencies, coupled with and accelerated by the introduction of new technologies, has propelled this reorganization. Thus, the Everything-as-a-Service economy reflects an *evolution* rather than a *revolution*, although it may not always feel that way to the companies that are disrupted.

These forces also help explain why the *pace* of disruption has increased. As some layers of a firm's production process become disentangled, and other aspects of the process become more

standardized, it becomes easier to disentangle additional layers. In turn, this means that disentanglement speeds up and becomes cheaper to do over time. Put another way, as companies become more efficient, these gains not only encourage *further* efficiency gains, but achieving additional efficiencies becomes easier and less expensive to do over time.


In the pages that follow, we delve further into this theory to explain the accelerating pace of disruption occurring across so many industries today.

The accelerating pace of disruption: the high-level theory

We believe an ongoing process of *disentanglement* is a critical factor behind increasing disruption. Over time, firms have gradually disentangled their production stacks – meaning the entire process of bringing a good or service to market – in ways that have improved their operational efficiency. This process has been enabled by technological advancements and, increasingly, by on-demand business services. The basic idea is that traditional operating structures for businesses, such as vertically-integrated ones, have largely broken apart, as Exhibit 1 shows.

Exhibit 1: The new corporate production stack

Older entangled operating structures have been replaced



In the past, firms' production stacks were often consolidated and vertically integrated because this structure allowed for greater operational control and efficiency. A single company was likely to manage nearly every part of the process of bringing its own goods or services to market. This may have included investing in and managing the collection of inputs and tools, the manufacturing or production of goods, the provisioning of services, the sales process and back-office operations – all on a relatively granular basis and largely from start to finish. This older model is represented by the “before” portion of this exhibit. Over time, as new technologies and third-party services are brought into the fold, less of each layer needs to be handled by firms internally or on a bespoke basis. As a result, many of these functions thin and standardize and thus disentangle from the adjacent layers. As a result, the old production stack, which was one thick layer, has been replaced by a new one with two thinner layers, which is represented by the “after” portion of this exhibit.

Source: Goldman Sachs Global Investment Research


The resulting pieces then separate into two types of layers: those that are “consolidating” layers and those that are “fragmenting” layers. **Consolidating layers tend to capture activities that benefit from economies of scale.** These activities are often capital-intensive and not particularly innovation-driven, and given the emphasis on economies of scale, smaller players that focus on these types of activities are likely to be absorbed by larger firms or may simply go out of business.

In contrast, **fragmenting layers tend to capture activities that are characterized by diseconomies of scale.** These activities are often innovation-driven with a narrow market focus and are not particularly capital intensive, but typically emphasize *expertise*. The smaller market segments that emerge at the fragmenting layer may be easier to defend because they often require specialization, and they may provide higher returns given that they tend to have relatively limited capital requirements.

As each layer separates from the adjacent ones, it becomes free to reach its optimal structure and scale without the constraints of having to conform to, operate with or fund the other layers or parts of the stack. As new technologies and third-party services are brought into the fold, less of each layer needs to be handled by firms internally or on a bespoke basis, which results in what we refer to as the “thinning” and standardizing of these layers.

Ultimately, today’s Everything-as-a-Service economy is the end result of the ongoing process of disentanglement. The existence of standardized layers allows third-parties with particular areas of expertise to provide most production stack related services on a stand-alone basis. And, as a result, many traditional competitive barriers have eroded or been eliminated – such as the benefits of vertically integrating to operate at scale – in part because it’s easier to leverage other providers to start or operate a business with far less investment (such as in capital and people) than in the past. See Exhibit 2.

Exhibit 2: How disentanglement reshapes industries and has resulted in the Everything-as-a-Service economy



As firms' production stacks break down, the resulting pieces separate into two types of layers: those that are "consolidating" layers and those that are "fragmenting" layers. Consolidating layers tend to capture activities that benefit from economies of scale. These activities are often capital-intensive and not particularly innovation-driven, and given the emphasis on economies of scale, smaller players that focus on these types of activities are likely to be absorbed by larger firms or may simply go out of business. In contrast, fragmenting layers tend to capture activities that are characterized by diseconomies of scale. These activities are often innovation-driven with a narrow market focus and are not particularly capital intensive, but typically emphasize expertise. The "before" portion of this exhibit shows an industry before disentanglement, while the "after" portion shows how the industry is reshaped through disentanglement, highlighting the resulting consolidating and fragmenting layers.

Source: Goldman Sachs Global Investment Research

Five key observations about disentanglement

- **First, disruption is now more likely to result in the displacement of *specific functions within a firm* – rather than the displacement of a well-positioned firm as a whole.** This is because well-positioned firms can rent what's new to enhance their existing areas of strength or to expand into new areas. These dynamics help to explain why, despite the pick-up in the

rate of disruption, the Everything-as-a-Service economy has not ushered in significant changes in terms of market leadership positions across many sectors. In other words, companies that have historically held top market share positions may still do just that – across many parts of the broader economy. Another contributing factor to this dynamic is that the firms that have emerged over the last two decades have largely entered market niches that did not exist before. To illustrate this point, consider which established companies Google, Facebook, or Stripe have disrupted, other than a few early entrants?

- **Second, from an economic standpoint, fragmentation is just as important as consolidation.** The creation of finely-tuned products is typically done by specialty firms on a narrow basis, which largely exist in fragmenting layers. In comparison, the activities that gain efficiencies from being done at scale tend to exist in consolidating layers. In our view, neither one can exist without the other. Instead, they interact in mutually reinforcing ways.
- **Third, specialization is the new norm and many firms are likely to continue to refine their focus, though the nature of specialization may look quite different relative to the past.** In the new business environment, successful firms often do less (not more) than their predecessors. Again, this is because in an environment where third-party services are used more often, firms can narrow their focus quickly. From the outside, many larger successful firms today do appear to have numerous business lines. But, when these firms are examined more closely – with an eye toward the areas where they are uniquely successful – a different story emerges: their differentiated business lines tend to be quite limited. Furthermore, by examining what is outsourced versus what is done in-house, it's easier to see that successful firms often have relatively narrow areas of focus. Consider Amazon's e-commerce business as an example of narrow specialization: Amazon covers an astronomical number of SKUs, but it is important

only in categories of items that can be shipped in a box and delivered in a day or two.

- **Fourth, a firm's area of competitive advantage may not necessarily be apparent or even linked directly to the products or services it sells or to its overall structure, and competitors can also be customers; this is a fundamental feature of the Everything-as-a-Service economy.** Unlike in the past, businesses that seem to compete with each other today may in fact have little to no overlap in terms of their actual areas of competitive advantage; they may also rely on each other as customers. The Samsung, Apple and Google smartphone example that we touched on at the outset of this publication underscores this dynamic. What's more, many firms are likely to have hybrid business models – meaning more than one area of focus or expertise. Such hybrid models are neither better nor worse than pure-play business models, as long as the combination is mutually reinforcing and beneficial.
- **Fifth, existing regulatory and competition policies and frameworks are likely to miss the mark in the new economy.** As we will discuss, consolidating layers collect capital and shed jobs, while fragmenting layers shed capital and collect jobs. As a result, in the new Everything-as-a-Service economy, financial and labor policy that favors physically large organizational structures over smaller ones is unlikely to help promote business growth and formation. And, since what a company does may not be obviously tied to what it sells, to be effective, regulation should focus more on *types of activities* rather than on *types of organizations*. Lastly, given the rise and importance of smaller, narrowly-focused businesses, competition policy may fail to achieve its aim. These are topics we explore in greater detail at the end of this section.

Disentanglement and disruption: a deeper dive

To understand disentanglement, it is important to consider how firms' production stacks have evolved over time. To reiterate, the production stack consists of each of the elements necessary to bring goods or services to market – design, manufacturing, distribution, advertising, human resources and payrolls, among others.

In the past, firms' production stacks were often consolidated and vertically integrated, because this structure allowed for greater operational control and efficiency. There were also few efficient alternatives to vertical integration, particularly for large firms producing and selling products over wide geographies or across different markets. As a result, a single company was likely to manage nearly every part of the process of bringing its own goods or services to market. This may have included investing in and managing the collection of inputs and tools, the manufacturing or production of goods, the provisioning of services, the sales process and back-office operations – all on a relatively granular basis and largely from start to finish.

Over time, as we noted earlier, many of these functions have thinned and standardized and thus disentangled from the adjacent layers. As a result, the old production stack, which was one thick layer, has been replaced by a new one with two thinner layers, which has a much flatter, more dispersed structure. This is largely because *the economic benefits from operating each component of the stack at its own optimal structure often make them more efficient in the aggregate than the entangled stack once was.*

As we have said, disentanglement isn't new. *What is new* is that the results have accumulated and the process has accelerated to such an extent that businesses of nearly any size can now more easily and inexpensively offload to third-parties many more non-core functions than was possible in the past.

In the early stages of disentanglement, the process itself was

costly; thus it only made sense to undertake when the potential economic gains were significant. As a result, disentanglement was rare, narrow in scope and slow. But as an increasing number of layers of the stack standardized over time, the cost of re-engineering declined, which is at least partly because incremental disentanglement becomes easier to do as expertise is gained. This has allowed for faster and wider disentanglement and explains why each new round of disentanglement is economically worthwhile, even as the economic gains decline.

As a result, each new round of disentanglement has been faster, causing more and more change to unfold for smaller and smaller economic gains. This dynamic increases the pace of change; it sits at the core of modern business disruption and helps to explain why disruption is likely to continue – and may even accelerate further.

The decline of traditional competitive barriers

Disentanglement has eliminated or reduced many traditional competitive barriers because third-party services are readily available, which often makes starting or running a business a less capital-intensive endeavor. With these resources in hand, a new product or service can take off quickly, and can become disruptive with astonishing speed. The structural and competitive implications can be significant for companies and their investors.

The case studies we explore next touch on the key developments in the process of disentanglement over time, and are meant to be illustrative rather than comprehensive. We start by considering the production process in the narrowest sense – meaning for a single firm. But we quickly move to considering how related entities in a vertically organized chain might behave, and then assess ecosystem-based efficiencies among otherwise independent firms today.

With case study 1, we don't aim to present the first case in the history of the process. Rather, we begin with an instance many readers are likely to find familiar.

Case study 1: Ford

When studying how firms' operations have been restructured over time to achieve greater efficiencies, Ford Motor Company is often cited as the archetypal example. However, in the context of the Everything-as-a-Service economy, *Ford isn't actually a particularly good example* of the kind of re-engineering businesses can do today. This is simply because the baseline level of operational inefficiency that was the norm when Ford implemented the assembly line was far greater than what most companies experience today.

Thus, Ford aimed to make process improvements in manufacturing to gain efficiencies, but it did not truly re-engineer its business structure in the way companies do today. The firm's operations remained largely vertically integrated – increasing in scale and scope over time – which is far from the thin operating structures that are now more of the norm.

Let's consider Ford's early business structure and approach with this in mind. Each automobile Ford once produced was created as a single unit, with groups of skilled employees building each vehicle by hand. Production was limited, as was adoption since relatively few people could afford these expensive vehicles. In 1913, in an effort to increase production and to lower prices, Ford made an aggressive push in the manufacturing of its Model T vehicles with the introduction of the assembly line.

As one example, rather than having a single person assemble a magneto – a component essential to the engine – from start to finish, the process was divided into nearly 30 distinct tasks, each of which could be handled by a different employee along the production line. As workers became more efficient at completing their individual assignments, Ford was able to reduce the average build time for each of these magneto units by more than 50%. In practice, this type of change in the manufacturing process was based on ensuring better matching between the labor pool and a


specific skill, as opposed to an actual re-modeling of the production stack.

Eventually, the company tried to optimize an increasing number of steps in the manufacturing process, in a similar effort to improve employee-skill matching. The results were beneficial: Ford was able to reduce the total production time for its Model T vehicles from more than 12 hours to around one and a half hours via these efforts. But Ford's production operations remained labor-intensive and largely entangled, with little technological automation. In other words, as Exhibit 3 shows, Ford's production stack was still very much vertically integrated and it had not truly re-engineered in the way firms can today.

Between 1913 and 1927, when the Model T was discontinued, Ford was able to increase the number of vehicles it could manufacture, eventually surpassing 15 million units given the process improvements the company made. Over the same period, the price of the Model T declined by more than 60% as some operational efficiency and the benefits of scaling-up some activities began to take effect. Nevertheless, Ford remained far from the types of re-engineered manufacturing companies we see today (Costa, Apr 2016).

Exhibit 3: Ford's use of the assembly line was an early form of disentanglement

Process improvements yielded meaningful benefits, but Ford's production stack remained largely entangled



The “before” portion of this exhibit shows Ford’s early production stack. At the time, each automobile Ford produced was created as a single unit, with groups of skilled employees building each vehicle by hand. Production was limited, as was adoption since relatively few people could afford these expensive vehicles. In an effort to increase production and to lower prices, Ford made an aggressive push in the manufacturing of its Model T vehicles with the introduction of the assembly line. Ford’s process improvements in manufacturing yielded benefits, but the firm did not truly re-engineer its business structure in the way companies do today. As the “after” portion of this exhibit shows, the firm’s operations remained largely vertically integrated and were far from the thin operating structures that are now more of the norm.

Source: Goldman Sachs Global Investment Research

Case study 2: McDonald's

McDonald's serves as a better early example of modern disentanglement. Like Ford, McDonald's disentangled its production processes, but – unlike Ford – it also *reorganized its operational and capital structures as well*. By disentangling both its processes and its organization, McDonald's more closely resembles the type of disentanglement we see today.

Around the time McDonald's began operating, food preparation was done within restaurants with little to no automation. What's more, the high turnover and local nature of food-service labor and real estate kept the restaurant model quite local, as one might expect. Over time, however, a variety of technical advances allowed food items such as uncooked fries and hamburger patties to be prepared offsite, frozen and then delivered to restaurants, where these foods were then prepared and distributed to local customers.

Starting with its very first restaurant, McDonald's had the specific aim of addressing problems associated with the then prevalent drive-in fast-food model, where service could be slow and inefficient and the quality of the food itself fluctuated. To accomplish these goals, the founders limited the restaurant's menu, implemented a somewhat Ford-like assembly-line system for food production, leveraged available technology and automation where possible – electric milkshake mixers, for example – and built their own customized tools.

But the real organizational change occurred when McDonald's shifted to a franchise structure beginning in the mid-1950s. By adopting this structure, the firm could market and advertise broadly (even on a global basis) and centralize the development of food preparation technologies, while procuring standardized ingredients on a regional basis and allowing restaurants to continue to operate locally. Perhaps most importantly, franchising changed the company's *financing model* – enabling funding and labor to stay local even as the company itself, its brand and its

The Everything-as-a-Service Economy

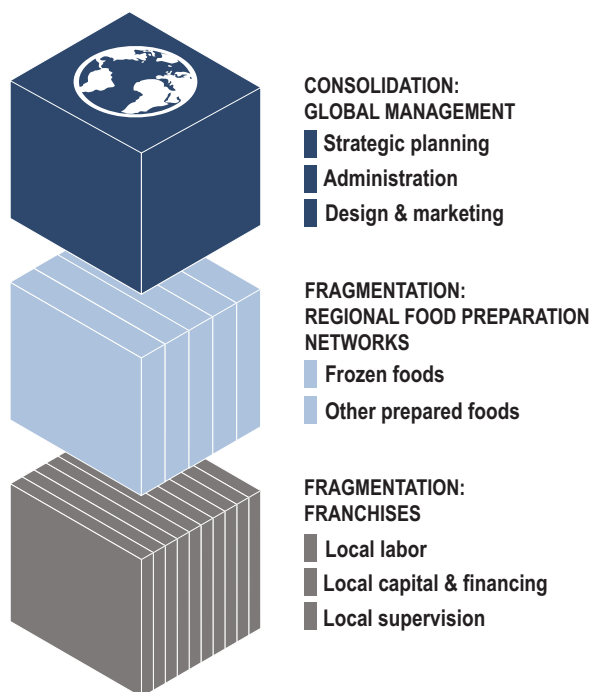
food went global. Today, more than 90% of McDonald's locations are franchises.

As shown in Exhibit 4, the three-layer system McDonald's created – principally via franchising – consisted of corporate global management, design and marketing; quasi-independent networks of food preparation on a regional level; and local franchises with local capital, labor and supervision. This three-layer system was far more efficient at each level than nearly any other food organization had been before. In this way, McDonald's did more to structurally reorganize itself than did Ford, although it isn't often cited as a "base case" example in this regard.

Not only does the McDonald's business model continue to exist today, but it was replicated by competitors and it even underpins some of today's modern "sharing" platforms, as in the case of ride-hailing services like Uber, which we discuss later.

Exhibit 4: The McDonald's production stack is characteristic of the Everything-as-a-Service economy

A three-layer system including consolidating and fragmenting layers



The McDonald's production stack more closely resembles the type of disentanglement we see across firms and industries today. As this exhibit shows, the three-layer system McDonald's created over time – principally by adopting a franchising model – consisted of corporate global management, design and marketing; quasi-independent networks of food preparation on a regional level; and local franchises with local capital, labor and supervision. This three-layer system, inclusive of consolidating and fragmenting layers, was far more efficient at each level than nearly any other food organization had been before.

Source: Goldman Sachs Global Investment Research

Case study 3: firmware (IBM)

IBM's implementation of firmware in the early 1960s was in many ways the first meaningful demonstration of modern disentanglement. It led to some of the broader business ecosystem dynamics we see today.


The original impetus for IBM's introduction of firmware was to allow its customers to more easily upgrade their computing hardware. Until IBM's System/360 computers (S/360) and the introduction of firmware, if customers wanted to upgrade their hardware, they also had to uplift the software as well – often at a high cost – which hampered hardware sales.

By standardizing how software accessed hardware, firmware allowed the same software to be used across an entire series of IBM mainframes. This meant that hardware changes could be made without also necessitating software investment, making it easier for IBM to sell hardware upgrades. It also made it easier for corporate computer users to invest in software.

Perhaps the biggest change ushered in by firmware was that it made it far more economically sensible for third-party vendors to begin developing software for corporate customers. This was a major step in creating the software industry as we know it today. As it turns out, hardware is a natural consolidating layer (meaning it benefits from economies of scale). In contrast, software is a natural fragmenting layer (meaning that development occurs on a narrow basis associated with specific use cases, reflecting the differing needs of various users). See Exhibit 5.

Exhibit 5: IBM: the computer industry before and after firmware

How firmware enabled consolidating and fragmenting layers to emerge in the computer industry



As the “before” portion of this exhibit shows, prior to the introduction of firmware with IBM’s S/360 machines, the computing industry was largely entangled, with hardware and software tightly coupled. But, by standardizing how software accessed hardware, firmware allowed the same software to be used across an entire series of IBM mainframes, which meant that the two assets could be addressed independently. This separation made it easier for IBM to sell hardware upgrades. But, as the “after” portion of this exhibit shows, perhaps the biggest change ushered in by firmware was that it made it far more economically sensible for third-party vendors to begin developing software for corporate customers. And, as it turns out, hardware is a natural consolidating layer (meaning it benefits from economies of scale). In contrast, software is a natural fragmenting layer (meaning that development occurs on a narrow basis associated with specific use-cases, reflecting the differing needs of various users). Ultimately, the introduction of firmware was a major step in creating the software industry as we know it today.

Source: Goldman Sachs Global Investment Research

Over time, as we mentioned earlier, other advancements in technology enabled further organizational disentanglement to unfold, and in many ways even accelerated the trend of disentanglement. These include, as examples, the introduction of the technologies that underpin software-as-a-service and cloud-computing capabilities, which reflect the complete

decoupling of hardware from software that has gradually occurred – a process that began with firmware (Bellini, Jan 2015).

Case study 4: user standards – Windows and iOS

As well, over time, advancements in user interface technology – the means by which users interact with software and hardware – became a driving force underpinning growth in the personal computing industry. To that end, the release of Microsoft's Windows 95 operating system was essential since it meaningfully simplified and standardized the user interface associated with interacting with personal computers.

As is well-known, Windows 95 included features like a “start” menu, which listed the software applications resident on the machine, as well as a taskbar with basic features (showing the time and the date, for example) that quickly became the standard in personal computing and are iconic even today. While these might seem like small technological changes, they actually reflected significant advancements in software graphic design, serving to make personal computers simpler and more intuitive to use. These changes also set the stage for the erosion in switching costs that characterizes the Everything-as-a-Service economy today.

What's more, with Windows 95, Microsoft effectively kicked off the process of disentangling the functional (or technical) layer of software from the user interface; while not new technology, Windows 95 brought such interfaces into the mainstream. This process may have reached its fullest expression with Apple's user interface standards developed in the 2000s, as well as in Apple's App Store, which we discuss further next. This trend is also evident in the development of HTML5 and other approaches to web development and design we see today.

While Microsoft set the standard for simplified user interfaces, Apple introduced further innovation through its mobile

touch-screen devices, for example with the natural scroll feature. This approach to scrolling required users to scroll up to move down a page, or to scroll down to move up the same page; while the action may be physically and functionally intuitive, the written version certainly isn't. While natural scroll was originally designed for its touchscreen devices (the iPhone and iPad), the firm incorporated the technology into its traditional line of computers in 2011, reflecting a widespread consumer shift in favor of intuitive design that has continued (Cabral, Jun 2014).

By producing devices that are functionally intuitive for users to operate, Apple helped to lower barriers to entry in software, as well as switching costs, contributing to the increasing disruption that is now the norm.

Improved user interface standards also made it possible for firms to begin to enlist users to participate in their "production" processes. In the past, firms would have shied away from providing users with direct access to their information systems since this not only represented a security risk but was also generally inefficient – with low take-up rates and high costs (in the form of training and monitoring). It also resulted in rigid systems, since those who did learn to use the system and were "good" customers would need to be retrained to address any changes. In rare instances, however, this tactic could result in product "stickiness," serving as a competitive barrier.

Ultimately, new user interface standards created new norms for how users interact with software – with intuition becoming an important underpinning element. Much as firmware made it easier for users to switch between hardware platforms, user interface standards made it easier for users to move between software systems. As a result, there was significant growth in software development as lower switching and development costs allowed increasingly narrow products to be widely adopted.

Compare the past to today. By layering a modern user interface on top of its information systems, a firm can now provide users with direct systems access – albeit at an abstracted level – with far less security risk and much greater efficiency. What’s more, as user interfaces have standardized, software has evolved to accommodate users’ expectations that these applications can be adopted with little to no training. Consider the prevalence of self-service user interfaces in travel booking applications, banking applications and e-commerce sites – just to name a few examples.

Case study 5: ride-hailing services







Likewise, modern ride-hailing services are beneficiaries of the user interface improvements we have described. In effect, modern user interfaces underpin ride-hailing companies’ capital-light operating models.

Traditional private car services have historically been limited by the extent to which each operator could invest in owning and maintaining a fleet of vehicles or vet a cadre of steady drivers with their own vehicles, with all of the fixed costs and complexities associated with employing drivers, such as insurance requirements. These factors inherently limited these firms’ scale. To illustrate this point, consider that some of the largest private car services in New York City, where the industry is well-established, are estimated to have fleets with fewer than one thousand vehicles each.

As Exhibit 6 shows, modern ride-hailing services – like Uber, Lyft and Didi Chuxing – have been able to overcome these limitations. These services typically rely on drivers sharing their privately owned vehicles with passengers in exchange for income, with the ride-hailing company providing the technology and other required business infrastructure to users and drivers through a clean user interface.

Exhibit 6: Ride-hailing players

A snapshot of the industry

 <ul style="list-style-type: none"> • Founded: 2009 • Global operator • 15mn daily trips • 10bn completed trips • 75mn riders • 3mn drivers 	 <ul style="list-style-type: none"> • Founded: 2012 • US, Canada • 1bn completed trips • 23mn passengers • 1.4mn drivers
 <ul style="list-style-type: none"> • Founded: 2012 • Southeast Asia • 2.5mn daily rides • 2bn completed rides • 1mn driver partners 	 <ul style="list-style-type: none"> • Founded: 2010 • India, Australia, UK, New Zealand • 9mn vehicles • 10mn driver-partners
 <p>Didi Chuxing</p> <ul style="list-style-type: none"> • Founded: 2012 • China, Australia, etc. • 30mn rides per day • 450mn users • 30mn drivers 	 <p>Careem</p> <ul style="list-style-type: none"> • Founded: 2012 • Middle East and North Africa • 300mn completed trips • 30mn+ users • 1mn drivers

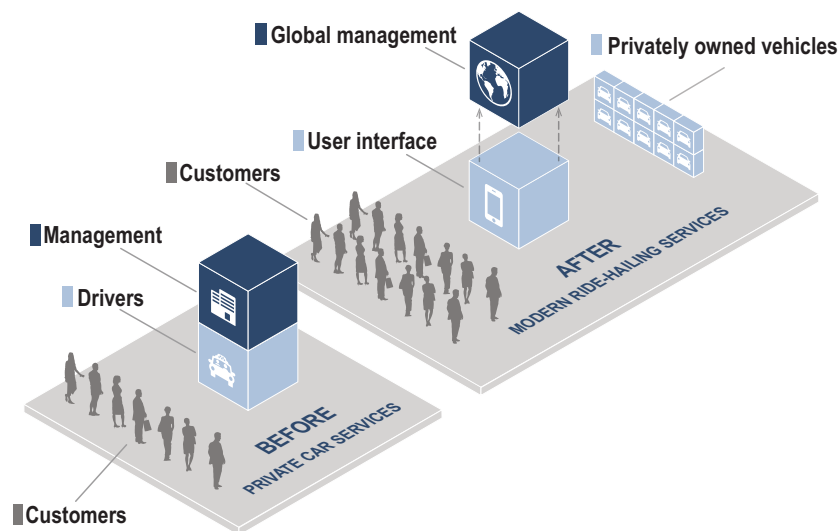
Source: Company data, Goldman Sachs Global Investment Research

These firms also rely on users' willingness to leverage their software applications to reserve rides and to rate drivers (who, in turn, can also rate their passengers), rather than offering a centralized reservation service. The rating system allows drivers to decide which passengers they'd like to provide their services to, and also protects customers by screening out drivers with consistently low ratings much more efficiently than if these protections were managed centrally.

This model – which involves drivers selecting their customers, being responsible for providing their own vehicles and managing the related expenses – is tantamount to a sort of “hyper franchising system,” as Exhibit 7 shows. The breadth that this model enables is significant and well beyond what McDonald's was able to achieve, as we discussed in case study 2. Uber, for example, is associated with tens of thousands of drivers in New York City alone and three million drivers globally (Burgstaller, May 2017).

Exhibit 7: The ride-hailing model - before and after disentanglement

How the industry has benefited from a hyper franchising system



Traditional private car services - shown in the “before” portion of this exhibit - have historically been limited by the extent to which each operator could invest in owning and maintaining a fleet of vehicles or vet a cadre of steady drivers with their own vehicles, with all of the fixed costs and complexities associated with employing drivers, such as insurance requirements. These factors inherently limited these firms’ scale. Modern ride-hailing services – depicted by the “after” portion of this exhibit – have been able to overcome these limitations. These services typically rely on drivers sharing their privately-owned vehicles with passengers in exchange for income, with the ride-hailing company providing the technology and other required business infrastructure to users and drivers through a clean user interface.

Source: Goldman Sachs Global Investment Research

Case study 6: ISO 9000

While the last few examples have largely centered on disentanglement driven by technological advancements, there have been other drivers as well. In the manufacturing industry in particular, ISO 9000 quality assurance standards were an early enabler of the Everything-as-a-Service economy, improving production processes and yielding significant gains. By obtaining ISO 9000 certification – which was in and of itself a costly endeavor – firms could verify that their operations produced


sufficiently standardized and high-quality output that they could be relied upon by other firms looking to use these vendors' output in their own production stacks.

In effect, these standards gave manufacturing firms the ability to begin consolidating and fragmenting layers of their production stack by outsourcing non-core functions. The result was a net improvement in their overall efficiency and productivity – despite the high initial investment costs necessary to ensure compliance with these standards. Thus, in some ways, ISO 9000 standards did for manufacturing (and, over time, for other industries) what accumulated software standards did for the computer industry.

Case study 7: Netflix

Netflix is an example of what disentanglement has wrought, as well as the emergence of the Everything-as-a-Service economy. By taking advantage of services provided by other firms, Netflix is able to distribute its proprietary content (as well as others' content) to a global marketplace, easily and more efficiently than it could have in the past. This is despite the fact that the firm does not actually create, warehouse or deliver much of the streaming media it sells. Instead, today, it is largely an organized collection of others' goods and services – from the bulk of its content library to much of its IT infrastructure, as Exhibit 8 illustrates.

Exhibit 8: Netflix is an example of a nearly virtual company



Netflix is an example of what disentanglement has wrought, as well as the emergence of the Everything-as-a-Service economy. By taking advantage of services provided by other firms, Netflix is able to distribute its proprietary content (as well as others' content) to a global marketplace, easily and more efficiently than it could have in the past.

Source: Company data, Goldman Sachs Global Investment Research

In the past, media companies were often viewed as natural monopolies, particularly since the high cost of delivering a complete bundle served as a significant barrier to entry. Today, in the Everything-as-a-Service economy, Netflix has been able to rent the bulk of its operational services at a sufficiently low cost that it is able to both offer a low-priced subscription service and focus its resources on acquiring a range of content for a variety of customer segments. While it is not yet clear whether Netflix will be able to collect enough viewership groups to become profitable on a cash basis, Netflix's capital-light and focused operating model has allowed it to become a major global competitor, despite the fact that the firm lacks certain content – like sports coverage – that was once viewed as essential to participating in the space (Borst, Oct 2015) (Terry, Jan 2017). See Exhibit 9.

Exhibit 9: The evolution of Netflix

Narrowly focused, capital light companies aren't necessarily small

THEN May 2002 (IPO)	
600,000+ subscribers	\$75mn+ revenues (CY 2001)
18mn+ daily content recommendations	11,500+ titles: movie, TV and other filmed entertainment
95%+ of titles selected by subscribers	3.3mn+ DVDs in stock
~70% of activity from back catalogue titles	In April 2002: 5.8mn+ DVDs shipped and received
	Standard plan: \$19.95 per month
\$22mn+ DVD amortization (CY 2001)	
260 full-time employees	

NOW 2017	
117mn+ streaming memberships	\$11.7bn+ revenues
Streaming service launched in 2007	First original series launched in 2013; 1000+ originals by the end of 2018
140mn+ hours of TV and movies per day	International expansion started with Canada in 2010; now available in 190+ countries
90%+ of customers watch original programming*	Global average monthly revenue per streaming customer: \$9.43
\$6bn+ streaming content amortization	
5,400 full-time employees	

(*) based on management's public comments.

Source: Company data, Goldman Sachs Global Investment Research

Regulatory disentanglement

As the theory and the examples we have explored thus far help to demonstrate, the changes ushered in by disentanglement are generally constructive in terms of their impact on industries and consumers, and more changes are likely to come. This is true on many levels, but particularly because the new economy is driven by the need to improve efficiencies across the economy, while also better fitting the needs of consumers.

But as we have said, *economic efficiency is not the only driver of the ongoing process of disentanglement*. For example, in some industries, regulatory or tax considerations can determine the boundaries of disentanglement rather than underlying economics. Finance provides a useful example in this regard, as with new payment systems, for example, where new entrants are providing more efficient services to consumers, but in many cases the boundaries are actually regulatory.

The often discussed non-bank versus bank split in lending may be considered a form of regulatory arbitrage. Non-banks that act in some ways as banks can be similar to a modern category of companies that serve as matching agents in niche markets (what we will next call “organizer companies,” which are associated with economies of fit), though in this case part of the reason non-banks may be structured this way is to avoid regulatory constraints. For example, much of the subprime lending market has this structure, where firms’ structures may allow them to avoid regulations (for example, providing New York consumers payday loans from western states to avoid usury laws). We also see this dynamic in the geographic organization of European banks. National considerations often prevail over economic efficiency both from the public and private perspectives. Medical care has some of the same obvious inefficiencies arising from artificially created boundaries due to licensing and regulatory boundaries. These organizational structures typically decrease

rather than enhance both fit and efficiency — reducing the value of services and increasing costs.

These dynamics suggest that *it would be beneficial to develop a clearer notion of regulating activities – rather than regulating types of organizations*. This would better match the modern structure of the Everything-as-a-Service economy, and would allow the broader business environment to continue to re-optimize for economic gains rather than to arbitrage regulations.

Competition policy

The emergence of the Everything-as-a-Service economy has significant implications for competition policy as well. This is because most anti-competitive concerns are reduced in the new business ecosystem, for three key reasons.

First, lower barriers to entry across industries mean that firms will often find it more difficult to successfully engage in anti-competitive behavior, regardless of their own size or market position.

Second, given the role that most large firms now play as part of a more cooperative business environment, they now have strong incentives to support rather than exploit others.

And third, the plug-and-play aspect of the Everything-as-a-Service economy lowers switching costs and makes it easier to displace “bad” actors, further limiting the scope for anti-competitive activity.

However, there is one way in which the Everything-as-a-Service economy may not mitigate anti-competitive concerns. Specifically, in small markets – for example, narrowly used drugs or patented technologies – it has become easier to set up capital-light, narrow-purpose entities that can exploit existing local markets. We call this “niche exploitation.”

The Everything-as-a-Service Economy

Consider, for example, a special-purpose company that purchases the rights to specialty drugs that are important to a small group of patients and then dramatically raises the prices of those drugs. The firm can manufacture and market these drugs by leveraging the benefits of the Everything-as-a-Service economy. Accordingly, this type of company can exploit the fact that it has only narrow market power, which is usually not a target of anti-trust officials. It can also exploit the fact that it can distribute its earnings unless or until this arrangement is challenged, leaving few assets behind to seize.

Thus, perhaps somewhat paradoxically, *in the Everything-as-a-Service economy, smaller businesses and smaller markets may have the largest inherent anti-trust risk as opposed to larger businesses and larger markets.*

In the next part of this publication, we focus on the competitive implications of the Everything-as-a-Service economy, with an emphasis on what business structures and strategies we think are likely to prevail, and what traps to avoid.

Part 2

New Rules for New Business Models

As we discussed in the previous section, in the new disruption-driven Everything-as-a-Service economy, where just about any business function can be outsourced, barriers to entry are generally lower across many markets. **Here we focus on the competitive implications and the strategies we think are likely to be successful in the new business environment.**

Focus on core areas of strength

To begin, we suggest a key tenet we believe companies should consider when operating in the modern business environment: **identify and focus investment in areas of competitive differentiation, and rely on the business ecosystem for all else.**

As we have said, in just about any market, a competitor – even a new entrant – can now scale quickly, with little to no capital and with little to no staff. And, since many firms now operate in a more modular structure and are generally more adaptable than in the past – which is an effect of the ongoing disentanglement we discussed in part 1 of this publication – the cost to transition to a new product or service is generally lower. This is also broadly true

The Everything-as-a-Service Economy

for consumers in light of new technologies and simplified user interfaces. Said another way, just as barriers to entry are lower across markets, switching costs are lower too.

Thus, for a firm to become a disruptor it simply needs to offer an *incrementally* better product or service than its competitors – and not even by that wide of a margin. Their customers can and will switch to new products and services to realize incremental benefits – and they'll do so for smaller gains than in the past. This is both an entrepreneur's dream and an incumbent's worst fear, since firms that fall behind risk rapid displacement.

Faced with the threat of disruption, it can be easy – but also potentially fatal – to default to what seems to be the only solution, namely speed and disrupting oneself.

Instead, we suggest that **companies should focus on and reinforce their key areas of strength, rather than attempt to shift gears and become something else entirely.** Simply put, being better at something that is *already* an area of strength is easier and more achievable than trying to be better by transforming into something entirely new. When viewed through this lens, operating successfully in the Everything-as-a-Service economy is more straightforward: build on strengths and whenever possible “rent” what's new.

The question for companies then becomes: how can a firm create a sustainable competitive advantage today, when just about any competitor – even a new entrant – can rapidly upend a market simply by leveraging third-party services, potentially even the ones that incumbents may also be using? We address this topic next.

New business models

Three of the four business models that are likely to prove successful in the new economy are essentially “classic” in nature, in that they are based on the well-understood economic models of

sustainable competitive advantage: economies of scale, scope and fit (or monopolistic competition).

First are what we call “*platform companies*,” which build their competitive advantages by creating *economies of scale* through higher asset utilization and optimization. This classic business model benefits from consolidation as the operational stack is thinned.

Second are what we call “*servicer companies*,” which build their competitive advantages by providing other firms with access to specific areas of expertise, usually referred to as *economies of scope*. Here the benefits come from consolidation as well, but only in the particular market segment that the servicer is addressing.

While the first two models have historical progenitors, the modern versions differ enough that the *definition of competition for these types of businesses is different today than it once was*. Newer versions of these entities can avoid vertically integrating their production stacks, unless doing so offers significant advantages. They can also achieve the necessary scale of operations by offering services to businesses that would have historically been considered competitors. Doing both well means that they must thin their operations – and this means that sometimes they are buying services from competitors while at other times they are selling services to competitors.

These dynamics affect well-known versions of competitive analyses, which often fail to address *the blurred lines that now increasingly exist between competitors and competitors that are also customers*. As a result, competition itself now relates directly and only to the firm’s areas of competitive advantage, which may not be evident based on the type of product or service the firm sells.

The third are “*organizer companies*,” which are focused on matching products with customers. In classic economics, this model is referred to as *monopolistic competition*, though what is

The Everything-as-a-Service Economy

different today is that there are many more distinct market niches that firms are able to target.


Organizers build their competitive advantages through *economies of fit*², which necessitates knowing their end-market and using this knowledge to create products or services that appeal specifically to their target customers. Additionally, customers must value these items in excess of the cost of production; organizer firms build their advantage by providing a better fit rather than by operating at a lower cost. In the context of modern disentanglement, this is a naturally fragmenting layer given the emphasis on specialization.

The fourth and final business model is new and is characteristic of the Everything-as-a-Service economy; it is underpinned by modern technologies that support information and matching economics.

We refer to entities in this space as “*learning companies*,” which focus on data collection and interpretation and *economies of learning*. These kinds of firms are often involved in online search, artificial intelligence (AI), big data analysis and some forms of software and social networking, as examples. We discuss when these firms can and cannot use data to build a sustainable competitive advantage through the notion of a “learning curve,” which provides a conceptual framework for assessing the scale-based economics of learning. Whether this is a naturally consolidating or fragmenting layer is case-specific. The four business models we have outlined are depicted in Exhibit 10.

² Historically, economies of fit were referred to as monopolistic competition (see Chamberlin, *The Theory of Monopolistic Competition*).

Exhibit 10: Business models that should prevail in the Everything-as-a-Service economy



Here we depict the four types of business models that are likely to prevail in today's economy. Each model is distinct in terms of its source of competitive advantage, meaning whether through economies of scale (platform companies), economies of scope (servicer companies), economies of fit (organizer companies) or economies of learning (learning companies). While three of these sources of competitive advantage have well-understood economic roots (economies of scale, scope and fit), the notion of economies of learning – which leans on data – does not. As such, we provide a distinct new framework for companies and investors to use to determine when data can serve as a source of competitive advantage.

Source: Goldman Sachs Global Investment Research

Of the four models we have described, thus far the Everything-as-a-Service economy has particularly enabled growth in the organizer category. Because just about any corporate operating function can be obtained from a third-party as a service, companies can focus on a single market. For organizers, the key to success is connecting with and understanding the appropriate target community. The more uniform the community is in its needs, and the more differentiated those needs are, the more protected the organizer is from displacement.

With these considerations in mind, we reiterate that each of the four types of business models is distinct in terms of its source of competitive advantage, meaning whether through economies of

scale, scope, fit or learning – and each is unique in terms of how it creates value.

What's more, for a firm to successfully leverage each model, its operations need to be organized and run in ways that optimize that competitive advantage. But, as we noted earlier, many companies do not neatly fit into any one specific category. Instead, they may be hybrids as various parts of their businesses may fit into one model while other parts may fit into another. There's nothing inherently wrong or unusual about this structure – what matters is how the whole fits together.

For the remainder of this discussion, we examine each of these four principal business models in further detail. We focus on key drivers of success as well as potential limits.

Platform companies: exploiting economies of scale

Platform companies' competitive advantage is built on effective asset management – whether the asset is physical or financial – and their principal aim is to achieve economies of scale through higher capacity utilization.

There are two primary formats for platform companies: *hosting* companies and *holding* companies. Hosting companies manage capital assets directly and focus on having a diversified base of customers. This allows hosting companies to achieve high levels of asset utilization by spreading the use of assets across a wide base of customers. In comparison, holding companies focus on owning a diversified portfolio of assets and primarily manage the portfolio rather than the underlying assets themselves. This helps them to generate a superior risk-return ratio than the average portfolio may generate. Regardless of the format, the source of competitive advantage for hosting and holding companies is the same: greater economies of scale.

For both types of platforms – hosting and holding companies – the

underlying math is essentially identical: larger diversified portfolios tend to be more predictable, and more predictable outcomes allow firms to plan and operate with tighter parameters. Taken together, these two factors increase their capital efficiency.

For holding companies, the construction of efficient asset portfolios is central to their strategy. For financial holding companies in particular, this concept is well-understood and based on standard portfolio theory. Nevertheless, even for these entities, the notion of efficient portfolio construction has been altered by the Everything-as-a-Service economy. This is because each entity within a financial holding company can now be thinned and optimized to focus only on those activities that benefit from the financial holding company structure, while offloading the activities that do not. In this way, as portfolio companies restructure to refine their focus, the holding company itself becomes both more efficient and better able to produce predictable returns.

For hosting companies, however, there's a concept of "load balancing" that is more subtle and generally less commonly discussed than efficient asset portfolio construction. Load balancing – or attempting to maximize asset utilization across as broad a swath of assets and over as long a period of time as possible – can be essential to achieving efficient economic outcomes, both for a particular firm and within a particular industry.

When market-share swings are an important and inevitable aspect of a particular sector – media is an obvious example – then shared infrastructure that supports the industry can be vital. Thus sharing assets across more customers, that may have different needs and may otherwise be competitors, over different time periods, can actually be the optimal outcome. In the end, this is precisely how sufficient economies of scale can be achieved.

The Everything-as-a-Service Economy

One way to better understand load balancing is to consider that it helps to explain why, in the Everything-as-a-Service economy, it's not unusual for a firm's competitor to also be its customer. Netflix operating on top of Amazon's cloud-based infrastructure is a simple example of how this works in practice: Netflix helps to raise the utilization rates of Amazon's cloud, even while Amazon competes with Netflix in streaming video.

In some markets, market share rather than total demand is the key factor determining firm-level asset utilization rates. When this is the case, the best way to reduce risk and increase average utilization rates is to share infrastructure – meaning either by hosting competitors or by renting capacity from competitors.

Another subtlety of the need to construct efficient user portfolios for hosting companies is that customer acquisition strategies can be priced based on usage complementarity. What this often means is that the hosting company can offer lower prices to customers that are flexible with their usage or that have usage patterns that naturally complement the hosting company's own usage (for example, the two companies don't share the same periods of peak demand). At the same time, customers that leverage the platform at peak times should be charged more to compensate for the risk that their demand exceeds the available capacity, potentially necessitating additional investment on the part of the platform provider.

For example, if two cloud-services companies have divergent baseloads – as Amazon and Google likely do given their different core businesses (e-commerce versus online search) – each firm is likely to evaluate potential customers differently and to charge them accordingly.

Consider live television streaming services in this context. The magnitude of peaks can be difficult to predict. Mass media events are often limited to a single service like broadcasts of the Super

Bowl or the Olympics, for example. When such events occur on one service, other service providers are often deterred from trying to orchestrate concurrent mass events that compete for the same audience.

As such, a hosting company that supports multiple live television streaming services gains efficiencies from consolidating varied peak activity across streaming services onto its own platform; a larger number of peaks, many of which are scheduled in advance and which are spread out over time, all improve the hosting company's asset efficiency. Following this same logic, the hosting company may be reluctant to support other entities with similar usage patterns as its live television customers – social networks, for example – since doing so would intensify peak usage, rather than diversify it.

Case study 8: Amazon's retail services business

Amazon's retail services exemplify the kind of platform hosting business we have described. By extending these services to third-party retailers – who are also its competitors – Amazon has been able to meaningfully improve its asset efficiency, well beyond what it otherwise could have accomplished.

To clarify, Amazon's retail services include its e-commerce website and the underlying IT infrastructure, as well as its expansive warehouse and logistics system. While these assets underpin the firm's own retail operations, they also support a large and growing network of independent sellers; in fact, of the billions of items that were sold on Amazon in 2017, more than half were from third-parties.

Amazon's retail services business is inherently capital intensive with natural scale economies. Starting with the company's mid-1990s launch as an online bookseller, Amazon began making significant IT infrastructure investments to improve the customer

The Everything-as-a-Service Economy

experience associated with e-commerce, given slow network speeds and limited website functionality at the time. The firm also began building a dynamic warehouse and logistics system that could efficiently support its rapidly growing e-commerce business.

To that end, Amazon has said that within its first two years in business as a bookseller, if it had a physical store instead of a virtual one, it would have occupied the equivalent of six football fields. In a move that helped improve its asset efficiencies, the company expanded into retail categories outside of books, to include CDs, DVDs, videos, home goods, among other items over time. While diversifying its own retail inventory would have improved Amazon's asset efficiencies, the extent of such activity would have been limited by the capital investments and the carrying risk involved.

By shifting to a platform hosting model, and encouraging third-party vendors to sell through its e-commerce site and leverage its logistics services, Amazon has been able to further optimize its asset efficiencies. Said another way, as Amazon's retail business has supported a growing number of individual retailers (a natural layer of fragmentation) it has benefited from higher utilization rates of its e-commerce and logistics assets (which are natural layers of consolidation) (Terry, Jan 2013). See Exhibit 11.

Exhibit 11: Amazon's e-commerce business, then and now

Growth as a platform hosting company

THEN May 1997 (IPO)	
2.5mn+ book titles for sale	\$16mn+ total revenues (CY 1996)
Carried most of the ~1.5mn English-language books in print	~340,000 customer accounts
1mn+ out-of-print titles	Available in 100+ countries
Carried a limited number of other items: CDs, videotapes and audiotapes	80,000 average daily visits
	40% of orders from repeat customers
"The Earth's Biggest Bookstore"	

NOW 2017	
100mn+ unique products for sale	\$177bn+ total revenues (CY 2017)
5bn items shipped with Prime globally	100mn+ paid Prime members globally
>50% of units sold came from third-parties	~60% of total revenues from the online store
Millions of small businesses sell through Amazon	Millions of Amazon's hardware devices sold
\$150bn invested between 2011 and 2017 in Amazon's global fulfillment networks, transportation capabilities and technology infrastructure	

Source: Company data, Goldman Sachs Global Investment Research

Despite the fact that the firm now handles billions of unit sales, Amazon has fewer than 800 fulfillment centers globally. The placement of each one and the inventory management within are

done strategically to ensure efficient delivery. Amazon is able to leverage the data it collects on the retail sales that occur on its platform to drive greater asset efficiencies across its warehouses and its logistics system more broadly. While the firm has been able to use data to enhance its operating strategy and structure, the information Amazon has amassed about customers' past purchases has not given Amazon an edge in retailing relative to others, as we will later discuss.

By operating as a hosting platform company, Amazon has been able to take broad-assortment retailing to the extremes and to take share in an established marketplace, against long-standing and well-established market participants like Walmart (Collett, Aug 2015). At the same time, consider that Amazon's e-commerce business is largely successful as it relates to the sale of goods that can fit inside cardboard boxes and can be delivered by trucks in the span of a few days. As a point of comparison, consider that in China's largest cities, e-commerce businesses – like Alibaba – are increasingly popular to support the sale of fast-moving consumer goods and fresh food, which can be transported using bicycles or motorcycles in the span of just 30 to 60 minutes from the time of ordering (Keung, Jul 2018).

Case study 9: Berkshire Hathaway

Unlike Amazon, which represents a typical hosting company, Berkshire Hathaway represents a typical financial holding company. It is based on achieving financial asset efficiency – and maximizing returns for a given level of risk – by maintaining a diversified portfolio of investments.

As is widely known, Berkshire acts as a holding company that invests its capital in businesses across sectors and markets, with a focus on portfolio companies it considers to be strong financial performers and market leaders. Berkshire's efforts are bolstered by its focus on portfolio companies that produce regular dividends, providing a reasonable source of cash flow that allows

it to make further investments in softer markets – which is of course the best time to buy. These dynamics help the company to, over time, generally generate a superior risk-return ratio than the average portfolio.

For its portfolio companies, the Berkshire holding company serves as a more stable and less expensive source of funding than the alternatives, such as bank or public market financing. For some of these portfolio companies, a lower cost of capital relative to competitors can serve as a valuable source of differentiation and competitive advantage – potentially reducing the need for these firms to engage in riskier endeavors to achieve the same outcomes. In this way, Berkshire is not the only beneficiary of its strategy.

What's more, as we mentioned earlier, each portfolio company can now be thinned and optimized in the Everything-as-a-Service economy, such that they can focus only on those activities that benefit the Berkshire financial holding company, offloading those activities that don't. As Berkshire's portfolio companies restructure to refine their focus, Berkshire itself becomes both more efficient and better able to produce more predictable returns, often including a steady stream of cash flows.

Case study 10: biopharma

The pharmaceutical (pharma) industry has developed a similar holding structure for slightly different reasons: large pharma companies can be thought of as holding companies of a portfolio of drugs.

There are a number of operating functions, such as marketing and distribution, which have some scale economies that help engender this structure. The risk management gains, however, have a different logic than in most financial holding companies. Specifically, drugs tend to compete in disease categories, and market share within these disease categories can be an important

source of risk (particularly as new drugs come to market); thus assembling portfolios of related drugs (meaning within the same disease categories), is more efficient in practice than establishing a portfolio of drugs that is diversified across diseases.

This is because, put simply, disease incidence (or demand) is far more stable than the prescribing habits of doctors. Therefore, as research and experience shift demand from one drug to another, large pharma firms can maintain more stable revenues by having a portfolio of related drugs, and by capitalizing on operating scale efficiencies in marketing to doctors that share the same disease specialty. To assemble these drug portfolios, large pharma companies tend to need to be acquisitive, collecting efficient assets much as a typical holding company would, but with different risk patterns (since concentration in a particular disease rather than diversification – perhaps counterintuitively – reduces risk) (Richter, Jan 2018) (Richter, Apr 2018).

Consider Johnson & Johnson, for example. Roughly half of its business is made up of strategic acquisitions and in-licensing deals, while the remainder is dependent on internal sources of research and development. The company has successfully engaged in early stage in-licensing and tuck-in merger and acquisition activity – including Cougar Biotechnology, Pharmacyclics and Genmab, to name a few – which were done when these firms' compounds were in the early phases of development. At the same time, Johnson & Johnson's larger deals, including its purchase of Actelion, provided the firm with access to a fully de-risked in-line portfolio of rare respiratory-disease drugs (Rubin, Nov 2016). See Exhibit 12.

Exhibit 12: Johnson & Johnson is an example of a platform holding company
Examples of strategic in-licensing deals and acquisitions



Source: Company data, Goldman Sachs Global Investment Research

A second source of efficiency, which is perhaps more interesting, has also led to this structure, namely: diseconomies of scale in research. For a variety of reasons, small and narrowly focused firms (often ones that focus on a single drug or disease area) are more efficient at drug development than are large firms. When viewed this way, late stage drug acquisition strategies can be the most efficient, rather than “in-house” development for large biopharma companies, even if this concept feels counterintuitive at first glance.

The oil industry has seen a similar pattern with respect to the development of shale, where smaller companies appear to be better at finding and developing assets, but larger companies appear much better at exploiting capital and logistic efficiencies as the shale assets become more established (Della Vigna, Nov 2017) (Della Vigna, Mar 2018), (Della Vigna, Apr 2018). Software also has a similar pattern: small firms’ innovations are often collected by larger firms that then leverage these innovations across a broad customer base; thus the established patterns of merger and acquisition strategies in these sectors.

Servicer companies: exploiting economies of scope

Servicer companies are the capital-light version of platform companies and they solve a new(er) problem that has arisen from the Everything-as-a-Service economy as an increasing number of companies are able to conduct commerce across multiple regulatory, tax, legal, communication and other technical environments with highly specific business requirements. This creates a need for entities that offer deep but narrow process expertise to manage the critical business issues that would otherwise prevent many companies from being able to scale geographically.

Servicer companies typically provide narrow, well-defined functions that are based on dynamic standards. Their offerings are comprehensive and are provided to customers in a way that allows the customer to essentially ignore the technical complexities, but in the end still achieve its specific aims. As part of this, servicer companies facilitate connectivity with their customers often competing not only on price, but also on ease of use.

In effect, servicer companies are enablers of the Everything-as-a-Service economy: the outsourcer can hire technical services from a third-party. These technical services can be complex, inherently dynamic and scale optimally, but the outsourcer can treat the service as though it were simple and static. This allows outsourcers to ignore complex non-core processes, focusing instead on their own core competencies.

For example, multi-jurisdictional payrolls and taxes, financial connections across multiple firms or entities, recruiting for multiple specialties, global marketing and logistics, specialized production, as well as a host of other functions, can all be managed to local standards. And the outsourcer doesn't have to concern itself with these functions, relying instead on servicers.

Servicer companies invest in intellectual capital to ensure ease of use while connecting outsourcers with providers; platform companies, in contrast, depending on the type, invest in physical capital or financial assets and focus on asset utilization rates or portfolio diversification. As we have said before, in the new economy, competitors can also be partners. For example, Stripe (a servicer company) uses Amazon's cloud infrastructure (a platform business) to power its payments technology even though Amazon competes in the space with Pay. At the same time, Amazon uses Stripe to handle some of its own payment transactions.

In terms of classic economics, *servicer companies are focused on economies of scope*. Servicer companies sell a form of expertise and invest in broadening their portfolio to meet the needs of new types of customers, potentially even in new geographies – rather than scaling specific business functions or maximizing physical or financial assets. In doing so, the servicer firm can extend its expertise and even add to its knowledge base, allowing it to deepen its specialization. To put a finer point on the notion, servicers invest in connectivity, knowledge and customer ease of use, rather than in physical or financial assets, where platform companies focus.

For example, because Amazon Pay is primarily structured to keep users on Amazon's platform, capital investments beyond these purposes may not be economically necessary or efficient. In contrast, Stripe is designed to work across platforms. Thus, even when the platform and servicer are providing the same offering, the service that the platform company provides is often tied to usage of an actual asset, with an emphasis on economies of scale, while what the servicer company provides enables the user to be asset indifferent, with an emphasis on economies of scope.

Case study 11: ADP in payrolls

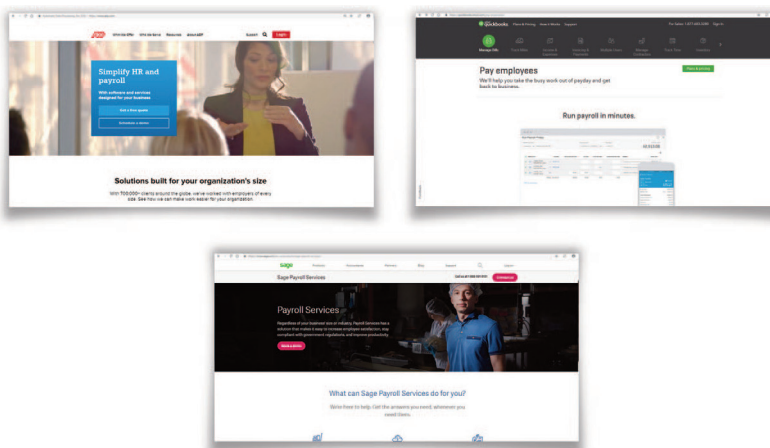
ADP is a classic example of a servicer in the various ways we have just discussed. While the company offers a broad suite of human capital management tools, its payroll services in particular are used by hundreds of thousands of businesses around the world, of varying sizes, in a range of industries and with different types of employees.

The company's payroll offering is intended to be an end-to-end solution that allows customers to offload the payroll function so that the outsourcer need not expend resources on complex payroll related matters. ADP's payroll solution can calculate employees' pay, assess tax withholdings, create paychecks and manage direct deposits, and also produce payroll reports and prepare firms' payroll tax returns as examples.

Managing other firms' payrolls necessitates a wide range of expertise – and this is precisely where servicers can come in handy in the Everything-as-a-Service economy (and why they tend to focus on economies of scope). Consider the complexities associated with a US-based employer paying a mix of traditional and freelance employees who are based overseas. Doing this effectively necessitates having a comprehensive understanding of federal, state and local payroll and tax requirements, both in the US and abroad. What's more, from a functional perspective, payroll providers must ensure that their technology can interact with a wide range of systems, which includes being able to connect to their clients' human resources systems to obtain the latest employee records, or to banks' infrastructure to facilitate direct deposits.

In summary, ADP and businesses like it, including Intuit and Sage as examples, benefit from the economies of scope associated with having deep payroll expertise that can be leveraged across an expansive customer base. See Exhibit 13.

Exhibit 13: Payroll providers like ADP are examples of modern servicer companies
Website screenshots from ADP, Intuit and Sage reflect simple user interfaces



Source: ADP, Sage, Goldman Sachs Global Investment Research. The Intuit screenshot was reprinted with permission © Intuit Inc. All rights reserved.

At the same time, ADP's customers benefit from being able to offload this core business function to an expert third-party so they can bypass the associated in-house investments, time and effort, and instead focus on their own core competencies.

Case study 12: payment processors

Likewise, digital payment processors are servicers that benefit from economies of scope. These firms provide payment services across many different types of users – consumers, merchants and financial institutions, for example – and across many different types of platforms.

Consider PayPal, for example. The firm, which is a leader in digital payments, connects millions of merchants and customers around the world with its technologies and facilitated more than seven billion payment transactions in 2017 alone in a range of currencies.

PayPal's users are able to leverage their accounts to receive and

The Everything-as-a-Service Economy


transmit payments for goods and services and to transfer and withdraw funds, among other capabilities. The functional, technical and regulatory complexities associated with facilitating these interactions are significant and necessitate deep expertise.

From a technical standpoint, PayPal has made significant investments in the technology infrastructure that underpins its payments solutions and has made it easy for developers to build its payment solutions into their mobile or web applications through standard APIs.

The complex inner workings of PayPal's payments system are not evident to the merchants that incorporate the service into their production stack, nor are they evident to end-users – because it leverages the modern user interface standards that have now become the norm. This means that PayPal's customers need not understand or concern themselves with the underlying workings of the product, again freeing them to focus on other core areas. As Exhibit 14 shows, there are a number of businesses that provide similar services, including Stripe, Square, Google Wallet, among others (Schneider, Aug 2017) (Ramsden, May 2018).

Exhibit 14: Payment processors - servicer company examples

Providers include PayPal, Stripe, Square and Google Pay



Source: PayPal, Stripe, Square, Google, Goldman Sachs Global Investment Research

Case study 13: Qualcomm

Qualcomm is another example of a servicer that enables outsourcers to focus on their own core competencies instead of non-core functions. The firm sells microchips that it designs and licenses its systems software to manufacturers, which then build these technologies into the devices they produce. Qualcomm's technology enables these devices – smartphones, tablets and laptops, for example – to connect with cellular networks.

Enabling reliable cellular-network connectivity across devices, networks and regions is a complex endeavor. It requires mastery over an expansive technology landscape, which includes evolving hardware and networking requirements, down to the protocol level. Qualcomm, which is an established expert in this field, is able to build this functionality into a limited number of chips that fit a wide range of devices.

Device designers and manufacturers can rely on Qualcomm's expertise to enable mobile internet connectivity across all of the devices they conceive of and produce. At the same time, Qualcomm benefits from extending its deep expertise across a wide range of customers and devices, improving its operating leverage.

Organizer companies: exploiting economies of fit

The next type of business in the Everything-as-a-Service ecosystem is the organizer. Firms that are organizers are paradoxical in that they represent both the most historically-consistent and historically-divergent business model today.

At the most basic level, organizer firms best match their products or services with the people who want to buy them. Doing this well has always been a defining characteristic of a successful firm – and

The Everything-as-a-Service Economy

it still is. In the Everything-as-a-Service economy, however, organizers can be narrowly focused, potentially free from the burden of having to address physical production or distribution in-house, as examples.

By tactically leveraging services provided by third-parties, organizers are able to have two primary areas of focus: knowing who their customers are and knowing which products or services best meet these customers' needs; all other activities are optional.

As we discussed in the first part of this publication the advantages that come from having a fully-disentangled business model are numerous – perhaps more so for organizers than for other types of business models. What's more, organizers can be both narrow in scope and capital-light, while also being global and scalable.

Because organizers are defined by their ability to match customers with the products or services that best meet their needs, in order to establish a competitive advantage, organizers must understand and maintain the match between their offerings and their target audience. This means that *their vulnerability to competitive displacement is largely determined by the strength or fragility of their matching abilities*. As we will discuss next, the ability of an organizer company to identify or create a “community” where it can specialize in matching members of that community with the goods and services they need serves as a form of “brand protection” that can replace old barriers to entry.

Consider Apple as an example of an organizer. The company sold nearly 280 million iPhones, iPads and Mac computers in 2017 to customers around the world. However, the firm relies entirely on third-parties for device manufacturing and assembly. As is often the case for organizers, by disentangling device production from design and distribution, Apple not only operates with far less capital than if it were instead vertically integrated, but it is also able to focus its resources on the activities where it can offer differentiated value to its customers.

Accordingly, organizers have two methods for creating potentially sustainable competitive advantages: product superiority and community loyalty and fit. As in the past, establishing product superiority often requires the use of proprietary inputs or processes. But, in the Everything-as-a-Service economy, the organizer can limit the scope of its operations to focus on its core product advantage, while also locating and serving the target customer with ease. In this way, the organizer can quickly reach optimal scale across a global marketplace, while realizing higher margins and a higher return on equity than would have otherwise been possible given its lower capital requirements.

Next we discuss how organizers can build a potentially sustainable competitive advantage through community relationships.

Organizers and the role of communities

Long before the Everything-as-a-Service economy, firms' target markets were often regionally constrained given production and distribution limitations as well as advertising reach. Today, these limitations have largely dissipated, not only because of firms' disentangled organizational structures, but also because of the proliferation of internet-based communities – and simply because of the walls that technology has broken down.

These communities are effectively groups of individuals who freely share their opinions about products and services. They are often the best source of information for firms to learn about what existing or potential customers truly value for the obvious reason that such communities are typically comprised of their target customers.

Although communities have always existed, they have increased in number and their scope has been refined in the last two decades – thanks in large part to the rising popularity of social networks and other online channels that support these groups (Twitter,

Facebook, YouTube, Amazon and Yelp, to name a few). Likewise, membership and participation in communities has dramatically increased in light of how easy it is to join one and since technology has largely removed previous geographical and accessibility limits.

In fact, any group with a common need – and a name for that need – can quickly form a community today. The community then becomes a defined market segment, and companies can create custom products to match the community, without being limited by the natural scale of production or of distribution. *This ability – to find, create and serve communities – is a defining characteristic of the Everything-as-a-Service economy – and for organizer companies more specifically.*

Given the importance of communities, there are several essential features that are worth noting. For instance, the community is the arbiter of success of the products and services that are produced to serve it. This has important competitive implications, since it's now a lot easier for consumers to coalesce, verify quality and assess relative value – and to do so quite publicly, quickly and with little or no input from the company (Terry, Mar 2015).

What's more, communities decide their own scope by common consent – growing, shrinking and splitting in ways that can affect how companies can engage with them. Sometimes the evolution of a community can lead to a niche market. When a niche market forms and is well-defined (meaning its needs are specific and understood) companies can create products and services to address those needs. Profitability isn't dependent on size: some profitable niche markets become large; others stay small.

In practice, even while organizers rely on communities for valuable customer insights, these firms cannot own, control or legally protect these groups. The fit and loyalty that the organizer engenders from any community must be earned and then

re-earned on a daily basis – which is precisely why product and service quality are key.

For organizers, capturing and keeping consumers' mindshare and wallet-share is therefore about community-building and product or service focus. By building an online presence, leveraging search-engine optimization as well as social media adeptly, companies can create globally-recognized brands more quickly, more easily and at a lower cost than in the past. But to do so effectively, they must develop and maintain bonds with their user communities. Failing in this regard can quickly erode brand value.

From the organizer's perspective, there are strong incentives to protect its position in its core market, and few incentives to invade others' territory, as markets tend to segment into well-defined communities.

Expanding into new, different or potentially incompatible communities is a risky endeavor since doing so can have the deleterious effect of damaging the organizer's ties to its core community. The inherent contradiction associated with mass-market luxury offers a reasonable example; it also helps to explain why brand segmentation is no longer as effective a strategy as it once was.

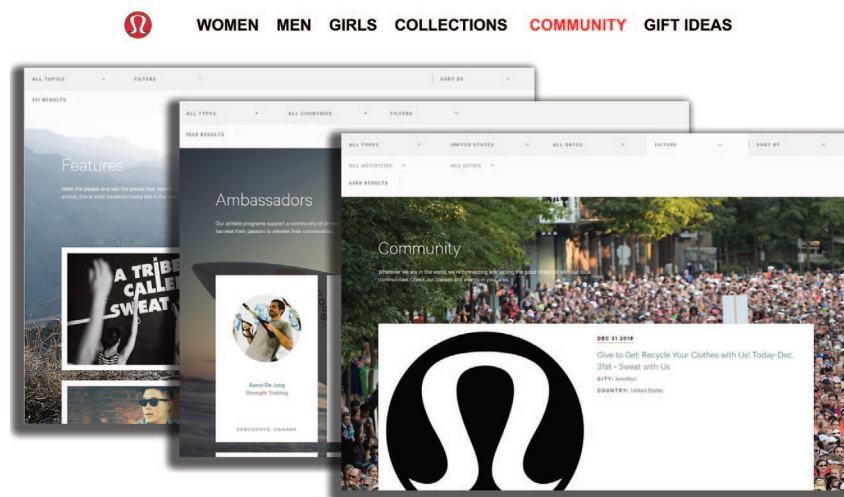
To that end, exclusivity often matters. This creates strong self-reinforcing patterns where firms diverge from one another in terms of their core areas of focus, rather than overlapping with one another. In economic terms, this is what is referred to as monopolistic competition. For organizers, the key is to find or to create self-identified communities with differentiated needs that are also profitable markets. The more distinct a community is from the other groups, and the more uniform the participants within that community are in terms of their wants and needs, the easier it is for an organizer to both cater to this group and to defend itself against displacement.

Case study 14: lululemon athletica

Consider lululemon athletica as an example of an organizer. The firm designs, distributes and sells premium-priced fitness clothing and gear to individuals pursuing an “active, mindful lifestyle.” In many ways, lululemon’s business model reflects the modern economy: the firm relies on third-parties to supply the fabrics for its apparel and to manufacture its products, choosing to direct its resources to overseeing these operations and to maintaining its own retail operations, both through physical stores and a growing e-commerce presence. By operating in this manner, the firm is able to focus on design, distribution, inventories and pricing – while also being able to connect directly with its community of users and iterate on its goods and services based on community feedback.

What’s more, lululemon leverages its salespeople and in-store community boards, brand ambassadors and other grassroots initiatives to bolster its “identity” and its appeal. Lululemon’s community-focused feedback loop is essential to its ability to provide its customers with the products that best fit their needs, and for the firm to defend against displacement. To that end, digital marketing and social media are critical to the firm’s community interactions, for example, given its more than two million Instagram followers. See Exhibit 15.

Exhibit 15: As an organizer, lululemon's community focus is key
Screenshots from the "community" section of lululemon's website



Source: lululemon athletica, Goldman Sachs Global Investment Research


Execution missteps can be costly. For example, in early 2013 lululemon was forced to recall its signature product – premium-priced black yoga pants – due to a failure in quality control. Initial attempts to remediate the problem were not effective, eroding community trust. The issue affected the firm's profits, reduced its market value and ultimately prompted leadership changes. Despite this issue, over time, as the athletic apparel market has fragmented, lululemon has benefited by catering to narrow but profitable markets (Walvis, Jun 2018).

We should note that pure organizer firms are likely to remain rare. This is because organizer firms often need to maintain some control over their production processes in order to be able to create the products or services that best meet the needs of the community (or communities) that they serve as their primary end market(s). See Exhibit 16, which shows lululemon's approach in

The Everything-as-a-Service Economy

this regard. In the Everything-as-a-Service economy, however, such firms should keep this level of control to a minimum, and that minimum is likely to decline over time.

Exhibit 16: In many ways, lululemon's business model reflects the modern economy



In many ways, lululemon's business model reflects the modern economy: the firm relies on third-parties to supply the fabrics for its apparel and to manufacture its products, choosing to direct its resources to overseeing these operations and to maintaining its own retail operations, both through physical stores and a growing e-commerce presence. By operating in this manner, the firm is able to focus on design, distribution, inventories and pricing – while also being able to connect directly with its community of users and iterate on its goods and services based on community feedback.

Source: lululemon athletica, Goldman Sachs Global Investment Research

Learning companies: the benefits of information

The last type of businesses are called learning companies. These entities build a competitive advantage through the effective utilization of data. More specifically, these firms collect and analyze data and leverage what they learn to create competitive differentiation, through organizational or output optimization. *Learning companies often have hybrid business models, since directly*

monetizing data-based insights can be difficult to do in practice and – in some cases – may be almost an after-thought once a company has already been established.

We see four types of learning firms: data-smart companies, data-asset companies, data-feedback companies and data-network companies.

Data-smart companies use internally-generated data as the foundation for their data-based insights – or what can be thought of as learning – which they then use to optimize both their operations and their output.

Data-asset companies tend to purchase or build propriety datasets from secondary sources (for example, data collected from sensors, genetic labs or satellites). These companies then use these datasets to provide data-driven services to others. To that end, data-asset firms are effectively platform companies – but ones that are dependent on data-driven asset efficiencies.

Data-feedback companies collect the data that are generated by users who are *already* leveraging the company's products or services. These companies analyze these data and leverage the resulting insights to improve their output; said another way, these companies create a feedback loop between their users and the goods or services they sell to those users (think of Spotify's playlist suggestions, Google Maps or even Amazon's product recommendations).

Data-network companies are similar to data-feedback companies in that they collect data generated by users who are already leveraging their output – but they use these data for a different purpose: *to connect their users to one another*. Examples of data-network companies include Uber, Lyft, Airbnb and Facebook.

While the economic models that underpin each type of learning company are unique, they share the common characteristic that each one requires data accumulation to drive learning, which then serves as their source of competitive advantage.

Learning companies and “the learning curve”

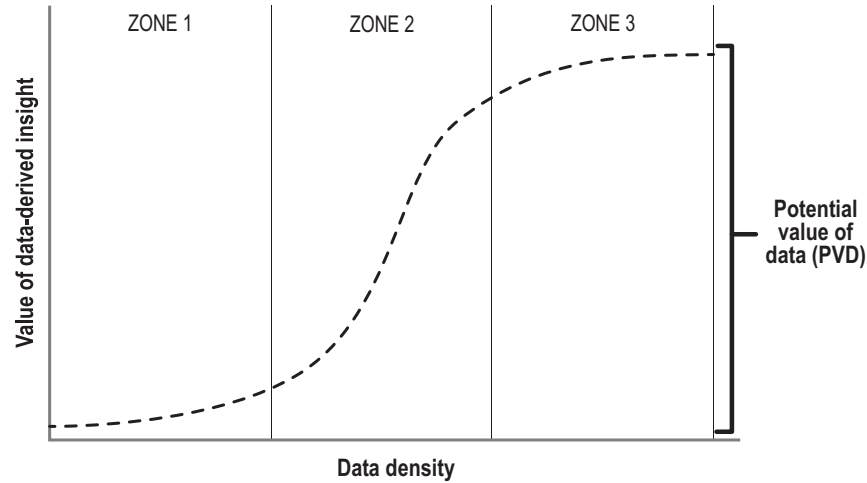
The learning curve in Exhibit 17 shows the potential value of data – or the total value of what can be learned from data – as a function of the amount of usable data a firm possesses³.

Specifically, each unit on the y-axis represents the incremental value derived from analyzing data related to a specific question or problem, while the x-axis represents the density (or volume) of usable data, which is dependent on the rate of data collection and the rate of data decay. The potential value of data (PVD) represents the *total* potential value that can be created through data analysis.

³ The learning curve can be derived a number of ways. In network theory (perhaps the most interesting derivation in the current context), the learning curve is the probability of being able to connect any two randomly chosen individuals in a large fixed population, where the x-axis is the number of individuals that have been linked randomly and pairwise to each other prior to the calculation. Another derivation of the learning curve comes from finance: it is the hedge ratio that provides optimal protection for an option against an increase in price (the market's summation of information). Perhaps the most general derivation of the learning curve is the central limit theorem from statistics, where the curve is the asymptotic distribution of the mean of any data-generating process, regardless of the underlying statistical distribution. The precise functional form describing cumulative probability in all of these derivatives is the cumulative normal probability curve. The key requirement for the learning curve to function as described is that the data generating function needs to be stable (although the precise notion of stability differs by application). In the context contained herein, data decay serves this purpose. You cannot learn from data if each new piece arises from a different random process; thus, data decay permeates our discussion.

Exhibit 17: The learning curve: the potential value of data (PVD) as a function of the amount of usable data a firm possesses

A conceptual framework for assessing the scale-based economics of learning



From an economic standpoint, the central point of the learning curve is that data-derived knowledge does not increase without bounds as the volume of data increases. Each unit on the y-axis represents the incremental value derived from analyzing data related to a specific question or problem, while the x-axis represents the density (or volume) of usable data, which is dependent on the rate of data collection and the rate of data decay. The potential value of data (PVD) represents the total potential value that can be created through data analysis.

Source: Goldman Sachs Global Investment Research

From an economic standpoint, the central point of the learning curve is that data-derived knowledge does not increase without bounds as the volume of data increases. This is for the simple reason that once there are sufficient data to answer the question or problem at hand, additional data only confirm what's already known – and the value of additional data and analysis is trivial. The potential value of a learning advantage is thus constrained by the nature of the question (or questions) at hand.

Thus, for each type of learning company, the uncertainty related to the PVD is a central question. The PVD must be large enough to justify the expense of building, buying or collecting the data. But, as is often the case, the actual value of data-based insights is largely

unknown until after the underlying database is built and the analysis has been done.

Another central question all learning companies must answer is related to data scarcity. On the one hand, if there aren't enough data available to analyze, deriving data-based insights simply isn't possible and businesses can get trapped in zone 1. On the other hand, if there are a lot of data and the data are widely accessible, everyone can use the data to support their businesses. In this case, data and learning are a cost of entry (rather than a competitive advantage) and all participants end up in zone 3, where data-based analysis does not provide meaningful competitive differentiation.

With this in mind, consider that the learning curve has a fairly specific shape that is common to all learning problems, and that it is comprised of three specific zones:


- *In zone 1, the learning curve is flat and the incremental value associated with data analysis is low. This means the gains associated with additional data analysis and data density are limited. The slow learning is due to the fact that a certain amount of data must be collected before it can be effectively modeled.*
- *In zone 2, the curve begins to slope upward and becomes steeper, typically very steep. At this point, the nature of the data model has become clearer and is better defined, so the incremental value of data-derived information is high. As a result, in this zone, accumulating more data – particularly relative to competitors – can result in a maintainable data advantage or MDA and can generate significant incremental value, as Exhibit 18 shows.*
- *In zone 3, the learning curve flattens since additional data accumulation and analysis no longer result in significant incremental value. In this zone, the learning process is nearly complete since most of what can be learned from data to address a specific*

question or problem has already been learned. Firms in the same market segment that reach zone 3 are in essentially the same competitive position.

While not technically precise, it can be helpful to think of zone 1 as the model specification search, zone 2 as the model estimation and zone 3 as the model verification.

Thus, companies can use different strategies to take advantage of the differing economics of the learning curve. As we noted earlier, we see four unique types of data-based learning companies: data-smart, data-asset, data-feedback and data-network firms. Each of these types of firms can be understood in the context of where they are located on the learning curve.

Exhibit 18: The learning curve - when data density is sufficient and there is a maintainable data advantage (MDA)



The learning curve has a fairly specific shape that is common to all learning problems. It is comprised of three specific zones. In zone 1, the learning curve is flat and the incremental value associated with data analysis is low. In zone 2, the curve begins to slope upward and becomes steeper, typically very steep; in this zone, accumulating more data – particularly relative to competitors – can generate significant incremental value. In zone 3, the learning curve flattens since additional data accumulation and analysis no longer result in significant incremental value; in this zone, there is little competitive differentiation between firms.

Source: Goldman Sachs Global Investment Research

Data-smart companies

It's a popular refrain to believe that all companies must become data-smart, but in practice it may not be possible or relevant to pursue this strategy. By using the learning curve to frame the issue, however, companies can begin to assess whether it's possible or worthwhile to pursue a data-smart model.

As we discussed above, strategies to move along the learning curve often have high fixed costs associated with capturing the necessary data and constructing the required analysis. At the same time, the gains from such investments are typically unknown until the strategy is fully implemented.

Furthermore, individual companies frequently have difficulty producing enough data on their own to be able to implement big-data types of analyses. Modern approaches to big data, AI and the like require vast quantities of data to produce meaningful insights that can move a firm from zone 1 to zone 2. Thus, in many cases, being data-smart simply proves impossible and a single firm on its own ends up stuck in zone 1 with little to show for its efforts.

However, if an individual company is able to generate enough data to successfully reach zone 2 or even zone 3, it is likely that the data will be related to highly repetitive tasks, as in the case of logistics, simple customer support or other basic operations. The risk-to-reward associated with making significant investments in collecting and analyzing such data – based on the notion that doing so will reveal hidden or unknown insights – may be poor; said another way, the PVD may not be sufficiently high relative to the investment involved. Rather than using data to optimize a product or service, instead what may be a better strategy – with a more favorable investment outcome – could be to focus on operational optimization, using a high initial volume of data to run the first analysis, which can then be complemented by high ongoing usage that allows even small improvements in efficiency to accumulate with meaningful results.

Data-smart strategies are therefore often likely to culminate in zone 3, which means that data-smart strategies will generally be defensive in nature (they may be a cost of entry, for example). This is because a failure to realize the basic efficiencies associated with the data-smart strategy are likely to put those entities that fail to do so at a significant disadvantage relative to other firms that have realized those efficiencies.

Another implication associated with this type of business structure is that *it may actually be better to be a second-mover rather than a first-mover from an investment perspective.* Knowing another company has succeeded at uncovering meaningful efficiencies from a particular data-smart strategy *significantly* improves the related risk-to-reward ratio. Put another way, it may be better to mimic the strategy that's already proven successful, rather than to create a novel data-smart strategy.

Data-asset companies

Data-asset companies must build databases that allow them to offer a learning-based service. This is in contrast to data-smart firms, which *already* possess the data necessary to pursue a learning strategy.

For data-asset firms, constructing a database typically requires a significant upfront investment associated with acquiring the necessary data, as does the related analysis. What's more, at the point when these investments are made, the firm typically does not know how much data will be necessary to allow it to progress into zone 2 or zone 3, nor does it know the PVD of the data.

Thus, the risk-reward ratio of data-asset strategies is in many ways analogous to deep-water drilling for oil or to new drug development: there are high up-front costs and significant uncertainty associated with discovery, but there is also a long tail of payments if the endeavor is successful. Another similarity is that data-asset strategies also require significant capital and diversification efforts to create a reasonable risk-reward tradeoff.

Accordingly, it is not surprising that well-established firms like IBM with its Watson Health Imaging business (and to a lesser extent, Google and Amazon) have led the way in the data-asset space, though there are start-up firms that have made some inroads (as with Flatiron Health, for example, which was acquired by Roche Holdings).

It is also worth noting that there are a number of important *differences* between data-asset strategies and oil platforms or the development of new drugs. Perhaps the most important difference is that data-asset firms, unlike oil platforms or pharmaceutical companies that develop new drugs, must assess “copy risk,” since *potential* competitors (aka new entrants) face very different incentives and hurdles than the innovators themselves.

As we touched on earlier, second movers do not face the same level of uncertainty that first movers do, thus their investments are subject to a more favorable risk-reward tradeoff. This is because second movers already know that valuable data-based insights do exist. *They also have a general sense both for the volume of data necessary to extract these insights and for the magnitude of the related PVD. At the same time, the second mover faces the risk of lower potential profitability;* this is because when the second mover enters the market, the first mover is incentivized to cut prices well below the average cost for the simple reason that the marginal cost to deliver data-based services is lower than the fixed cost to develop the services in the first place.

This type of copy risk can be difficult to determine, particularly before a company knows how much value a particular data-asset will generate to address a specific problem or question; this only reinforces the need for data-asset firms to diversify and to have sufficient capital to experiment again.

Broadly speaking, however, as companies decide which investments they should make, there are two observations worth

considering. First, if it is likely that the *full investment* (in both the data and the related analysis) will need to be replicated to produce the results, the investment is likely safer from a risk-reward perspective. Second, and on the flip side, when it is likely that a second mover will be able to bypass the full investment and still arrive at the same results, the original strategy is more likely to be copied and the risk associated with the first mover's investment is high.

The nature of copy risk can be considered via examples. Consider a "safe" example first, meaning a case involving low copy risk. As in the case of IBM's Watson Health Imaging business, interpreting MRI data requires a large start-up database of interpreted images and significant ongoing technology investments, both to receive and to interpret new MRI data. Thus an ongoing build of cross-checked interpretations would make replicating this data-asset strategy difficult.

As another example, consider a data-asset enabled maintenance service for elevators, which is based on data collected from sensor arrays or histories of elevator maintenance – this example could go either way. On the one hand, if producing the maintenance service requires a complex assessment of the sensor input data, copying the strategy could be difficult. On the other hand, if the maintenance service could be approximated through simpler forms of analyses, for example by counting hours of service rather than calendar time associated with the service, it could be replicated at a lower cost – and the associated copy risk would be high.

It is worth highlighting one more difference between data-asset companies and somewhat analogous deep-water platform companies: namely, *economies of scope can easily play a significant role in driving data-asset efficiencies*. This is because the lessons learned and technologies developed in one data-asset project may result in new but related projects. Sensor-based data collection

and interpretation, or image-based data processing and interpretation, as examples, could easily represent natural projects with scope efficiencies. In either instance, the related skills could be applied to many different databases and therefore could allow a skilled data-asset company to become even better at both assessing the risks and lowering the cost associated with new ventures related to their particular area of expertise.

Thus, by organizing themselves as diversified data-learning firms, data-asset companies can combine the risk efficiencies inherent to platform holding companies with the scope efficiencies of servicers. However, as we discussed earlier, this hybrid entity would still be constrained by the underlying mathematics of the learning curve.

Data-feedback companies

The most complex but most talked about learning companies are those that rely on the collection of user data to refine the user experience – hence the name: data-feedback companies. For data-feedback companies there are two distinct but related challenges. The first is to find an advantage, and the second is to maintain it.

Historical efforts suggest that finding a true advantage based on customer data isn't easy. "Discovered behavioral patterns" related to individuals generally aren't complex or surprising. Amazon offers an illustrative example. The firm's internal use of consumer data for logistics and inventory management (data-smart strategies) has been helpful; the firm also has one of the largest customer databases ever amassed. Yet, the firm's product placement and sales strategies are often quite simple, to the point where third-party retailers have been able to mimic Amazon's strategies and outpace Amazon in terms of unit sales on Amazon's own retail platform.

As a simple example of the limits of the value of user data, a firm doesn't need to have Amazon's extensive customer database to

realize that a consumer who is searching for ovens may want to purchase one. While an advertiser can use this information to display ads of ovens (showing ones that are better or cheaper, but similar to what the consumer has already viewed), for a merchant to serve this customer well, more often than not, it will simply need to stock the most popular oven models, which does not necessitate extensive customer-specific data or analysis.

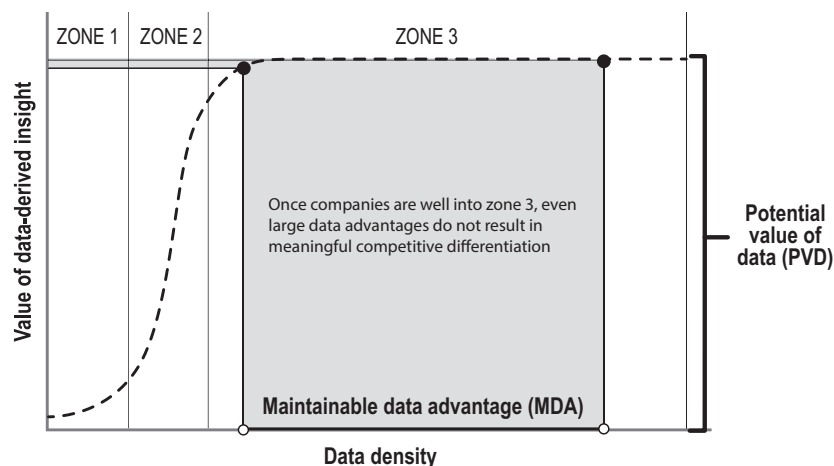
So, when an advantage can be found (put another way, when the PVD is high) the data-accumulation process must be sufficiently difficult that the data-feedback company is able to progress along the curve (and capture a significant portion of the PVD), while its competitors are constrained from doing the same.

In today's information economy, data are generally cheap and just about all companies have data in abundance. What this means is that, most of the time, in situations involving data analysis, all firms end up in zone 3 – meaning they largely end-up knowing the same things. *In these cases, data does not serve as a source of competitive differentiation. Instead, it is simply a cost of entry.*

Data density and decay

As a result, *the key to understanding whether user data can serve as a source of sustainable competitive advantage – and whether it cannot – is data density.* Data density is driven by two separate processes: the rate of data collection and the rate of data decay. If the rate of data decay is low, then eventually all data collectors (even those with slow rates) will eventually end up in zone 3 with little competitive differentiation – whether the data advantage is maintainable or not. See Exhibit 19. If the rate of data decay is high, however, then it becomes possible to build a competitive edge by collecting data faster than anyone else.

Exhibit 19: In zone 3, there is little competitive differentiation between firms regardless of the MDA



The key to understanding whether user data can serve as a source of sustainable competitive advantage – and whether it cannot – is data density. Data density is driven by two separate processes: the rate of data collection and the rate of data decay. If the rate of data decay is low, then eventually all data collectors (even those with slow rates) will eventually end up in zone 3, with little to no competitive differentiation, as this exhibit shows.

Source: Goldman Sachs Global Investment Research

Navigational maps provide a simple example of how data decay affects whether it is possible to sustain a competitive advantage. Depending on the precise nature of the map's use, the user's sensitivity to accuracy and to how recently the data were collected changes, which thus changes the effective rate of data decay.

Navigational maps that are used to locate places or roads generally have a slow rate of data decay since new places and new roads are relatively infrequent occurrences. For example, it's equally easy to locate the Grand Canyon on a map of the United States today as it was 50 years ago. In past generations, it was common to find 10-year-old maps in cars that could be used during navigational emergencies.

Accordingly, in the case of simple navigation, the slow rate of data decay made it possible and relatively easy for all map providers to reach zone 3 (where little or no competitive advantage could be derived from differences in accuracy or timeliness). However, if maps are applied to more demanding problems – for example, to find the fastest route home through a busy city during rush hour – the dynamic changes.

In the case of real-time traffic navigation applications, like Waze or Google Maps, the data that are accumulated are subject to very high rates of decay, such that reaching zone 3 is difficult; this is particularly the case in terms of side routes, or when traffic patterns are changing rapidly. In this situation, the best vendor has a significant and self-reinforcing advantage. This is because these services often collect and analyze user location information to provide real-time navigation guidance; thus the more users any one service can attract, the faster their rate of data collection and the more accurate their insights, allowing them to move up the curve in zone 2 and to stay there as users congregate around the best provider. What's more, the concentration of users on one vendor lowers competitors' data collection rates and reduces the value of their data-derived insights (keeping them in zone 1), further reinforcing the lead vendor's edge (even on less used routes where data collection is more difficult).

A similar dynamic can be observed in web-based search. Early on, when web crawlers – a tool for indexing web pages to support search engines – were viewed as central to a vendor's competitiveness in the space, many vendors were willing to invest in developing the technology; the rate of change in web pages was sufficiently slow that reaching zone 3 was viewed as widely achievable. As it became clear that the searches themselves – particularly recent searches with a short-lived relevancy – were more useful for producing the most relevant search results, a clear self-reinforcing dynamic took hold. This was especially true in the case of popular or trend-based searches.

As a result, Google – which pioneered the use of its repository of searches to improve the applicability of its search results – has been able to develop a sustainable advantage in online search. Google’s ability to anticipate users’ keystrokes, highlight “hot” places to go or feature “trending” stories are examples of features that incorporate a large volume of data with a high rate of decay.

Ultimately, *this means that unless the rate of data decay is high for a given problem, all firms addressing that problem are likely to end up in zone 3 with little to no competitive differentiation. However, when the rate of data decay is high, a lead in data collection can become a self-sustaining and self-reinforcing advantage.* Data-density advantages only translate into competitive advantage in zone 2 and can only be maintained if competitors (particularly the runner-up) don’t make it to zone 3. We believe this is why, despite the general perception of the importance of user data, there are more examples of successful data-smart and data-asset companies than of successful data-feedback companies. It just isn’t that easy to find examples where the runner-up doesn’t eventually make it to zone 3.

Data-network companies

Data-network companies are similar to data-feedback companies in that they too leverage user data in ways that reinforce the value of their products or services. The primary difference is that *data-network companies use data to connect users to each other, which compares with data-feedback companies that use data to create output tailored to each user.*

For data-network companies, this means that *data density is defined by the number of active users, and the key driver of data decay typically has more to do with activity levels (rather than a change in the data).* Examples of data-network companies include Uber or Lyft, Airbnb and Facebook.

The competitive issues facing data-network companies are similar to those facing data-feedback companies. For both types of strategies, progressing out of zone 1 involves significant hurdles. After doing so, the firm's ability to build a sustainable competitive advantage is dependent on whether competitors are able to reach zone 3 – where differentiation is likely minimal at best.

One essential determinant of a data-network company is what defines an active user in the space. Another important consideration is whether being an active user in one network precludes or interferes with the user's ability to be active in another. If the networks are competing for users' time (as with Netflix, Facebook or the Fortnite Battle Royale game), there is a natural constraint that forces the system toward dominant vendors. However, if the service is consumed based on specific needs (as with Uber and Lyft or Airbnb and VRBO), the market is more likely to have multiple vendors that are in ongoing competition, and the network alone is unlikely to create a persistent advantage.

Communities of users – and the relevant boundaries – play an important role in driving the economics of data-network companies, which we addressed in more detail earlier in the organizer section. In some circumstances networks naturally divide into communities in which there is an advantage in specializing in providing network services *within a community* rather than to the general population. Modern dating applications – like Bumble, Tinder, Coffee Meets Bagel and e-Harmony – are simple examples of data-network businesses whose success is determined by active users.

For data-network companies, the ability to monitor and regulate membership can become a sustainable advantage; the ability to offer high-quality drivers, rental spaces, vendors, or other specific community affiliations may represent a key competitive strength.

The Everything-as-a-Service Economy

In such cases the data-network company is essentially mixing two types of learning-company data models: data-network strategies that are based on the directory of users with data-asset strategies (using reviews, background checks) to police the service. The mix can yield a hard-to-replicate business model.

In summary, analysis of the learning curve leads to a four-part test companies can take to determine whether data-based strategies can create a sustainable competitive advantage:

- First, are there sufficient data to analyze?
- Second, are the insights gained from such data analysis novel enough to create significant value?
- Third, is the implementation of those insights complex enough to prevent competitors from simply copying the approach?
- Fourth, are the data scarce enough that a competitor cannot repeat the same analysis?

If each of these questions elicits an affirmative response, building a sustainable competitive edge through data is possible. However, more often than not, this is unlikely to be the case, which means that data-based strategies tend to be a cost of entry rather than a source of sustainable competitive advantage and that robust second-mover strategies may be more cost effective than first-mover ones.

■ Concluding Thoughts

As the business environment continues to evolve, it is important that companies and their investors consider the many ways that disruption has and will continue to reshape the modern business landscape.

In this publication, we have explained some root causes of the increase in disruption we see today – namely, the ongoing process of disentanglement, which is often enabled and accelerated by the introduction of new technologies. We expect this process to continue until the most optimal and efficient outcomes are achieved. In other words, continued disruption is highly likely.

We also addressed new forms of businesses that are likely to thrive in the new economy – which is characterized by the fact that nearly anyone can purchase virtually any service from a third-party. We identified four sources of ongoing economic advantages that can be exploited by companies to achieve long-run success in today's business ecosystem: economies of scale (platform companies), economies of scope (servicer companies), economies of fit (organizer companies) and economies of learning (learning companies). We also provide a new "learning curve"-based framework for companies to

The Everything-as-a-Service Economy

determine when data can serve as a source of competitive advantage – and when data cannot.

In the end, the new business environment is inherently more complicated than the old one. Companies that are competitors in some areas are often customers in other areas with greater frequency than in the past. Smaller businesses operating in smaller markets may be the most likely to breach standards around appropriate competitive behaviors. And what a company actually does isn't always apparent from looking at its revenues. As well, capital-light companies with limited footprints may still be big – meaning they may generate substantial revenues and be global.

With all of this in mind, we believe that companies can be successful going forward by finding their competitive advantage and sticking to it, while outsourcing the rest.

■ Bibliography

Bellini, H. (Jan 2015). *Cloud Platforms - Volume 1: Riding the Cloud Computing Wave; RHT down to Sell*. Goldman Sachs Global Investment Research.

Bezos, J. P. (Apr 2018). *2017 Letter to Shareholders*. Retrieved from <https://blog.aboutamazon.com/company-news/2017-letter-to-shareholders>

Borst, D. (Oct 2015). *Americas: Entertainment: Media Restack: TWX off CL; Downgrade VIAB, SNI; Upgrade OUT*. Goldman Sachs Global Investment Research.

Buffett, W. E. (Feb 2018). *2017 Letter to Shareholders*. Retrieved from <http://www.berkshirehathaway.com/letters/2017ltr.pdf>

Burgstaller, S. (May 2017). *Rethinking Mobility: The 'pay-as-you-go' car: Ride hailing just the start*. Goldman Sachs Global Investment Research.

Cabral, M. (Jun 2014). *Commentary: Apple 2014 WWDC: Improved integration across iOS and Mac OS; platform extends into Health and Home*. Goldman Sachs Global Investment Research.

Chamberlin, E. (1933). *The Theory of Monopolistic Competition*. Cambridge Harvard University Press.

Collett, M. (Aug 2015). *The Infinite Shelf: E-commerce threatens market fragmentation for consumer staples*. Goldman Sachs Global Investment Research.

Costa, D. (Apr 2016). *Factory of the Future: Beyond the Assembly Line*. Goldman Sachs Global Investment Research.

Della Vigna, M. (Apr 2018). *Top projects 2018: Forget about the rocks. It's all about consolidation*. Goldman Sachs Global Investment Research.

The Everything-as-a-Service Economy

Della Vigna, M. (Mar 2018). *A new era for oil investing: The Age of Restraint*. Goldman Sachs Global Investment Research.

Della Vigna, M. (Nov 2017). *Big Oils rise again as the industry consolidates*. Goldman Sachs Global Investment Research.

How Netflix Works With ISPs Around the Globe to Deliver a Great Viewing Experience. (Mar 2016). Retrieved Nov 15, 2018, from Netflix:

<https://media.netflix.com/en/company-blog/how-netflix-works-with-isps-around-the-globe-to-deliver-a-great-viewing-experience>

Keung, R. (Jul 2018). *E+Commerce/Logistics of Things (LoT) in China: Shopping & Delivery Re-Imagined (II)*. Goldman Sachs Global Investment Research.

Love, J. F. (Aug 1995). *McDonald's Behind the Arches*. Bantam Books.

Mims, C. (Sept 2018). The Prime Effect: How Amazon's Two-Day Shipping Is Disrupting Retail. *The Wall Street Journal*.

Netflix Case Study. (2016). Retrieved Nov 15, 2018, from Amazon Web Services:

<https://aws.amazon.com/solutions/case-studies/netflix/>

Pisano, G. P., & Shih, W. (2012). Producing Prosperity: Why America Needs a Manufacturing Renaissance. *Harvard Business Review*.

Porter, M. E. (1985). *Competitive Advantage: Creating and Sustaining Superior Performance*. New York: The Free Press, A Division of Macmillan, Inc.

Ramsden, R. (May 2018). *Banking on technology: The shareholder benefits of a digital future*. Goldman Sachs Global Investment Research.

Reeves, M., & Deimler, M. (Jul-Aug 2011). Adaptability: The New Competitive Advantage. *Harvard Business Review*.

Richter, S. (Apr 2018). *The Genome Revolution: Sizing the genome medicine opportunity*. Goldman Sachs Global Investment Research.

Richter, S. (Jan 2018). *Vol. 2: Setting the stage for oncology M&A in 2018*. Goldman Sachs Global Investment Research.

Rothenberg, R. (Feb 2018). *The rise of the 21st century brand economy*. IAB.

Rubin, J. (Nov 2016). *Johnson & Johnson: In talks to buy ATLN*. Goldman Sachs Global Investment Research.

Schneider, J. (Aug 2017). *Future of Finance: Payment Ecosystems*. Goldman Sachs Global Investment Research.

Shelly, J. (Nov 2017). *2017 Curalate Consumer Survey: Social Content is the New Storefront*. Retrieved Oct 18, 2018, from Curalate: <https://www.curalate.com/blog/social-media-content-survey/>

Spangler, T. (May 2018). Netflix Content Chief Says 85% of New Spending is on Originals. *Variety*.

Terry, H. P. (Jan 2013). *Amazon.com Inc. (AMZN): Assuming coverage: Remain Buy rated on ecommerce leader*. Goldman Sachs Global Investment Research.

Terry, H. P. (Jan 2017). *Netflix: 2017 content payoff to drive acceleration in subscriber growth; Buy*. Goldman Sachs Global Investment Research.

Terry, H. P. (Mar 2015). *The Future of Finance: The Socialization of Finance*. Goldman Sachs Global Investment Research.

The Everything-as-a-Service Economy

Walvis, A. (Jun 2018). *Apparel and Accessories: Initiates with an Attractive coverage view*. Goldman Sachs Global Investment Research.

Welson-Rossman, T. (Oct 2018). Tutoring for the Modern Age. *Forbes*.

Wohlsen, M. (Jun 2014). A Rare Peek Inside Amazon's Massive Wish-Fulfilling Machine. *Wired*.

■ Disclosures

Third party brands used in this publication are the property of their respective owners, and are used here for informational purposes only. The use of such brands should not be viewed as an endorsement, affiliation or sponsorship by or for Goldman Sachs or any of its products/services.

This report has been prepared by the Global Markets Institute, the public-policy and corporate advisory research unit of the Global Investment Research Division of The Goldman Sachs Group, Inc. ("Goldman Sachs"). Prior to publication, this report may have been discussed with or reviewed by persons outside of the Global Investment Research Division. While this report may discuss implications of legislative, regulatory and economic policy developments for industry sectors and the broader economy, may include strategic corporate advice and may have broad social implications, it does not recommend any individual security or an investment in any individual company and should not be relied upon in making investment decisions with respect to individual companies or securities.

Distributing entities

The Global Investment Research Division of Goldman Sachs produces and distributes research products for clients of Goldman Sachs on a global basis. Analysts based in Goldman Sachs offices around the world produce equity research on industries and companies, and research on macroeconomics, currencies, commodities and portfolio strategy. This research is disseminated in Australia by Goldman Sachs Australia Pty Ltd (ABN 21 006 797 897); in Brazil by Goldman Sachs do Brasil Corretora de Títulos e Valores Mobiliários S.A.; Ombudsman Goldman Sachs Brazil: 0800 727 5764 and / or ouvidoriagoldmansachs@gs.com. Available Weekdays (except holidays), from 9am to 6pm. Ouvidoria Goldman Sachs Brasil: 0800 727 5764 e/ou ouvidoriagoldmansachs@gs.com. Horário de funcionamento:

segunda-feira à sexta-feira (exceto feriados), das 9h às 18h; in Canada by either Goldman Sachs Canada Inc. or Goldman Sachs & Co. LLC; in Hong Kong by Goldman Sachs (Asia) L.L.C.; in India by Goldman Sachs (India) Securities Private Ltd.; in Japan by Goldman Sachs Japan Co., Ltd.; in the Republic of Korea by Goldman Sachs (Asia) L.L.C., Seoul Branch; in New Zealand by Goldman Sachs New Zealand Limited; in Russia by OOO Goldman Sachs; in Singapore by Goldman Sachs (Singapore) Pte. (Company Number: 198602165W); and in the United States of America by Goldman Sachs & Co. LLC. Goldman Sachs International has approved this research in connection with its distribution in the United Kingdom and European Union.

European Union: Goldman Sachs International authorised by the Prudential Regulation Authority and regulated by the Financial Conduct Authority and the Prudential Regulation Authority, has approved this research in connection with its distribution in the European Union and United Kingdom; Goldman Sachs AG and Goldman Sachs International Zweigniederlassung Frankfurt, regulated by the Bundesanstalt für Finanzdienstleistungsaufsicht, may also distribute research in Germany.

Goldman Sachs conducts a global full-service, integrated investment banking, investment management and brokerage business. It has investment banking and other business relationships with governments and companies around the world, and publishes equity, fixed income, commodities and economic research about, and with implications for, those governments and companies that may be inconsistent with the views expressed in this report. In addition, its trading and investment businesses and asset management operations may take positions and make decisions without regard to the views expressed in this report.

© 2018 Goldman Sachs.

No part of this material may be (i) copied, photocopied or duplicated in any form by any means or (ii) redistributed without the prior written consent of The Goldman Sachs Group, Inc.

“ For a firm to become a disruptor it simply needs to offer an incrementally better product or service than its competitors – and not even by that wide of a margin. Customers can and will switch to new products and services to realize incremental benefits – and they’ll do so for smaller gains than in the past. This is both an entrepreneur’s dream and an incumbent’s worst fear, since firms that fall behind risk rapid displacement.

Faced with the threat of disruption, it can be easy – but also potentially fatal – to default to what seems to be the only solution, namely speed and disrupting oneself... Instead, we suggest that companies should focus on and reinforce their key areas of strength, rather than attempt to shift gears and become something else entirely. ”



The Everything-as-a-Service economy allows companies to better direct their efforts to their core areas of competitive advantage. Doing so can pay enormous dividends and create stronger more secure market positions, but failing to do so can end in quick failure.