

Goldman
Sachs

GLOBAL
MARKETS
INSTITUTE

July 2019

A SURVIVOR'S GUIDE TO DISRUPTION



Steve Strongin
steve.strongin@gs.com

Amanda Hindlian
amanda.hindlian@gs.com

Sandra Lawson
sandra.lawson@gs.com

Sonya Banerjee
sonya.banerjee@gs.com

Dan Duggan, Ph.D.
dan.duggan@gs.com

Table of Contents

Chapter 1: Survivor’s guide - the short form	3
Chapter 2: Disruption’s evolutionary roots	9
Chapter 3: Perfecting Platforms	19
Chapter 4: Niche after niche - Organizers	32
Chapter 5: The competitive value of data	44
Chapter 6: Concluding thoughts	56
Appendix A: Considering communities	59
Bibliography	61
Disclosure Appendix	62

The Global Markets Institute is the research think tank within Goldman Sachs Global Investment Research. For other important disclosures, see the Disclosure Appendix.

Chapter 1: Survivor's guide - the short form

We examine how companies can reshape themselves to better compete in today's Everything-as-a-Service (EaaS) economy¹. In this new economy, firms can use services provided by other businesses to grow faster, while using less capital and fewer people than would otherwise be possible. Industries are reorganizing in response to these dynamics, and companies must adapt or risk falling behind.

EaaS can be thought of as an extreme form of outsourcing. In the past, firms would selectively outsource business functions to reduce costs, for example by outsourcing ancillary functions like operating a cafeteria within an office or by outsourcing labor-intensive but simple manufacturing processes. Over time, however, the high degree of standardization that has emerged across manufacturing, communications, data systems and user interfaces, among other areas, has made it possible to outsource virtually any business function. As a result, firms are now able to decide precisely which functions to keep in-house, and which functions they should allocate to external parties instead.

By leveraging other firms to provide core business functions, EaaS companies can scale their businesses at a faster pace than before *and* access expertise or technology that would otherwise have been out of their reach. EaaS businesses can thus do what they need to do to compete, not only more cheaply, but also more quickly and – most importantly – better than they could on their own.

The most radical and sometimes confusing aspect of these new business models is that they allow firms to organize themselves around their sources of competitive advantage. Firms used to structure themselves around what they sold – around the cars, steel or computers that comprised the bulk of their sales, as examples. As a result, nearly all industry participants ended up looking quite similar from a structural standpoint. Today, the best firms are typically organized around what they do well and – to the extent that they can – rely on other firms to do the rest.

To understand how different this new system can be from the old model, consider the smartphone industry. Some firms, such as Samsung, still look much like yesterday's fully-integrated producers (except that Samsung produces about as much product for its competitors as it does for itself, a dynamic we also explore). In contrast, Apple – the most successful of the smartphone producers – is organized entirely around the customer experience and does not manufacture any part of the iPhone. In fact, Apple has so little to do with the actual production chain that when a customer purchases an iPhone it is quite possible that no Apple employee has touched any part of that phone.

Importantly, the EaaS economy has also changed the nature of competition. In the past, firms could become market leaders with inferior products by having superior distribution or services capabilities relative to their competitors. Business strategy was often based

¹ Some of this material was included in Strongin, Steve et al (Dec 2018), *The Everything-as-a-Service economy*, Goldman Sachs Global Markets Institute.

on creating and maintaining barriers to entry rather than on creating the best product or service.

Today, the firm with the best product or service can rely on other firms via widespread outsourcing to defeat barriers to success. Whether a new entrant needs help to ramp its production, distribution, marketing or any other business function, it can do so at a world-class level, at a competitive price and at a speed that would have been unthinkable just a generation ago, simply by leveraging other firms' resources. As a consequence, the only way for a firm today to defend its market position is to actually have the best product or service. Anything less is likely to result in quick displacement.

The potential for rapid displacement is why "disruption" has become such a significant preoccupation for managements, entrepreneurs and investors alike. In its most extreme form, this preoccupation with disruption causes a preoccupation with speed, which is often expressed with phrases like: "disrupt yourself." However, once the source of the disruptive threat is understood, it is clear that quality – not speed – is the key to survival.

Companies in the EaaS economy cannot stand still, but they also should not give up their existing advantages in search of what is "new" simply for the sake of doing so. Instead of looking to "disrupt yourself," firms' new mantra should be: "do what you do best – but do it better – and outsource anything that hinders your success."

This begs the question: how can a firm figure out where to draw the line between what it should do itself and what it should exit? And how is the accumulation of these choices changing corporate and industry structures? To answer these questions, it can be useful to examine how the EaaS economy came to be (which we cover in detail in Chapter 2, entitled "Disruption's evolutionary roots"). The key is that firms are reengineering themselves to separate the functions that can be run more efficiently independently than when they are run together.

In most cases, this means restructuring the functions that gain efficiencies from scale ("scale businesses") and that derive profits from cost advantages, and separating them from functions that benefit from having niche audiences and charging premium prices ("boutiques"). It often becomes necessary for these functions to evolve into independent businesses to fully realize their potential.

The split that exists between scale businesses and boutiques defines almost everything about how EaaS companies are run. This ranges from how each type of firm relates to its customers, to how each type of firm invests its time and its capital. Much of the confusion in the marketplace about what makes good or bad business models or practices is due to the fact that there is no one-size-fits-all approach – and what a good scale company should do is often the exact opposite of what a good boutique should do.

An interesting example of how differently scale and boutique businesses approach products – even when those products directly compete with one another – arises in the media industry. Traditional broadcast companies (scale businesses) want programs that attract the widest possible audiences in order to generate maximum advertising revenues per hour of broadcast. In contrast, new subscriber-based media firms

(boutiques) need incremental subscribers and thus search for niche programs that have distinct audiences and that might have little or no broad appeal (we look at this dynamic in detail in Chapter 4 in the section entitled “Finding the next niche – or why TV is getting weirder”).

Optimizing an EaaS business requires identifying the precise nature of its competitive advantage and focusing relentlessly on enhancing and leveraging that advantage, which necessitates avoiding misusing resources (intellectual or financial) on distractions. This is not a trivial problem because one firm's distraction is another firm's strategic opportunity.

Firms that are successful in this regard are likely to become extreme versions of traditional scale or boutique businesses, which we refer to as **“Platforms”** or **“Organizers.”** To that end, we delve into the economics underlying these EaaS business models that drive sustainable competitive advantages in the EaaS economy. We also examine how a firm can self-evaluate, define and focus on its core – with a particular emphasis on distinguishing between the activities that are likely to provide synergies and the activities that are likely to represent a waste of resources.

The EaaS transformation – from integrated producers to Platforms...

In the past, a firm's scale was driven by its ability to compete as a vertically integrated entity. In the EaaS economy, some scale companies, having been freed from the need to develop, sell or maintain their own products or services have morphed into **Platform** companies. For these new EaaS-enabled Platform companies, success is mostly dependent on their ability to construct efficient portfolios of activities (this is described in detail in Chapter 3, entitled “Perfecting Platforms”).

The new economics of Platform companies explain why large pharmaceutical companies now tend to acquire overlapping drugs within a disease category rather than investing broadly to develop a diverse array of non-overlapping drugs (because it lowers portfolio risk). They also explain why cloud-based IT services typically want customers with different demand patterns (because it allows the cloud-based services provider to improve its load balancing and attain higher capacity utilization rates).

The new economics of Platforms also explain why payment processing companies tend to avoid extensive client customization (because achieving economies of scope requires a high level of standardization). More importantly, they explain why modern Platform companies happily act as suppliers to their toughest competitors (LG and Samsung selling their displays to Sony and Apple, as examples) as a way to reduce their exposure to shifts in market share.

The deeper (and sometimes odd) implications of Platform companies' need for scale are easiest to see when considering the random assortment of ancillary services these firms offer. Platform companies' non-core offerings, such as Amazon's music service or Google's Maps application, are often priced in a way that cannot – and is not intended to – produce reasonable returns on investment as standalone businesses would be intended to do. Understanding the pricing of these ancillary services requires not only

understanding the cost structure of the ancillary service itself, but also its impact on the firm's overall efficiency.

...and the transformation from boutiques to Organizers

Likewise, enabled by EaaS, some boutique firms have been able to significantly narrow the scope of their operations while also expanding their global presence. These super-charged boutique firms, which we refer to as **Organizers**, represent an even greater break from the past than Platform companies do. Freed from the need to build their own operating infrastructure, Organizers are able to focus narrowly on identifying and understanding their target customers.

By leveraging other EaaS firms to provide key business functions, Organizers are able to reach customers anywhere in the world, to supply products of any type, at any volume and of any technical complexity – all at globally competitive prices. This increased reach has transformed Organizer firms' potential size and profitability, allowing them to move beyond being niche locally constrained boutiques to being global super-competitors.

Organizers are built around the economics of fit. They create their superior performance by segmenting off specific customer groups that are willing and able to pay a premium for well-matched products and services. These are the companies that have become fixated on social media and on their user communities. But, the fit-based business model isn't new. In fact, the economic model describing this type of market structure goes back to Edward Chamberlin's "The Theory of Monopolistic Competition" from 1933, which described when firms create pseudo monopolies by segmenting client bases with distinct tastes.

Lululemon and Apple are good examples of these new Organizer companies. Each firm has created niche products for specific customer communities for which they charge premium prices, while using other firms to actually produce those products.

Netflix is an extreme example of the Organizer business model. Netflix is essentially a virtual company – one that uses external production companies to provide content and Amazon's Web Services platform for sales and distribution – as is typical of an Organizer. What makes Netflix unique is the fact that its core business model is to identify multiple niche audiences and then to license and commission specialty programming for each of those niche communities. Usually a firm can bond to only one target group at a time without creating counter-productive brand confusion, but in Netflix's case, the community's bond is to the programming rather than to Netflix itself. This allows Netflix to repeatedly identify and then supply programming for one niche group after another.

For Organizers, success requires understanding what makes good versus bad community designs and how and when a firm can create and nurture such communities to their own advantage (this is covered in Chapter 4, entitled "Niche after niche - Organizers"). The key is finding communities with distinctly different needs from other communities, but where the needs are common within the single community.

As the target community becomes larger, it tends to become less distinct and less uniform in its needs. And as a firm's offerings expand to meet the needs of more

consumers, the needs of the target community tend to blur and become more general and the Organizer tends to become less profitable as a result. Ultimately, the Organizer profits by being narrow and targeted and by charging premium pricing, while Platform companies' profits are driven by operating efficiencies and thus are dependent on serving a broad and diverse customer base.

The fundamentally different economics of Platforms and Organizers help explain why Netflix (an Organizer) can run profitably on Amazon Web Services. Netflix can do this even while Amazon (a Platform) essentially gives away its competing media streaming offering – Amazon Prime Video – to incentivize consumers to become Prime members on its retail platform.

The competitive value of data

In Chapter 5, entitled “The competitive value of data,” we examine in detail how data creates value in the EaaS economy. After all, in a world where a company may never meet its customers or touch its products, data plays a special role in creating and driving coordination across firms. However, it is a very different thing for a company to be dependent on data than it is for a company to gain a competitive advantage from data. For most EaaS firms, data is a cost of entry – not the basis of a sustainable competitive advantage.

In Chapter 5, we present an analytic framework based on what we term the “learning curve” to help assess how and when data can go beyond being a simple factor to consider while running the firm and become an essential part of a sustainable competitive advantage. We use this framework to provide firms with a four-part test to assess the efficacy of data-based business strategies.

All four parts of the test must be true for data to serve as the basis of a sustainable competitive advantage:

1. Is there sufficient data to analyze?
2. Are the insights gained from the data novel and valuable enough to be of competitive benefit?
3. Is the data-derived strategy difficult to copy without the data?
4. Is the data sufficiently scarce or hard enough to collect that competitors cannot replicate the analysis in normal course?

The first three parts of the test are fairly straightforward: the data needs to exist, provide insights that are truly useful and those insights need to be hard to copy. The fourth part of the test is more subtle because data – particularly about technical issues that are not proprietary in nature – often accumulates over time, which means that all market entrants that are willing to wait (and can afford to do so) can eventually participate competitively. However, there is one case where non-proprietary data can create a powerful virtuous cycle of competitive advantage, which is when only recent data is useful.

To that end, there is a notion of “data decay” that is essential to many non-proprietary data-based strategies. For these strategies, rapid data decay means that better data collection will result in superior output, which attracts new users, further improving the related data collection and – perhaps most importantly – depriving competitors of that data. This virtuous cycle can result in highly profitable and difficult to copy business models.

But such applications are uncommon, which means that most data-based strategies cannot produce a sustainable competitive edge. Instead, data often serves as a cost of entry to a given market and second-mover strategies (copying a data-based strategy) rather than first mover strategies (creating a data-based strategy) are likely to be more efficient over the long term.

Final thoughts

For firms today, surviving disruption is largely about giving up functions best left to others. Failure is usually the result of attempting to preserve functions that have become dead-weight and that distract the firm from enhancing its core advantages.

Consider the thousands of merchants that now have global reach both through Amazon's retail services platform and through Google or Facebook for advertising. As well, consider how Walmart has refocused away from items that are easy to ship in a box and has levered its store-based logistics expertise to become a significant force in fresh food. The EaaS economy allows firms to leverage their core functions at speeds and reach that would have been unthinkable in the past. However, that also means that the core needs to be competitive at that same scale.

In this light, it is worth pondering exactly which firms the biggest disruptors have disrupted. Which companies precisely have Google, Facebook or even Amazon displaced? A close look at the marketplace shows that these companies forced other firms to adapt (sometimes painfully) to narrower and more efficient business models, but that many of their major competitors still exist.

Ultimately, in the EaaS economy, firms need to define their target functions and markets such that they are better and can stay better than their potential and actual competitors. This requires both greater and lesser ambition: greater ambition because firms need to be the best at what they do and less ambition because they also need to pick narrower niches in which to compete.

Chapter 2: Disruption's evolutionary roots

The focus of this chapter is on how the structure of the economy has evolved into today's EaaS form, where nearly any business function can be outsourced. We examine the ongoing process of "disentanglement" as the driving force behind EaaS business models – and of the accelerating pace of industry disruption – using a series of illustrative company-specific case studies.

As the case studies will show, while the emergence of these new business models may appear sudden, they are the product of a natural progression over years of companies reengineering their operations in pursuit of greater efficiencies. Each new reengineering effort enables firms to disentangle and standardize parts of their business activities. Such standardization, in turn, makes it easier for firms to both reengineer and to externalize other parts of their operations. As a result, the process of disentangling firms' production stacks has proven to be both self-reinforcing and accelerating.

Today, across industries, a firm's production stack can be almost entirely disentangled or standardized, which makes it both possible (and competitively necessary) for firms to keep in-house only the layers of their stacks that are competitively differentiated – while outsourcing the rest. In the pages that follow, we provide detailed examples of how disentanglement works and its economic drivers, with a particular emphasis on the strategies that firms can use to successfully adapt to these changes.

Why do we focus on disentanglement?

The reason we focus on disentanglement is because it allows – and in some ways forces – deep changes to the way firms and their industries are organized. As a firm disentangles its processes into separate layers, it becomes possible for the firm to better utilize the scale economies that are natural and optimal for each one.

The key point is that, when entangled, every part of a firm's integrated production stack has to operate at the same scale. This entanglement constrains the parts of the stack that have strong economies of scale, keeping them too small to achieve maximum efficiency. It also affects the parts of the stack that are subject to diseconomies of scale, pushing them to operate at too large a scale – beyond what is efficient.

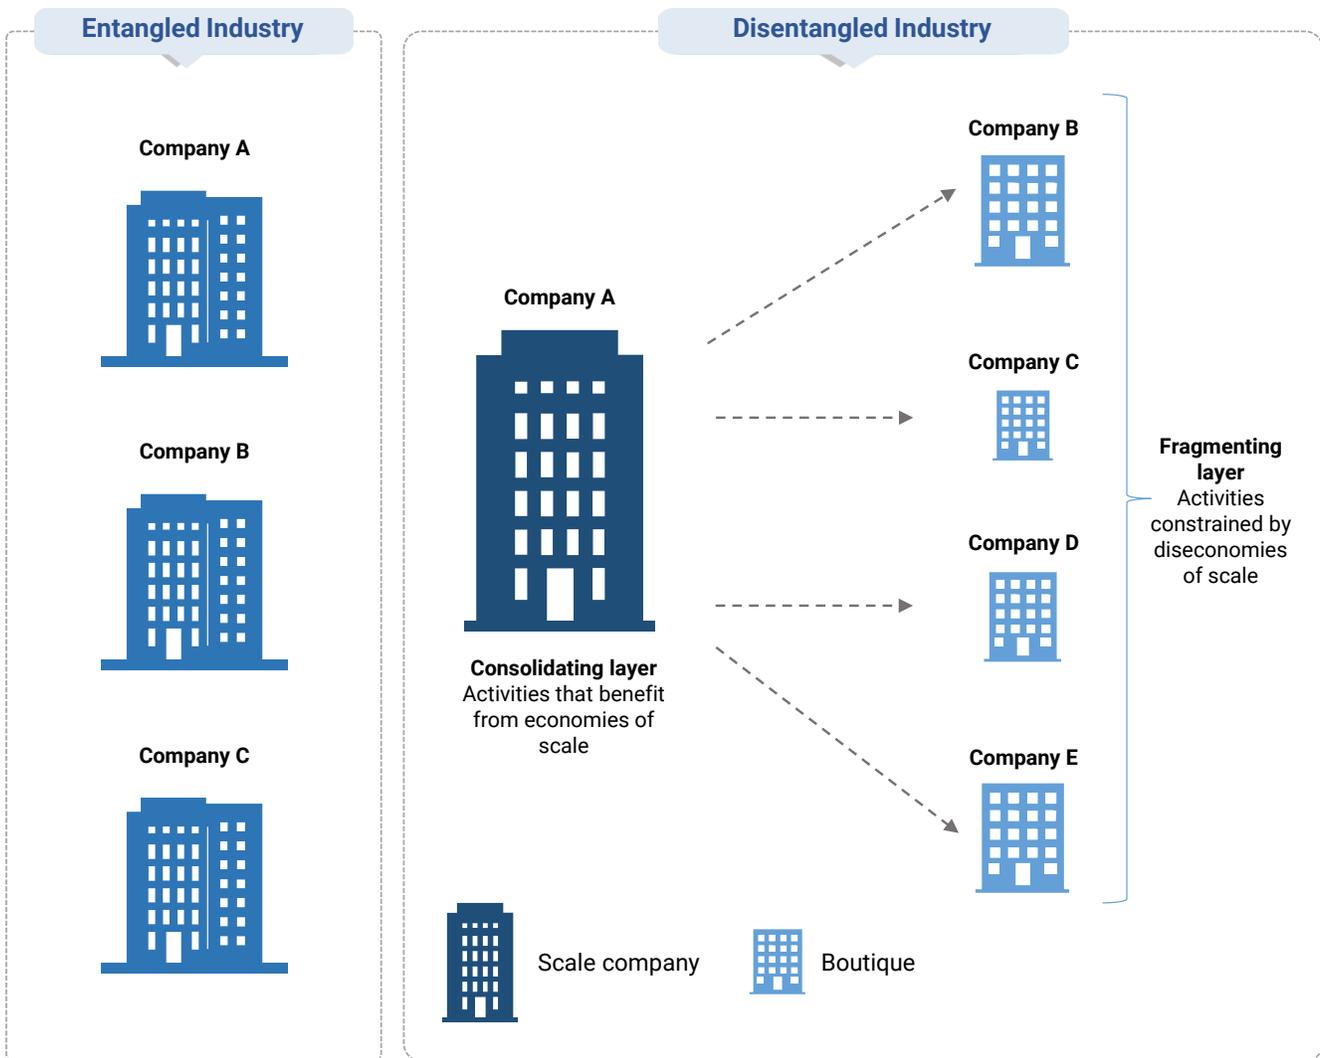
Once disentangled, the different layers can become independent businesses. This allows the prior entangled industry to reorganize into two disentangled industries, which we term the "consolidating layer" and the "fragmenting layer."

In the consolidating layer, firms grow by providing services to multiple other firms, consolidating the activity and gaining scale efficiencies. In the fragmenting layer, new entrants – which are now able to purchase key operational services from the disentangled scale suppliers – are in turn able to efficiently provide their products and services to new or increasingly narrow niches that previously would not have been efficient to service.

Net, the combination of the consolidating and fragmenting layers across industries has the effect of making companies more efficient. It has also led to the creation of a wider variety of high-quality products and services that are available to consume.

Exhibit 1 illustrates how disentanglement reshapes industries. The left side of the exhibit shows entangled firms within an entangled industry, where each firm is responsible for the entirety of its own production stack. When firms are integrated in this way, they are forced to operate all parts of the stack at a common scale, rather than operating each part at its own individual optimal scale. In comparison, the right side of the exhibit shows how the industry is reshaped when firms disentangle their operations and the industry splits into a consolidating layer (scale companies) and a fragmenting layer (boutiques).

Exhibit 1: How firm-level disentanglement ultimately reshapes industries



Source: Goldman Sachs Global Investment Research

Disentanglement and disruption: a deeper dive using case studies

McDonald's serves as an early example of modern disentanglement. At the time McDonald's began operating, food preparation was almost always done within the restaurants serving the food, with limited automation. The high turnover and local nature of food-service labor and of real estate kept the restaurant model quite local.

McDonald's was able to accelerate its growth when it reorganized to split the global operating company from the restaurants. The McDonald's franchise network now consists of thousands of independently owned and locally operated restaurants, while key operational functions – like marketing and franchise location decisions – are centrally managed.

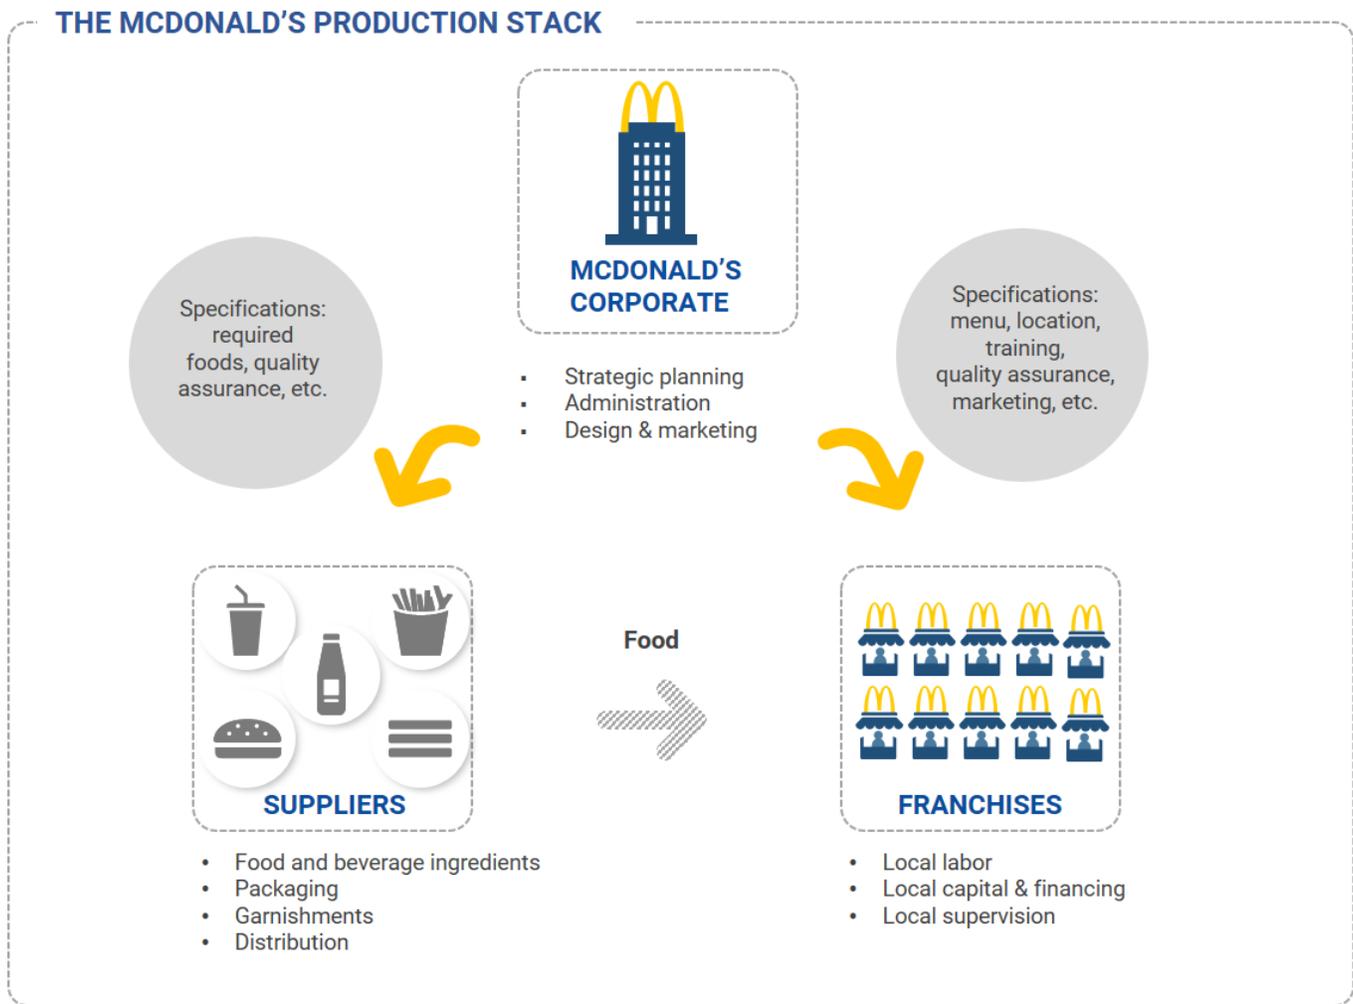
Case study 1: McDonald's

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 assembly-line system for food production, leveraged available technology and automation where possible – such as electric milkshake mixers – 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 (eventually 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 food went global. Today, more than 90% of McDonald's locations are franchises.

As shown in Exhibit 2, the three-part system McDonald's created – principally via franchising – consists of corporate global management, which centralizes functions like design and marketing; quasi-independent networks of food preparation on a regional level; and local franchises with local capital, labor and supervision. This three-part system is far more efficient at each level than nearly any other food organization had been before.

Exhibit 2: The McDonald's production stack demonstrates the benefits of disentanglement
 As McDonald's disentangled its operations, it enabled industry consolidation and fragmentation



Source: Company data, Goldman Sachs Global Investment Research

Each part of the McDonald's production stack is able to exploit very different scale economies. McDonald's corporate benefits from very high economies of scale related to its global organization, branding, and advertising. The food suppliers operate more regionally – thereby avoiding the problems associated with cross border, multi-product or other logistics diseconomies of scale, while still achieving greater economies of scale in production than an individual restaurant could achieve. Franchises leverage the global branding and regional suppliers and focus on the local labor and financing management issues that have the highest diseconomies of scale.

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

Industry disentanglement

The McDonald's case study is particularly instructive in that the efficiencies the firm gained through disentanglement extend beyond what would have been the case with simple process improvements.

McDonald's franchises operate more efficiently by utilizing franchise owners to access local lending and labor markets, which allows the global corporation to be relatively capital light and to avoid much of the human resource complexity that is typically associated with local recruiting efforts and employee training. This makes it easy to see how disentanglement made the parts of the McDonald's production stack more efficient.

At the same time, considering only the McDonald's case study could incorrectly give the impression that disentanglement operates within a set of related firms. In most cases, disentanglement leads to broader industry reorganization and to the entry of new competitors in the fragmenting layer.

A good example of the industry-wide implications of disentanglement is IBM's development of firmware in the 1960s. Firmware was largely developed to improve the efficiency of upgrade paths for IBM hardware, but it had the secondary effect of helping to create the modern software industry, which we discuss next.

Case study 2: firmware (IBM)

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.

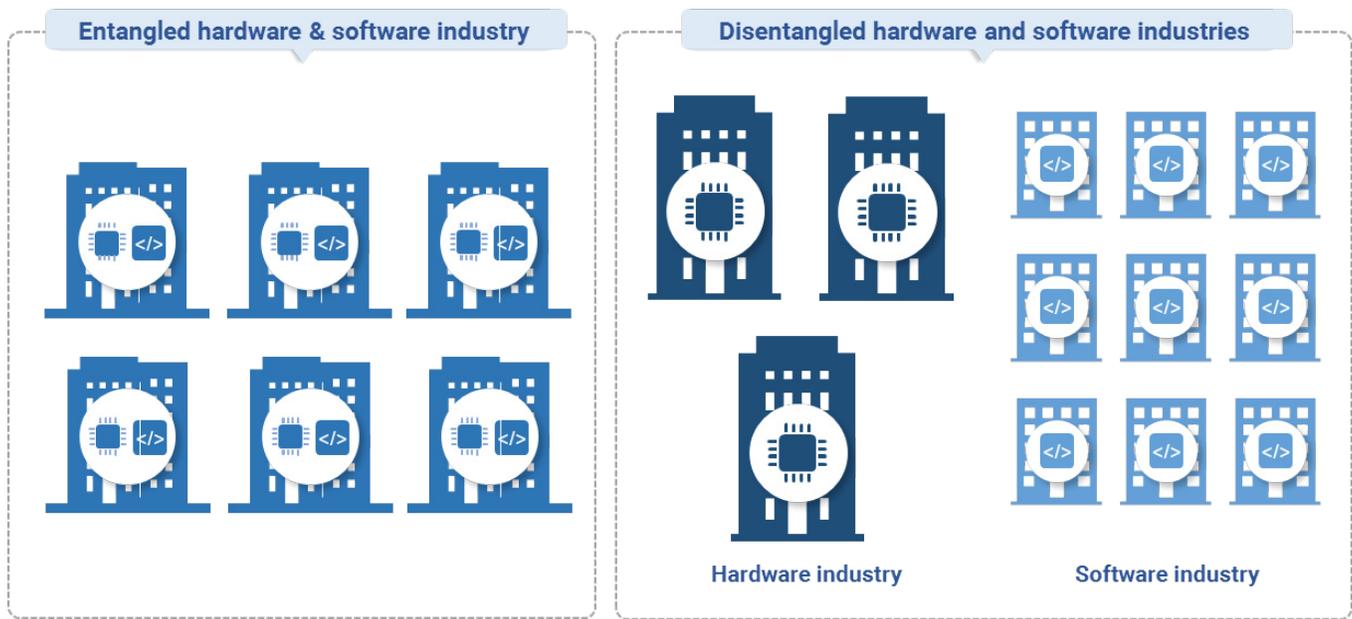
By standardizing how software accessed hardware, firmware allowed the same software to be used across an entire series of IBM mainframes. This meant that hardware changes could be made without also necessitating software investment, making it easier for IBM to sell hardware upgrades. It also made it 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 specialty software for corporate customers. This was a significant step toward creating the software industry as we know it today.

As Exhibit 3 below shows, 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).

Exhibit 3: The computer industry before and after firmware

Hardware is a natural consolidating layer while software is a natural fragmenting layer



Source: Goldman Sachs Global Investment Research

The super accelerant: making change cheaper

The development of firmware (and later of application programming interfaces or APIs) was an early sign of an entire era of standardization that was to come. Together, with the parallel process of standardization in manufacturing (see the box on ISO 9000 at the end of this chapter), goods production became more flexible and the speed at which companies could create and produce new products quickly and with limited capital increased substantially. However, better products alone do not create industry disruption – widespread adoption is also necessary.

As we discussed earlier, business strategy was in prior times often centered on using barriers to entry to limit competition and thus to retain customers, even with products or services that were marginally – sometimes even significantly – inferior to competitors' offerings. The innovation that largely eliminated barriers to entry across industries and truly accelerated industry disruption was the development of user interface standards.

User interface standards, unlike APIs or ISO standards, are not so much the product of committees or technical development groups, but rather have a distinctly more organic quality. Users – not developers – get to decide what works. Nevertheless, the various codifications of user standards were landmarks that have generated an order of magnitude reduction in the cost of change that has helped drive the rapid pace of disruption that defines the modern business environment.

From today's vantage point, it can be hard to remember (and even harder to understand), that users once required extensive training to operate most products or services, to shift from version 1.2 to 1.3 of a software package, or to accomplish simple tasks – to reserve a seat on an airline as one example (which was once only possible to

do through proprietary computer systems that required users to undergo weeks of onsite training).

Today, many firms claim credit for the innovations that made extensive user training unnecessary. But from the user's perspective, it was the widespread adoption of standards by most (if not all) providers that had the biggest effect – not their invention. Standardization made it possible for users to switch between systems, providers or even products and services without necessitating dedicated training and without significant costs. From that perspective, as we discuss in the next case study, operating systems from Microsoft and Apple clearly mark key shifts in the user experience.

Case study 3: user standards – Windows and iOS

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. In particular, the release of Microsoft's Windows 95 operating system was essential since it meaningfully simplified and standardized the personal computing user interface.

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, as examples) 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 EaaS economy.

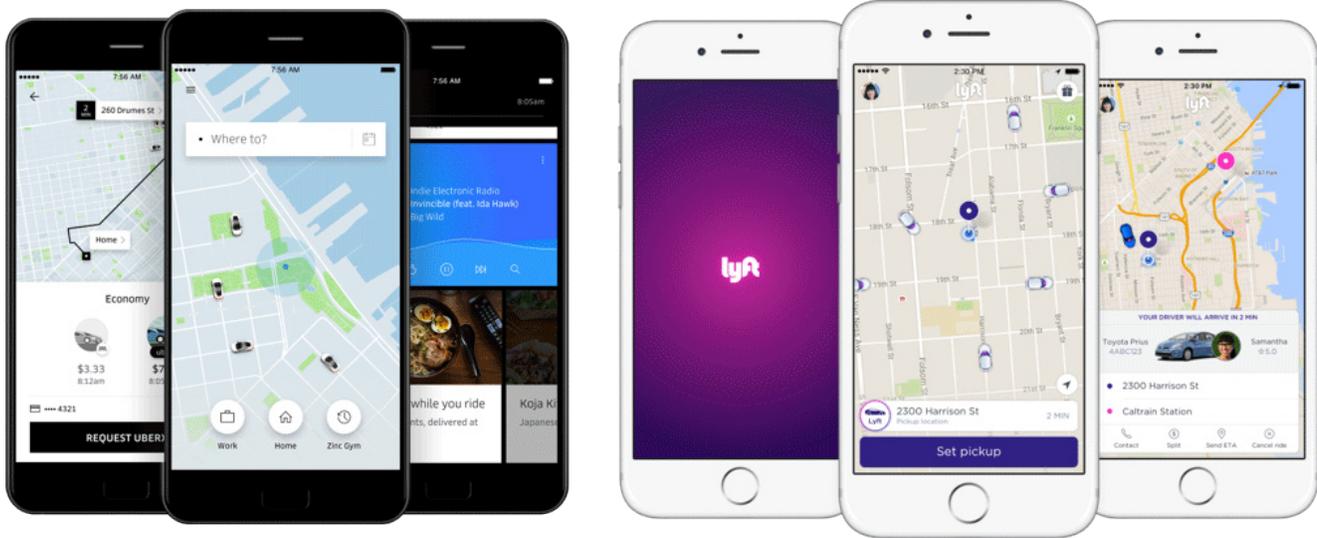
What's more, with Windows 95, Microsoft effectively kicked off the process of separating the functional (or technical) layer of software from the user interface. While not new technology, Windows 95 brought such interfaces into the mainstream. This process may have reached its fullest expression with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

While Microsoft effectively set the standard for simplified user interfaces, Apple introduced further innovation through its mobile touch-screen devices, as 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.

By producing devices that are functionally intuitive for users to operate, Apple helped to lower barriers to entry in a large variety of businesses – from banking to specialty retail – as well as switching costs between competitors' offerings in those industries. This ease of adoption has significantly contributed to the increasing pace of disruption that is

now the norm. Exhibit 4 below is an illustrative example of how standardized user interfaces make it relatively easy for consumers to switch between competing services, comparing two ride-hailing mobile application interfaces.

Exhibit 4: Intuitive user interfaces make switching between competing products easier
Snapshots of Uber's and Lyft's ride-sharing applications



Uber image: trademarked and owned by Uber Technologies, Inc.; Lyft image is from the Lyft press kit, June 2019.

Source: Uber, Lyft, Goldman Sachs Global Investment Research

From user standards to user-integrated production processes

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 navigate any changes. In some instances, this tactic could result in product “stickiness,” serving as a competitive barrier.

Ultimately, the evolution in 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. There has been significant growth in software development as lower switching and development costs have 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 applications should be usable 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.

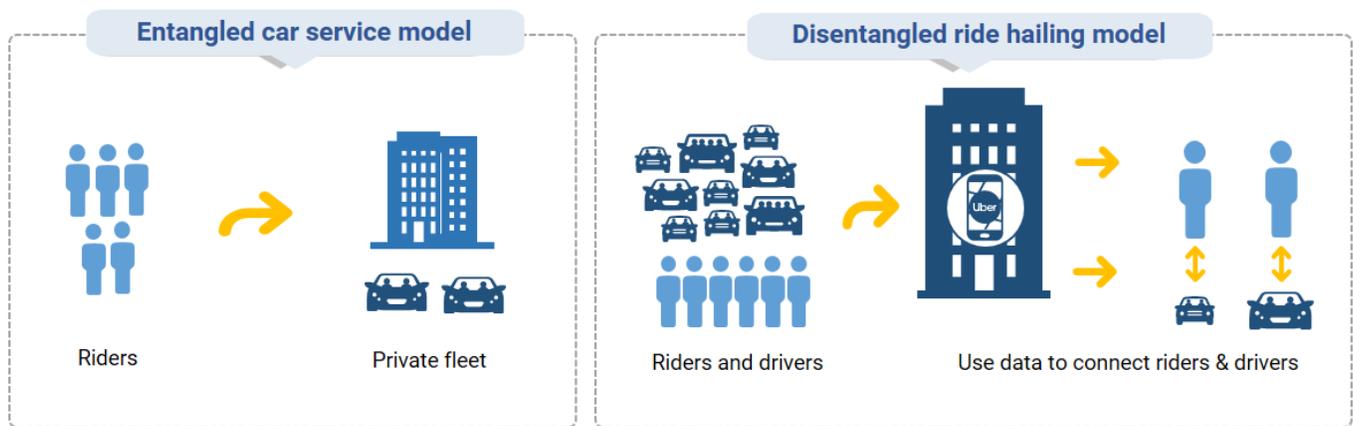
Likewise, modern ride-hailing services are beneficiaries of the user interface improvements we have described. In effect, modern user interfaces underpin ride-hailing companies' capital-light operating models.

Case study 4: ride-hailing services

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.

Modern ride-hailing services – like Uber, Lyft and Didi Chuxing – have been able to overcome these limitations by structuring themselves as matching platforms. These firms 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. We show how this new model compares with the old one in Exhibit 5 below.

Exhibit 5: How user interfaces reshaped transport



Source: Uber, Goldman Sachs Global Investment Research

These firms rely on users' willingness to leverage their software applications to reserve rides and to rate drivers (who, in turn, can also rate their passengers). 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." The breadth that this model enables is significant and is well beyond what McDonald's was able to achieve, as we discussed in the first case

study. Uber, for instance, has around 24,500 employees but is associated with nearly four million drivers globally and provides an average of 17 million trips each day.

ISO 9000 quality standards

While the last few case studies 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 EaaS 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 disentangling the layers of their production stacks, allowing consolidating and fragmenting layers to emerge. 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.

Chapter 3: Perfecting Platforms

Platform companies' source of competitive advantage comes primarily from the efficiencies that result from operating at scale – whether in terms of their production, marketing, risk management or capital management, among other capacities. Accordingly, the need to operate at scale drives many of these firms' strategic decisions and produces complex incentives.

There are two key steps to optimizing a Platform company: the first is to pinpoint the precise nature of the firm's scale-based competitive advantage, while the second is to structure the company to optimize the resulting efficiencies. For Platform companies, this inevitably means optimizing "the portfolio" as a whole rather than optimizing each business activity in isolation. To that end, the pursuit of greater scale often means sacrificing some efficiency at the micro level – potentially within a narrow business line – in order to gain efficiency at the aggregate level.

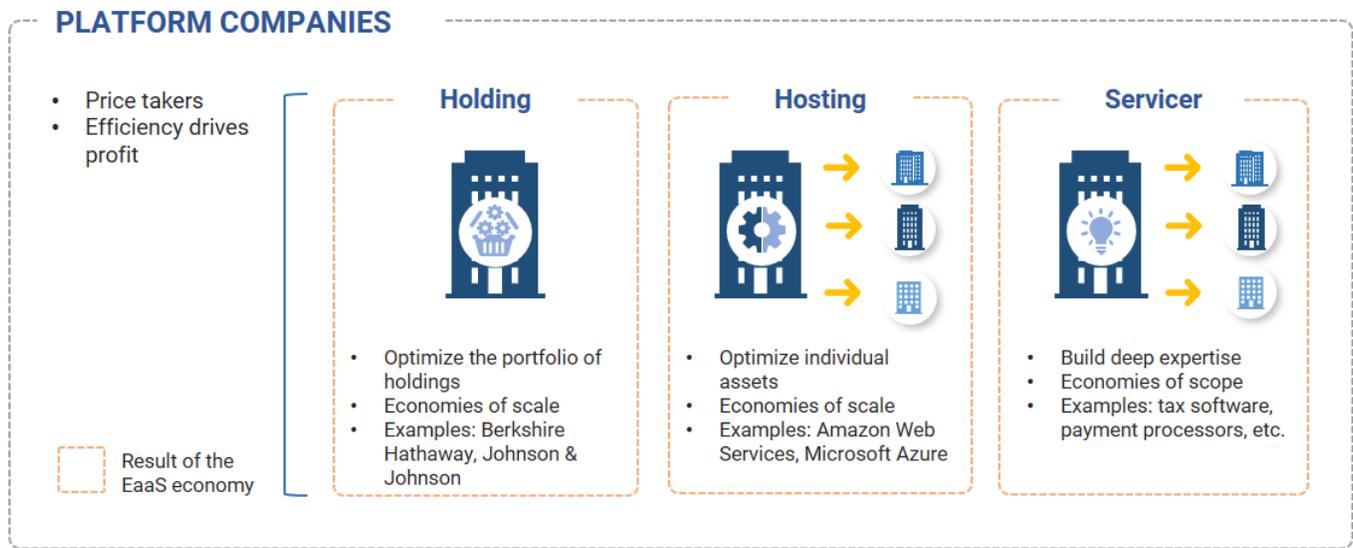
As an example of this "for the good of the whole" logic, consider a Platform company's decision to supply its goods and services to companies that are also competitors in some way – which is a fairly typical scenario in the EaaS economy. In enabling a competitor, the Platform company can hurt its competitiveness in its end markets but, at the same time, the Platform company also benefits from increasing the volume of its core business activity.

Whether engaging in this activity is rational depends on whether the net effect is positive. And, in a world where product standards make using other vendors easy and users' switching costs are low, the scale benefits that arise from selling to competitors typically do outweigh the costs. For the Platform company itself, the most difficult aspect may be to maintain its portfolio focus and to resist taking a more micro perspective.

Broadly, there are three archetypes for Platform companies: Holding companies, Hosting companies and Servicers. We will discuss each one in turn, but note that in practice these archetypes can overlap. Exhibit 6 captures key points for each model.

- **Holding companies** focus on owning diversified portfolios of assets and prioritize managing these portfolios in their entirety, rather than each of the underlying assets independently. Examples: Berkshire Hathaway, Johnson & Johnson.
- **Hosting companies** focus on having a diversified customer base, which allows them to manage their individual assets as efficiently as possible. Hosting companies seek to achieve high levels of physical asset utilization by distributing usage across as wide a customer base as possible. Examples: Amazon Web Services, Microsoft Azure.
- **Servicers** focus on providing their customers with some form of intellectual capital in a plug-and-play form, which allows Servicers to scale at very low marginal costs. Servicers seek to leverage this intellectual capital as much as possible and seek to expand their businesses by finding new customer types – but only ones that do not create diseconomies of scale by having unique customization requirements. Examples: tax software providers, payment processors.

Exhibit 6: The three types of Platform companies



Source: Goldman Sachs Global Investment Research

Regardless of the format, the source of competitive advantage for Platforms is the same: higher capital or capacity utilization rates. Scale – on its own – is not usually enough to create an advantage. Instead, it is the way that scale allows a firm to better utilize its assets that creates the efficiency-based advantage.

Thus, the underlying math for Holding companies, Hosting companies and Servicers is essentially identical: first, larger well-diversified asset portfolios tend to produce more predictable outcomes; second, firms that produce more predictable outcomes are able to plan more efficiently for their businesses and operate within tighter parameters. Taken together, these two factors increase Platform companies' overall capital efficiency.

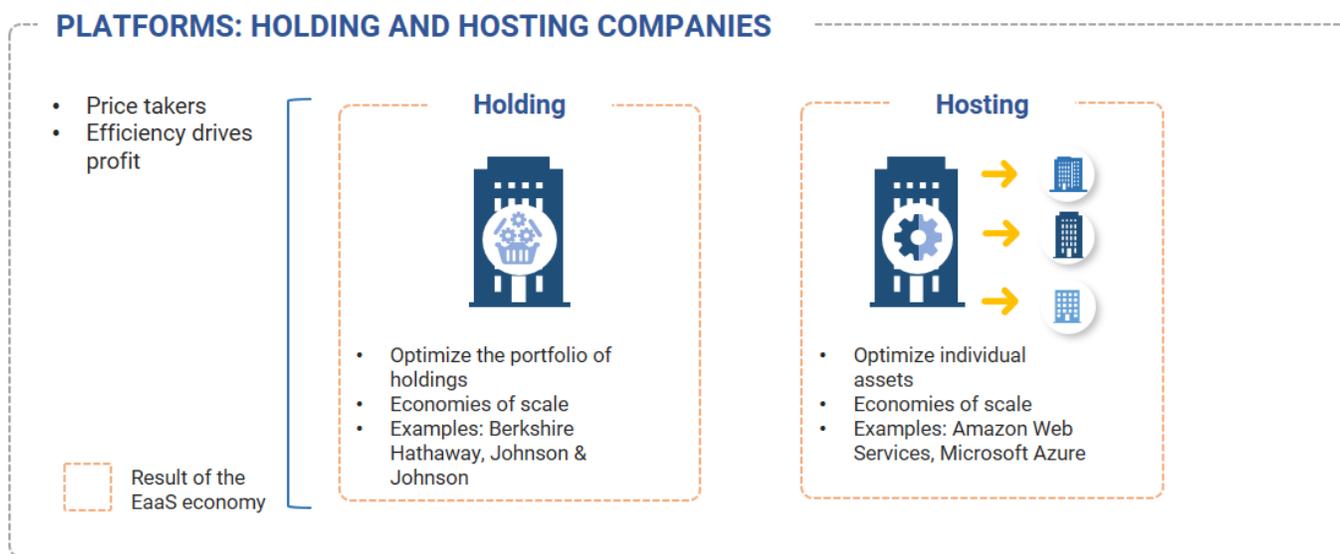
However, once the core business design of a Platform company is set, some types of incremental volume increase its efficiency while other types of incremental volume reduce its efficiency. Optimizing a Platform requires being able to distinguish between volume that is additive to the business and volume that is not, and weighting the business heavily toward the good.

Because of the relationship that exists between volume and efficiency, many Platform companies have a broad notion of complementary businesses that can be difficult to understand from an outsider's perspective. Accordingly, this often results in the appearance of haphazard design – even for the most disciplined companies. But, as we discuss further in the next section, this haphazard design actually has a rigid logic that good Platform companies apply ruthlessly.

Platforms: a closer look at Holding versus Hosting companies

We first look at Holding and Hosting companies – highlighted in Exhibit 7 below – as they share a common focus on constructing efficient portfolios of business activities. We address Servicers in the section that follows because optimizing Servicer companies tends to be more about optimizing the user populations (though it is worth noting that Hosting companies do sometimes need to worry about their user populations).

Exhibit 7: A closer look at Holding and Hosting companies



Source: Goldman Sachs Global Investment Research

For **Holding companies**, the role of portfolio selection is straightforward and is based on standard portfolio math. In the EaaS economy, however, the long-standing concept of efficient portfolio construction has been altered relative to the past. This is because each asset (or portfolio company) held by a financial Holding company can now be thinned and optimized so that it focuses only on the activities that benefit from the financial Holding company structure, while the remainder is outsourced. And as its portfolio companies restructure to refine their focus, the Holding company itself becomes both more efficient and better able to produce predictable returns.

In comparison, **Hosting companies** derive their primary cost advantages from “load balancing” – which involves distributing a given type of workload or activity across as broad a swath of assets with as predictable and level a load as possible. From the Hosting company’s perspective, this requires having a large number of customers (which may include competitors) with as diverse a demand profile as possible (particularly customers with different demand peaks), to generate sufficiently high utilization rates to create a meaningful cost advantage.

This notion of load balancing helps explain why, in the EaaS 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.

The precise way in which a Hosting firm can best organize its operations to maximize its capacity utilization depends critically on the uncertainty it faces. Consider markets in which market-share – rather than total demand – is the key factor in determining firm-level asset utilization rates. The best way to reduce risk and increase average utilization rates is to support the industry broadly, which may include selling to competitors (if their utilization patterns enable diversification) or by renting capacity from competitors (if doing so allows the business to scale up and down as needed).

For Hosting companies, another subtlety to constructing efficient portfolios of users is that their customer acquisition strategies can be priced based on customers' 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 (e.g. if 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 (in e-commerce versus online search, respectively) – 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 are, as examples. 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.

Continuing this example: a Hosting company that supports multiple live television streaming services gains efficiencies from consolidating varied peak activity on its own platform; put another way, being able to support a larger number of peaks – many of which are scheduled in advance and are spread out over time – can improve the Hosting company's level of 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 likely intensify peak usage, not diversify it.

Case study 5: Amazon's retail services business

Amazon's retail services exemplify the kind of 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.

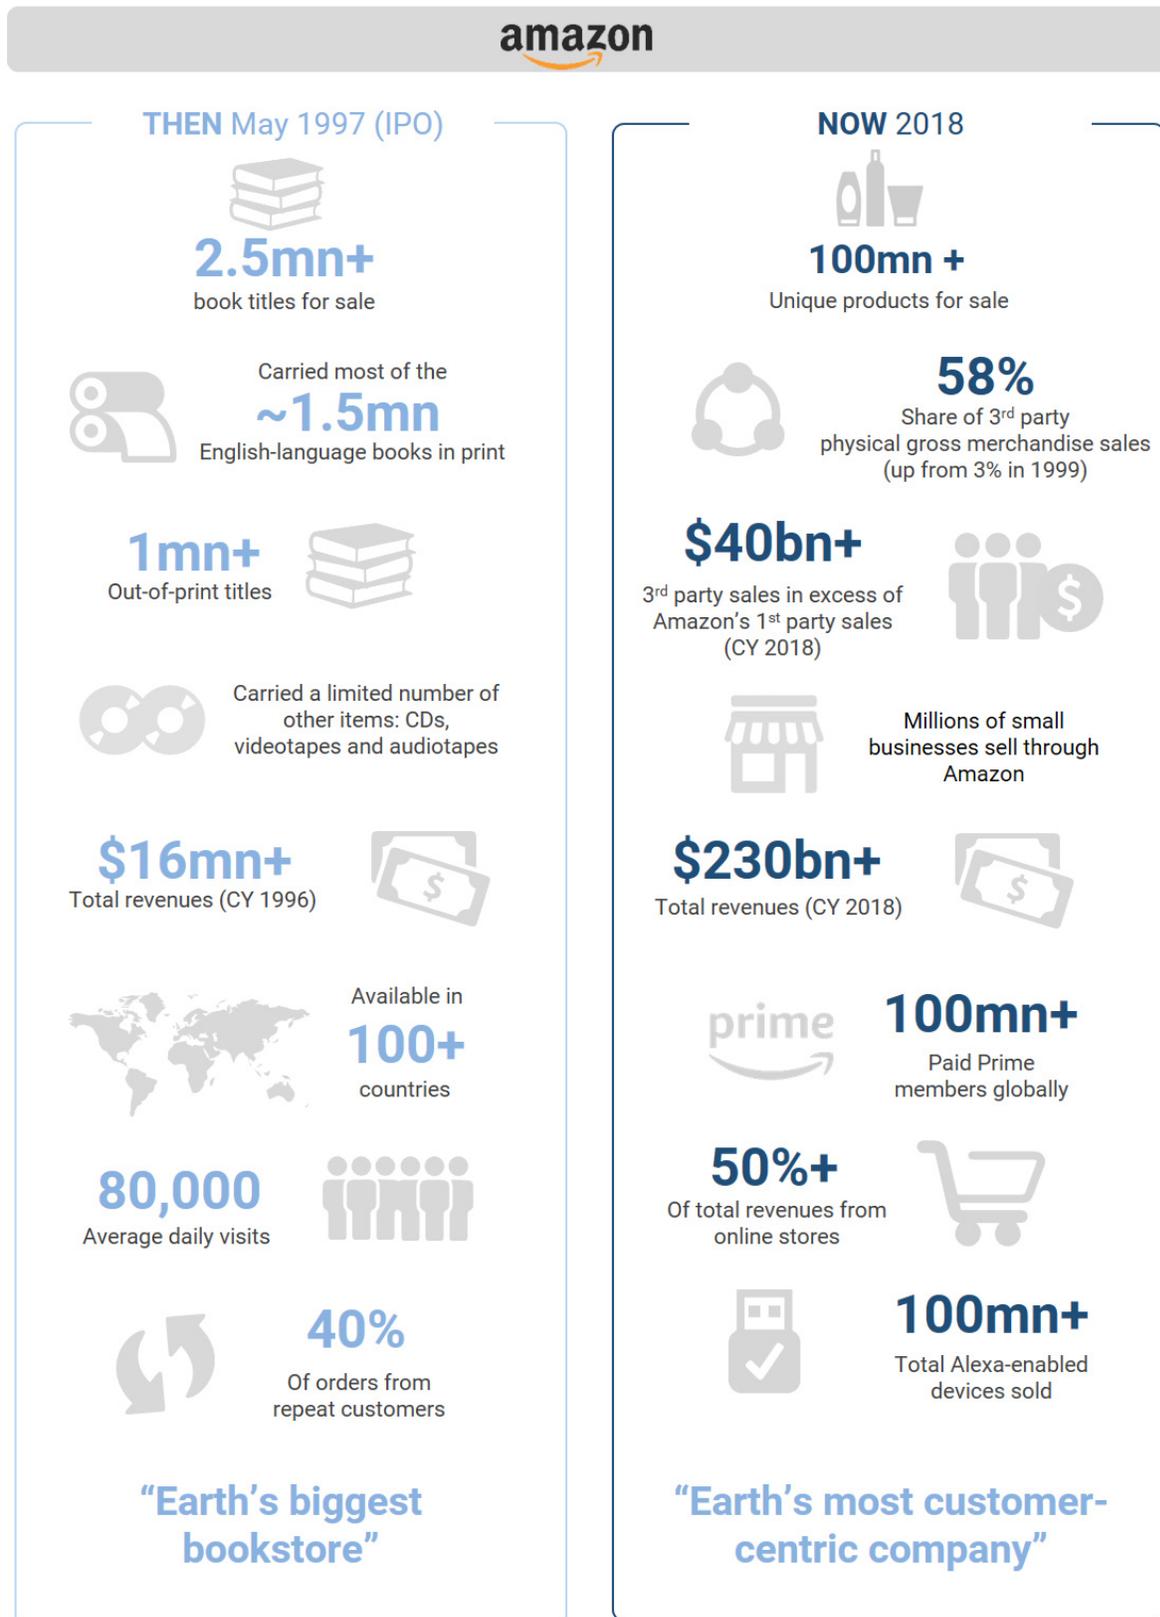
To clarify, Amazon's retail services include its e-commerce website and the underlying IT infrastructure, as well as its expansive warehouse and logistics system. While these assets underpin the firm's own retail operations, they also support a large and growing network of independent sellers. In fact, independent third-party sellers comprised 58% of total physical gross merchandise sales on Amazon in 2018, up from just 3% in 1999 (which is the earliest available metric).

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 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. Over time, the company expanded into retail categories outside of books, to include CDs, DVDs, videos, home goods, among other items. While diversifying its own retail inventory would have likely improved Amazon's asset efficiencies, the extent of such activity would have been limited by the capital investment and the carrying risk involved.

By gradually shifting to a Hosting model, and encouraging third-party vendors to sell through its e-commerce site and leverage its retail services, it is likely that 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 fragmenting layer) it has benefited from higher utilization rates of its e-commerce and logistics assets (which are natural consolidating layers). Exhibit 8 offers a snapshot of what Amazon's retail business looked like in its early days relative to today.

Exhibit 8: Amazon's e-commerce business, 1997 and 2018



Source: Company data, Goldman Sachs Global Investment Research

Despite the fact that Amazon now handles billions of unit sales, the firm has fewer than one thousand fulfillment centers globally. The placement of each one and the inventory management within are optimized 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 as a merchant relative to others, as we discuss in Chapter 5.

By operating as a Hosting company, Amazon has been able to take broad-assortment retailing to the extremes and take share in an established marketplace, against long-standing and mature market participants like Walmart. As well, an interesting illustration of how narrow and sensitive Platform companies' competitive advantages are can be seen by comparing Amazon with its Chinese counterparts Alibaba and Tencent.

Amazon leverages a well-developed truck-based delivery system and thus focuses on overnight shipment of boxes. In China, higher road traffic and cheaper labor make traditional and motorized bicycles more efficient for most delivery types. As a result, Alibaba and Tencent have found it much easier to deliver fresh groceries and prepared foods in China than Amazon has in the US, while Amazon has had a much easier time competing with traditional stores by delivering larger items such as mid-sized electronics and appliances in the US than Tencent and Alibaba have had in China.

Case study 6: Berkshire Hathaway

Berkshire Hathaway (Berkshire) exemplifies a typical financial Holding company. Berkshire's business model 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 predictable 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 EaaS economy, such that it 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 7: the pharmaceutical industry

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 portfolios of drugs.

There are a number of operating functions, like 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 experience has shown that disease incidence (or demand) is far more stable than the relative demand for particular drugs within a disease category. 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 counter-intuitively – reduces risk).

Consider Johnson & Johnson. 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 acquired 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.

A second source of efficiency 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 counter-intuitive at first glance. It is not completely clear if the diseconomies in research are due to operational or capital reasons. We look at the

possible reasons why capital costs for research may be lower in small firms than for large ones in our report “What the market pays for.”²

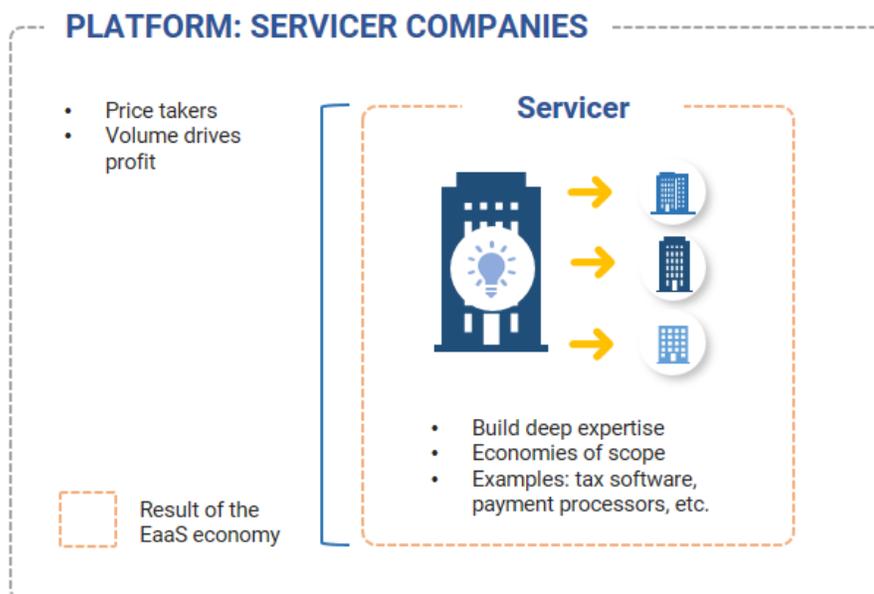
Likewise, the oil industry has seen a similar structural pattern with respect to the development of shale, where smaller companies appear to be better at finding and developing assets, but larger companies appear much better at exploiting capital and logistics efficiencies as the shale assets become more established. 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. These dynamics help explain the established patterns of merger and acquisition strategies in these sectors.

Servicer companies: exploiting economies of scope

Servicers are the intellectual-capital version of Platform companies and solve a newer problem that has arisen from the EaaS economy. Namely, given the barriers that the EaaS economy has brought down, an increasing number of companies are able to conduct commerce across multiple regulatory, tax, legal, communication and other technical environments – which necessitates being able to meet highly specific and often local business requirements. There is therefore more of a need for businesses that sell deep expertise in each of these areas.

This is where Servicer firms add the most value. These firms offer deep but narrow process expertise and can manage the critical business issues that would otherwise prevent many companies from being able to scale geographically. Exhibit 9 highlights a few of these points.

Exhibit 9: Servicers are the intellectual capital version of Platform companies



Source: Goldman Sachs Global Investment Research

² See Strongin, Steve et al (Apr 2019), *What the market pays for*, Goldman Sachs Global Markets Institute.

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 and therefore often compete with other Servicers on ease of use as well as on price.

In effect, Servicer companies are enablers of the EaaS economy. An outsourcer can hire technical services from a third-party Servicer firm, and while these technical services can be complex, inherently dynamic and scale on-demand, the outsourcer can treat them as simple and static (said another way, the outsourcer doesn't need to know how the Servicer's product works). This allows outsourcers to ignore complex non-core processes and to focus instead on their own core competencies.

As examples, 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 by outsourcing these functions to Servicers (the outsourcer doesn't have to concern itself with these functions).

Servicer companies invest in intellectual capital to ensure ease of use while connecting outsourcers with providers. Hosting and Holding companies, in contrast, 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. Consider that Stripe (a Servicer company) uses Amazon's cloud infrastructure (a Hosting business) to power its infrastructure even though Amazon competes in the payments space with its Pay offering. 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 expanding their customer base, 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.

Case study 8: Payroll software

Providers of payroll software – such as ADP, Intuit or Sage – are classic examples of Servicer firms. ADP, for instance, offers a broad suite of human capital management tools and its payroll services in particular are used by businesses around the world, of varying sizes, in a range of industries and with different types of employees.

These firms' payroll offerings are intended to be end-to-end solutions that allow customers to outsource payroll functions so that they need not dedicate in-house resources to doing the same. These solutions 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 deep expertise – and this is precisely where Servicers can come in handy in the EaaS 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, providers of payroll software must ensure that their technology can interact with a wide range of systems. This 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.

ADP and businesses like it, benefit from the economies of scope associated with having deep payroll expertise that can be leveraged across an expansive customer base. At the same time, these firms' 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.

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 services providers – and across many different types of systems.

Case study 9: payment processors

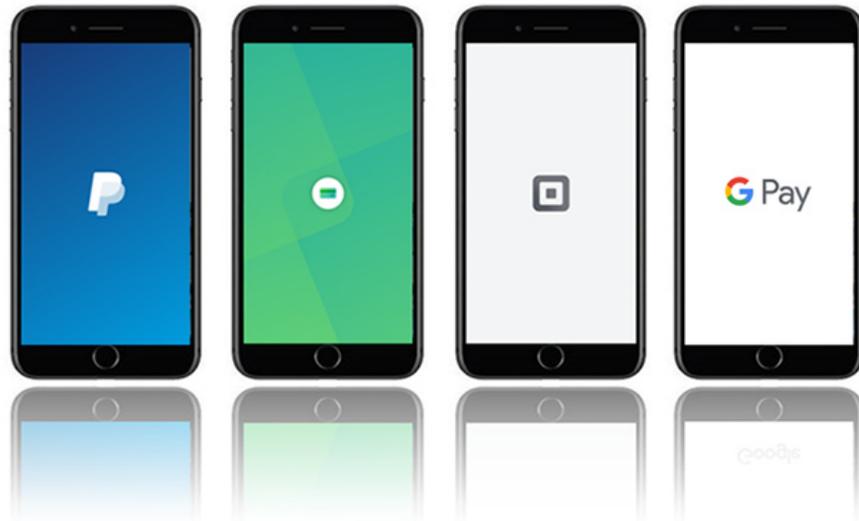
Consider PayPal as one specific example, given its market leading position, although there are a number of other firms that offer similar payment processing services including Stripe, Square and Venmo for instance. PayPal in particular connects millions of merchants and customers around the world with its technologies and facilitated nearly 10 billion payment transactions in 2018 in a range of currencies.

PayPal's users are able to leverage their accounts to receive and 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. To that end, PayPal has made significant investments in the technology infrastructure that underpins its payments solutions, allowing developers to build its payment solutions into their mobile or web applications through standard APIs.

The complex inner workings of PayPal's payment system – or of other similar firms' payments systems – are not evident to the merchants that incorporate these services into their production stacks, nor are they evident to end-users, since they leverage the modern user interface standards that have become the norm. This means that PayPal's customers need not navigate the underlying workings of the product, again freeing them to focus on other core areas.

Exhibit 10 illustrates the idea that standardized user interfaces obscure the underlying complexity and allow users to move easily between similar offerings.

Exhibit 10: PayPal, Stripe, Square and Google Pay are examples of Servicers
Standardized user interfaces obscure the underlying complexity



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

Case study 10: Qualcomm

Qualcomm is another example of a Servicer. 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 – 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.

Next we consider how cable companies – which are part of a mature and highly developed industry – are evolving in light of the EaaS economy. As programming has fragmented, some cable providers have begun to shed programming to focus more on their core scale businesses. There are significant implications for both the structures of these businesses as well as their potential profitability.

Case study 11: cable versus programming

For decades, cable companies – Comcast and Charter Communications, as examples – have sold their services as bundles inclusive of programming. Although content has long been critical to subscribers and live programming (for news and sports, as examples) remains a draw, the associated margins have meaningfully compressed over the last decade for even the largest distributors.

Underpinning this shift is the fact that the programming industry has fragmented. Not only is it now possible to view traditional TV content by subscribing to network-specific applications, but there are also numerous alternative services that provide a wide range of on-demand entertainment (Hulu, YouTube, Amazon Prime and Netflix among many others).

As programming has fragmented, some cable providers have begun to address programming separately from their infrastructure-based services, allowing them to focus more on optimizing their core scale businesses. Cable One, which is the seventh-largest cable company in the US (or the 12th largest pay TV video service across cable, telecommunications and virtual services), is an example of one such provider. The firm has begun to pivot away from the traditional bundle in order to sell its broadband services separately from content, which has been a low (even negative) margin business for the company.

Given how the cable industry is evolving, the key point is that the development of new technologies and the passage of time can dramatically alter firms' optimal decision-making strategies and can cause industries to reorganize. Consider the number of developments that needed to occur over time for the cable bundle to disentangle: programming contracts needed to come up for renewal, media streaming technologies needed to advance as did the array of network-connected mobile devices, and the content side of the industry related to programming needed to fragment with new pricing models, just to list a few examples.

Concluding thoughts

As we look at how successful Platform companies are managed, the aspect that is most striking is how narrow their business models actually are. Examples include: Amazon's focus on boxed products (Hosting company), Google's focus on search-based advertising (Hosting company), Stripe's focus on payments (Servicer), or Berkshire Hathaway's focus on capital-intensive but mature businesses (Holding company). Each of these firms has mastered the idea of the narrow, scalable layer.

At a high level, the rather random array of activities these companies engage in sometimes gives them the appearance of the worst of the old conglomerate models. But, what is often the case is that these ancillary businesses add volume to these firms' core profit-making activity – and these firms are indeed sharply focused on a narrow business activity. We are, of course, abstracting from the purely speculative investments some of these companies make either as potential defensive actions or as pure investments in future business models.

Chapter 4: Niche after niche - Organizers

Organizers, in general, are more simply structured than are Platform companies (of course, exceptions exist). An Organizer's primary focus is typically on defining, developing and maintaining its customer base. Accordingly, an Organizer's production stack is typically as simple and flexible as possible to allow it to shift with the needs or desires of its target customer base.

Organizers build their competitive advantage through economics of fit³, which necessitates knowing their end-markets and using this knowledge to create products or services that appeal specifically to their target customers. Additionally, their customers must value these items in excess of the cost of production because Organizer firms build their advantage by providing a better fit and charging a premium, rather than by operating at a lower cost. In the context of modern disentanglement, Organizers naturally tend to fall within the fragmenting layer given their emphasis on matching.

The EaaS 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 target community – even the narrowest niche markets – with a tailored product and quickly reach optimal scale across a global marketplace; these conditions are highly conducive to the Organizer model. At the same time, the firms that do so can realize higher margins and a higher return on equity than would have otherwise been possible had they needed to build their own supporting operating structure.

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. Exhibit 11 highlights key features of Organizers.

³ Historically, economies of fit were referred to as monopolistic competition. See Chamberlin, E. (1933), *The Theory of Monopolistic Competition*, Cambridge Harvard University Press.

Exhibit 11: Organizer companies

Source: Goldman Sachs Global Investment Research

How Organizers create sustainable competitive advantages

Organizer firms 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 it still is. In the EaaS economy, Organizers can narrowly focus on the match and freed from having to address physical production, sales, or distribution, can be both capital light and global.

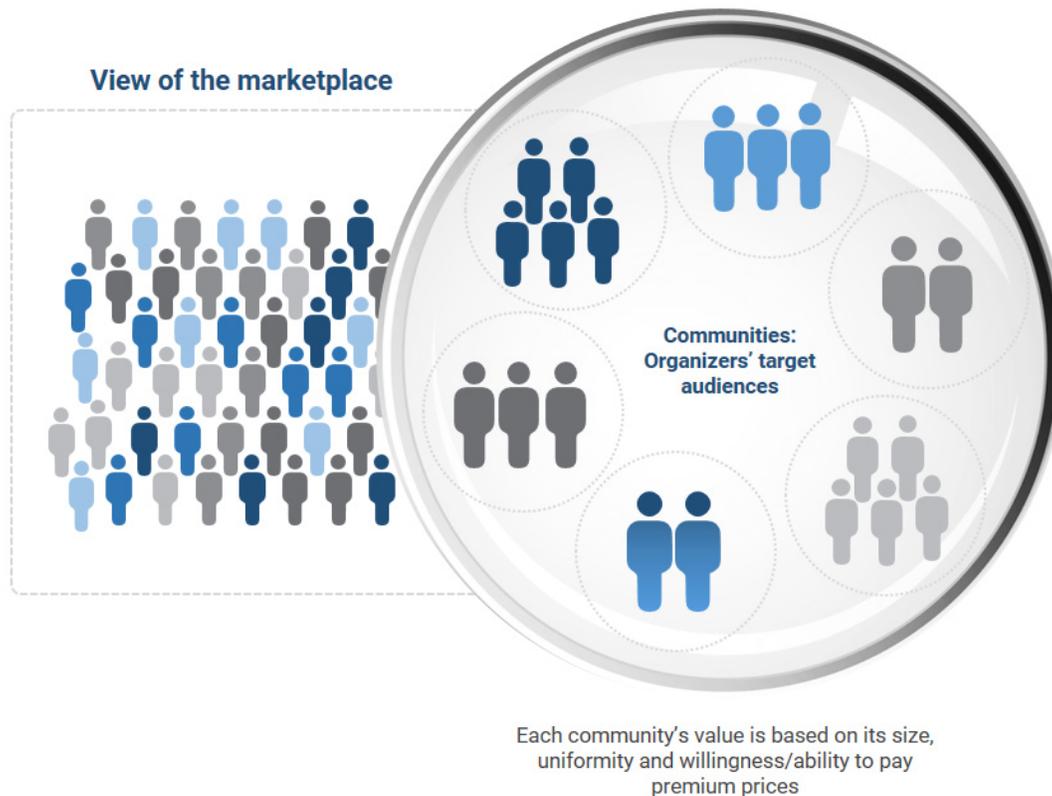
Organizers and the role of communities

Before the EaaS economy, firms' target markets were often regionally constrained; production and distribution limitations were factors, as were advertising budgets and 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.

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, Instagram, 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 can then become a defined market segment, and companies can create custom products to match that community. What’s more, in the EaaS economy, these companies need not remain limited by the natural scale of production or of distribution due to outsourcing. The ability – to find, foster and serve highly specific communities – is a defining characteristic of the EaaS economy – and of Organizer companies specifically. Exhibit 12 illustrates these concepts.

Exhibit 12: Organizers focus on serving specific communities



Source: Goldman Sachs Global Investment Research

Given the importance of communities in the EaaS economy, particularly to Organizers, 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.

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 or control these groups. The fit and loyalty that the Organizer engenders from any community must be earned and then re-earned – 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.

Case study 12: Netflix

Netflix may be the best example of community creation as a business model, but it is somewhat atypical in that it does not suffer from the cross contamination normally associated with community brand management. Netflix can provide its subscribers with kung fu movies and pacifist political fare without creating cognitive dissonance for its users. This allows the firm to pursue a particularly aggressive business strategy of identifying and serving multiple specialized communities and thus provides an excellent case study in how those niches are found or created.

Netflix was able to establish itself as an early entrant in media streaming by licensing some broadly watched mass market programs such as *Breaking Bad* and *The Office*. Netflix's current growth strategy now largely appears to be based on developing original content that appeals to specific communities of viewers. By matching its content to the many communities of viewers that it serves, Netflix is able to both build and reinforce its large and fast-growing base of paying subscribers, which is now approaching 150 million globally.

The key economic point is that Netflix – as a subscriber-based service – profits by increasing the number of its subscribers, rather than by increasing engagement with its subscribers as an advertising-driven company would.

Finding the next niche – or why TV is getting weirder

To illustrate how Netflix's library of content drives subscriber growth by adding new communities, we analyzed comments made on Reddit discussion forums since 2010 about 50 programs available on Netflix. Some of the programs that we screened for are Netflix Originals (as with *Daredevil* and *Russian Doll*), while others are content licensed from third parties (as with *Breaking Bad* from AMC and *The Office* from NBC). In total, we considered more than 500,000 comments posted by more than 100,000 Reddit users.

Exhibit 13 demonstrates the difference between how programs rank when viewed on an advertising basis (where the focus is on total viewership) relative to how they rank when viewed on a subscriber basis (where the focus is on capturing incremental viewers). The first column – shown on the left-hand side of the exhibit – ranks programs by the number of Reddit users who comment on the program. The second column –

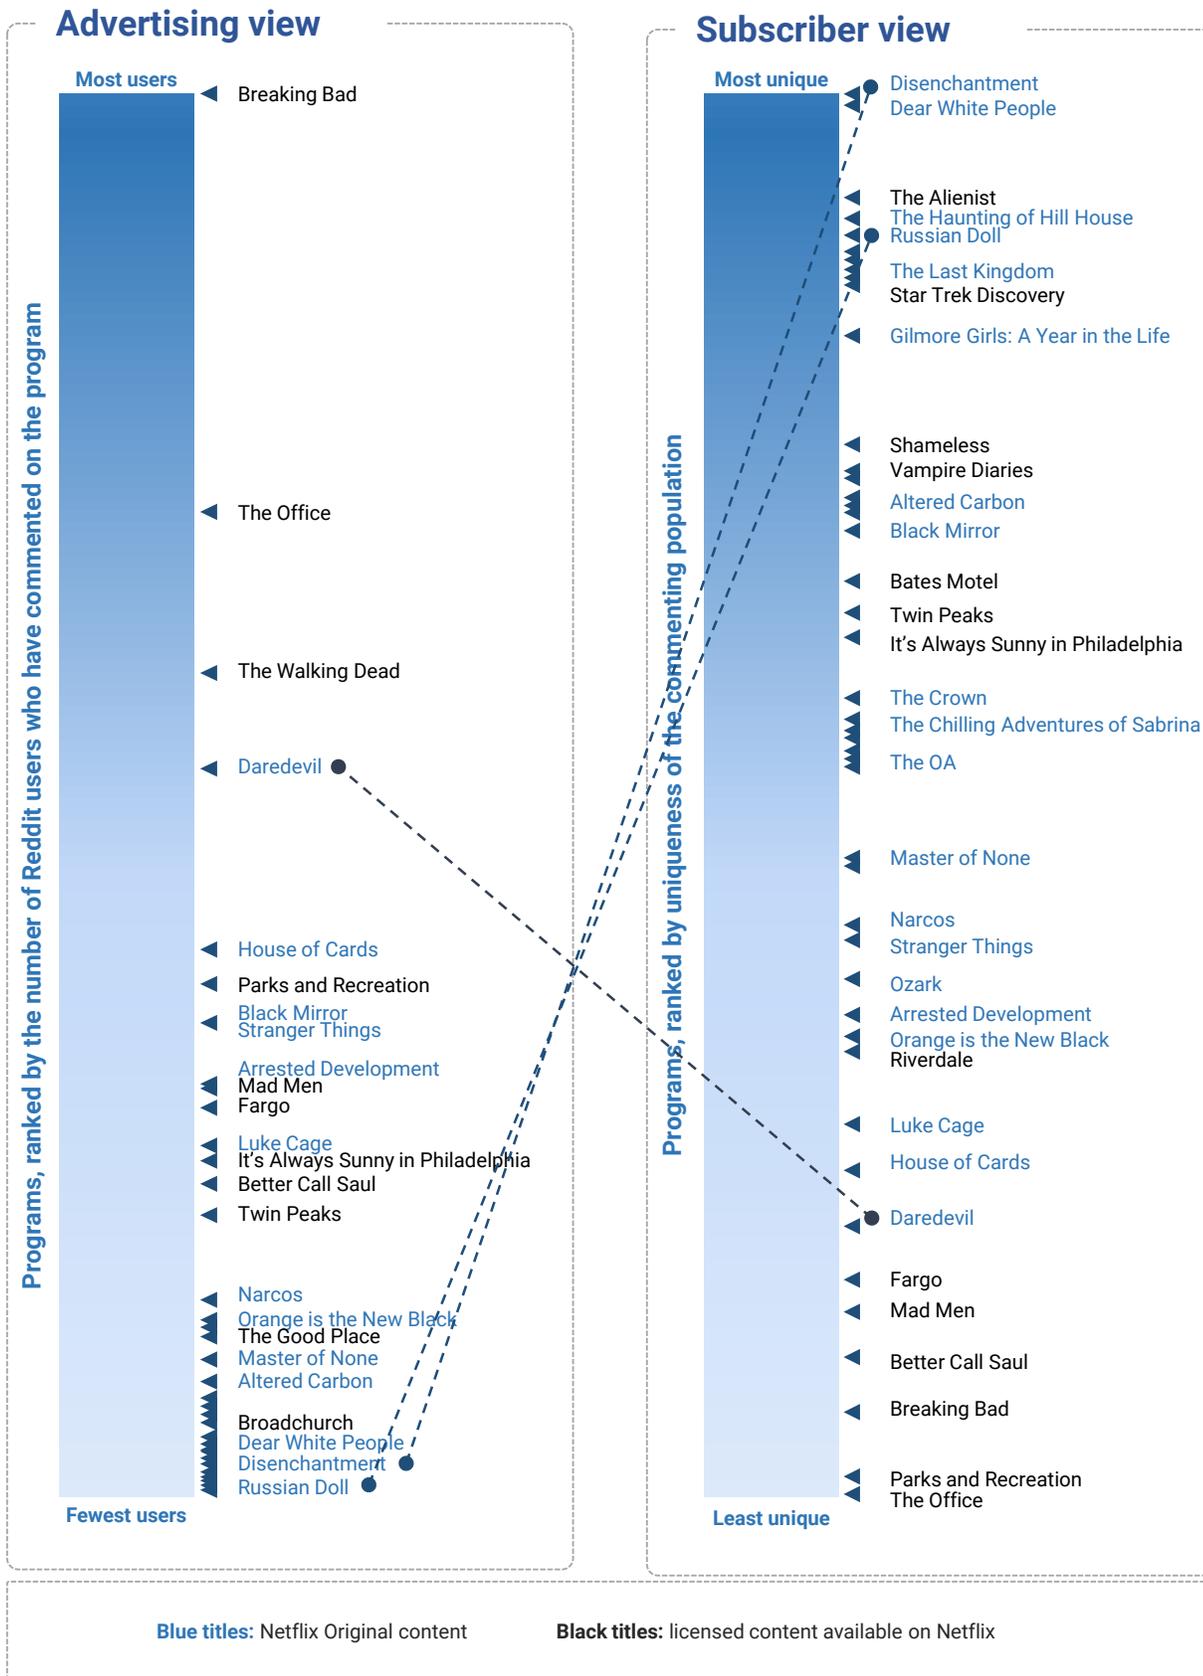
shown on the right-hand side of the exhibit – ranks programs by how “new” or incremental the commenting population is (or, said another way, by how unique the program’s community is). Rankings in the second column are determined by calculating how well a program’s community can be replicated by combining the five most similar programs already available on Netflix⁴; this can be thought of as a measure of uniqueness (see Appendix A for more detail on how this analysis was done).

The key point of Exhibit 13 is that the subscriber model has a radically different rank ordering of programs than does the advertising model. Programs like Disenchantment and Russian Doll – which very few Reddit users comment about – appear at the top of the subscriber rankings, while Daredevil – which is a more popular program – has a relatively low ranking per the subscriber model but a relatively high ranking per the advertising model. Accordingly, an advertising-driven network would much prefer to have Daredevil while a subscriber-driven firm would prefer to have Disenchantment or Russian Doll as these programs attract incremental audiences.

By examining these rankings, it is clear that Netflix-developed programs (titles are written in blue font) tend to appear at the top of the subscriber model rankings but at the bottom of the advertising rankings. One would also expect, over time, that generally popular programs, such as Daredevil, would migrate to advertising based platforms re-enforcing this high-subscriber/low advertising view pattern for Netflix original programs.

⁴ We looked at a variety of similar ranking algorithms, including best fit types, all of which produced similar answers.

Exhibit 13: A comparison of the advertising and subscriber views of programming



Source: Netflix, Crimson Hexagon, Goldman Sachs Global Investment Research

Below, Exhibit 14 offers more of a visual representation of these communities. As Graphic A in this exhibit shows, there are four broad “clouds” of users, as defined by their commenting habits. The first three clouds, which are light blue in color, capture Reddit users who commented on broadly popular programs such as *The Office* and *Breaking Bad*. The fourth cloud – which is gray in color – includes users who did not comment on broadly popular programs, but who did comment on other programs available on Netflix.

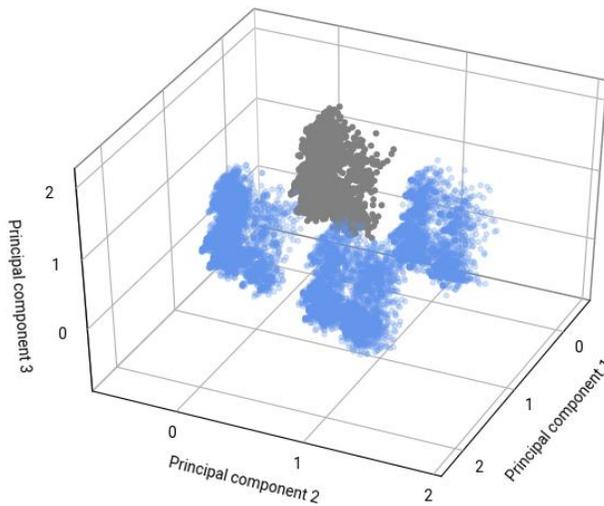
We build on this analysis in Graphic B of the same Exhibit to highlight: 1) the users who commented on *Daredevil* – shaded in dark blue, and 2) the combined users who commented on *Disenchantment* or on *Russian Doll* – shaded in dark red (the overlap in users that commented on both *Disenchantment* and on *Russian Doll* is less than 1% and thus is not visible).

Graphic B neatly demonstrates the difference between niche and popular communities. Reddit users who commented on *Daredevil* – the dark blue dots – are broadly scattered across all four clouds. In comparison, users who commented on *Disenchantment* or on *Russian Doll* – the dark red dots – are tightly bunched in the gray cloud, thereby reflecting these programs’ high efficiency in attracting incremental subscribers, but in much lower total numbers.

Graphic B also illustrates why, as programs find and attract new niches, each one is likely to be smaller and perhaps quirkier. Thus the well noted tendency for new TV programs to become more and more narrowly targeted or, as some have characterized it, weirder.

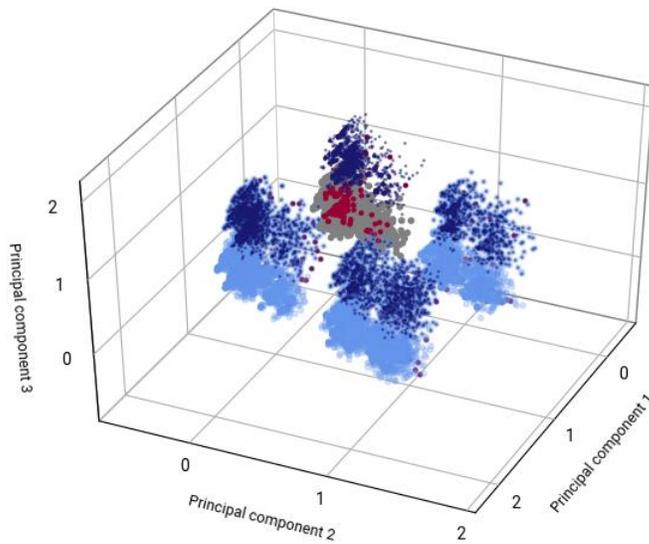
Exhibit 14: TV viewer analysis – analysis of Reddit comments

Graphic A: The fourth pillar – the untapped incremental audience



■ Reddit users who have commented on Breaking Bad and/or The Office
 ■ Reddit users who have commented on other programs

Graphic B: The unique appeal of Disenchantment or Russian Doll vs. the popular appeal of Daredevil



■ Reddit users who have commented on Disenchantment or Russian Doll
 ■ Reddit users who have commented on Daredevil

Source: Netflix, Crimson Hexagon, Goldman Sachs Global Investment Research

Community-linked Organizers

As we noted earlier, Netflix is a somewhat special example of Organizer firm in that it can collect multiple communities without creating brand confusion. The more normal case is when an Organizer company and its community are linked in some type of identity loop, where the community defines the firm's output and that output defines the community – which makes it difficult (if not impossible) for the Organizer to serve multiple communities at once.

From the Organizer's perspective, the strong link that exists between the community and the brand creates strong incentives for the firm to protect its position in its core market, and few incentives to invade others' territory. This is because expanding into new, different or potentially incompatible communities can damage the Organizer's ties to its core community. The inherent contradiction associated with "mass-market luxury," particularly at the high end of the market, offers a reasonable example. These dynamics also help to explain why brand segmentation is no longer as effective a strategy as it once was.

To that end, community identification can be a powerful positive force. It can create strong self-reinforcing patterns where firms diverge from one another in terms of their core areas of focus. In economic terms, this is what is referred to as monopolistic competition, though given how this concept has evolved in the EaaS economy we refer to it as the economics of fit. 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 13: lululemon athletica

Consider lululemon athletica as an example of a more common type of 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 EaaS 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 given its millions of online followers. Exhibit 15 illustrates lululemon's Organizer business model.

Exhibit 15: Lululemon's Organizer business model

Source: lululemon athletica, Goldman Sachs Global Investment Research

Given the firm's close relationship to its community and its promise of quality, execution missteps have proven to be costly. 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 senior leadership changes. Despite this issue, over time, as the athletic apparel market has fragmented, lululemon has benefited by catering to a narrow but profitable niche.

The lululemon example speaks to the potential pitfalls of an Organizer failing to serve its community. But the question of whether an Organizer should maintain a narrow focus on one niche community or pivot to a larger one is worth considering.

As we noted earlier, Organizers must cope with the risk that expanding into new markets – and serving multiple communities simultaneously – can ultimately damage its ties to its core community. Next we examine Coach as an example of an Organizer firm that has had difficulty serving the high and low ends of a market simultaneously, straining its community ties.

Case study 14: Coach

Founded in 1941 as a family-run business, Coach focused for decades on designing, producing and selling high-quality leather goods and fashion accessories for a narrow customer base (the firm describes itself as “the original American house of leather”).

Over time, the company expanded both its product line and its customer focus to offer affordable luxury, representing early efforts to reach a broader market. Eventually, the company began incorporating mixed materials (not just leather) and more ornate designs than it had in the past. This strategy allowed the company to expand its target market and increase the scope of its operations.

In the mid-2000s, the firm expanded its lower-priced offerings – including the range of items it sold carrying its logo – used a wider range of raw materials and grew its factory outlet presence. Although these changes allowed Coach to expand its revenues and capture a larger market for a time, these changes ultimately weighed on Coach's core

community – which favored quality craftsmanship – in favor of a larger market (for some core customers, this move cost the brand some of its cachet). Coach's dollar share of the US bags and luggage market was 15% in 2008, peaked at 17% in 2010 and fell to 13% by 2013.

In 2014 the company announced its intention to reposition itself as “modern luxury” from “accessible luxury.” The firm closed stores and adjusted its sales practices – reducing the frequency of promotions and altering its product strategy (e.g. by improving the quality of the products available through its outlet stores). The company also refocused on its own retail presence to further control discounts, however by 2016 the firm's market share had fallen to 8%.

As this example shows, as Coach expanded its target market, it faced challenges related to successfully preserving its competitive differentiation, particularly given the increased customer visibility in the EaaS economy.

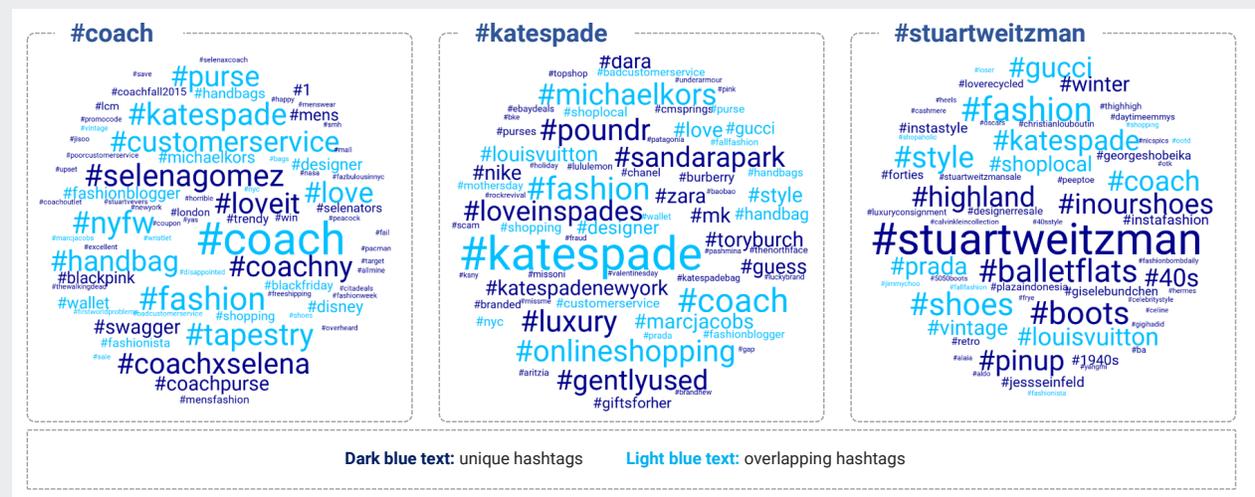
A glimpse into market perceptions of Tapestry using Tweets

As part of its repositioning efforts, Coach made two significant acquisitions – of Stuart Weitzman and Kate Spade & Company, which are also each Organizer firms but with different target communities – before shifting to a Holding company structure (with Tapestry, Inc. as the Holding company and Coach, Stuart Weitzman and Kate Spade as the portfolio brands).

In Exhibit 16 we offer a glimpse into market perceptions of Coach relative to Kate Spade and Stuart Weitzman. To do this analysis we considered Tweets referencing #coach posted since the start of 2015 and Tweets referencing #katespade or #stuartweitzman after these firms were acquired, filtered for relevance to each brand (we do not consider Tweets referencing sports coaches, for example). For further technical detail on how this analysis was done, see Appendix A.

Each word cloud is the product of analyzing millions of Twitter comments to identify affiliated topics. In this exhibit, the size of each word corresponds directly to how frequently it appeared in conjunction with the primary hashtag (meaning, relative to #coach, #katespade or #stuartweitzman). Likewise, the color of the text in each word cloud has meaning: words in dark blue are unique to the primary hashtag, while words in light blue are used in conjunction with at least one of the other hashtags.

Exhibit 16: Understanding communities through word clouds: #coach, #katespade, #stuartweitzman



Source: Crimson Hexagon, Goldman Sachs Global Investment Research

As the exhibit shows, while there is some overlap between the language used in conjunction with all three brands (reflected by the light blue text), there is more overlap between Coach and Kate Spade than there is with Stuart Weitzman. Nevertheless, the nature of the chatter around each brand varies meaningfully – related to the marketing for each brand (e.g. #katespade and #sandarapark or #coach and #selenagomez) or to their primary product focus (e.g. #stuartweitzman and #boots).

What's more, based on the other brands that appear in each word cloud, this analysis would suggest that #coach may still be perceived as a more accessible and affordable brand, while #katespade and #stuartweitzman may be perceived as being somewhat higher-end brands.

Chapter 5: The competitive value of data

Data is now the lifeblood of many firms, particularly in the modern economy in which companies tend to focus on their narrow area of expertise while outsourcing the rest. From organizing and optimizing complex multi-vendor production processes to customer acquisition, service and retention – these modern firms are almost entirely dependent on data. Naturally, trying to use data to establish a competitive edge has therefore become big business.

Anecdotes about data-driven successes abound, but experience suggests that it is actually quite difficult for businesses to use data to build a sustainable competitive advantage. In fact, pinpointing examples of companies that have successfully used data to maintain a competitive edge is a challenging task. This begs the following two questions: 1) why haven't more companies been able to build a sustainable competitive edge using data, and 2) when can data serve this purpose?

We address these two questions by building a conceptual framework that we refer to as the “learning curve.” The learning curve helps us assess the factors that underpin when a company can use data to create a competitive edge – and perhaps more importantly, when it cannot.

Using the learning curve, we analyze four types of data-driven learning strategies:

- **Data-smart strategies** rely on a business's internally generated data as the foundation for producing data-based insights – or what can be thought of as learning. These insights can be used to optimize both a firm's operations as well as its output. An example of a business that uses a data-smart strategy is Amazon's logistics service.
- **Data-asset strategies** are dependent on a business's ability to build a proprietary dataset using secondary sources, for example by collecting (free or purchased) data from sensors, genetic labs or satellites. These proprietary datasets can be used to produce data services that are sold to others. An example of a business that uses a data-asset strategy is IBM Watson Health.
- **Data-feedback strategies** are applicable to businesses that collect user data. To employ this strategy, businesses collect the data that is generated by the users of their products or services, analyze it and leverage the resulting insights to enhance their products or services. Said another way, data-feedback strategies describe when a company leverages user data to create a feedback loop between its users and the goods or services it provides to those users. Examples of businesses that use data-feedback strategies include Spotify with its playlist suggestions, Amazon with its retail product recommendations or Google Maps.
- **Network strategies** are also applicable to businesses that collect user data. However, the purpose of a network strategy is to leverage user data to connect users with one another. Examples of businesses that use network strategies include Uber, Lyft, Airbnb and Facebook.

While the economic models that underlie each type of learning strategy are unique, each one requires data accumulation to drive learning, which then serves as the primary source of potential competitive advantage. We also analyze the role of data decay and copy risk in determining the competitive value of a data-based advantage.

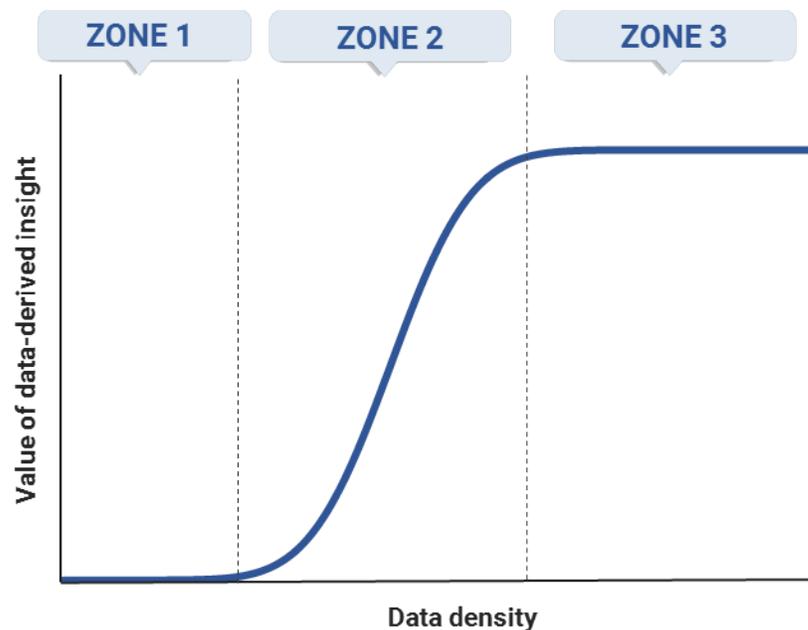
The learning curve

Exhibit 17 is an illustrative depiction of a learning curve, which can be used to assess the scale-based economics of learning as well as the potential competitive impact. More specifically, and as the exhibit shows, the learning curve is a depiction of the potential value of data (PVD) – or the total value of what can be learned from data – as a function of the amount of usable data a business possesses.

Each unit on the y-axis represents the incremental value derived from analyzing data related to a specific question. Each unit on the x-axis represents the density (or volume) of usable data, which is dependent on the rate of data collection as well as the rate of data decay.

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

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



Source: Goldman Sachs Global Investment Research

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

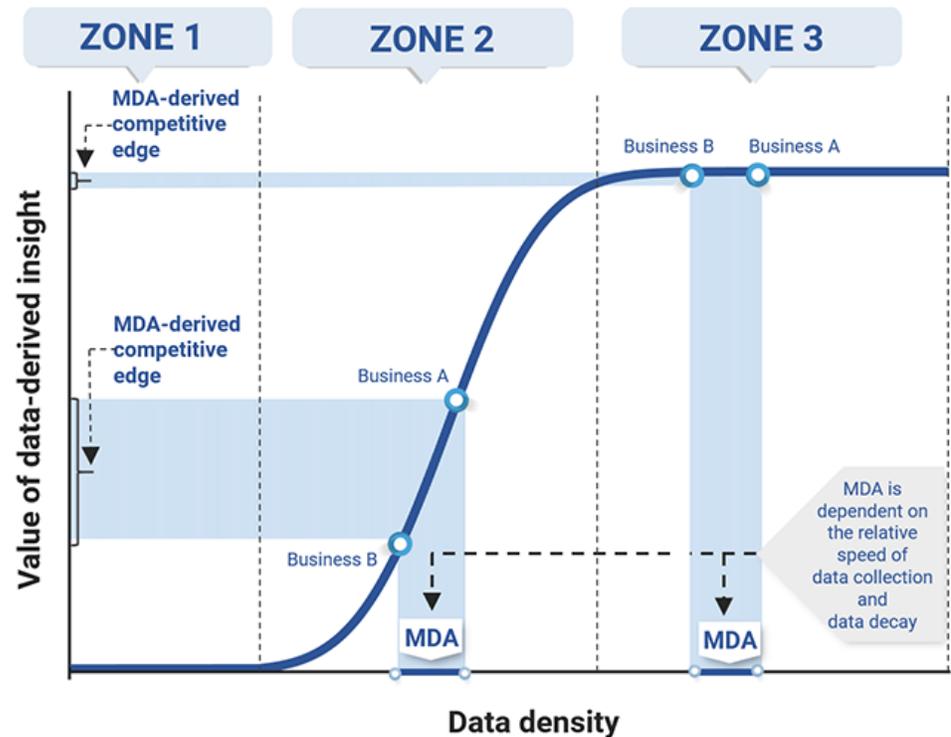
Thus, for each type of learning strategy, the uncertainty related to the PVD is a central question. The PVD must be large enough to justify the expense of building, buying or

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 issue for all learning strategies is related to data scarcity. On the one hand, if there isn't enough data available to analyze, data-based strategies are likely to get trapped in zone 1, producing little value. On the other hand, if the data isn't scarce in some way, all market participants will likely reach 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. Exhibit 18 illustrates these dynamics.

- In zone 1, depicted on the left-hand side of Exhibit 18, the learning curve is flat and the incremental value associated with data analysis is low. This means the gains associated with incremental data analysis and data density are limited. The fact that learning is slow in zone 1 is due to the fact that a certain amount of data must be collected before it can be effectively modeled.
- In zone 2, the learning 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 (MDA) and can generate significant incremental value (as the middle portion of Exhibit 18 shows). The MDA refers to the pure advantage in the amount of data one business can collect relative to another; the learning curve can then be used to map that MDA to determine a business's relative competitive position given the value of its data-derived insights.
- In zone 3, the learning curve flattens because incremental data accumulation and analysis no longer result in significant value, which can be seen on the right-hand side of Exhibit 18. 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; businesses in the same market segment that reach zone 3 are in essentially the same competitive position. Accordingly, as Exhibit 18 shows, the same MDA that resulted in a significant competitive advantage for Business A relative to Business B in zone 2 becomes a very small advantage if both businesses reach zone 3.

Exhibit 18: The value of a maintainable data advantage (MDA) differs significantly by zone

Source: Goldman Sachs Global Investment Research

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.

More broadly, the learning curve – with its characteristic S-curve shape – can be derived in a number of ways, but the derivation that is easiest to understand (both in terms of the underlying mathematics and the economics) comes from network theory. As a network is built by connecting nodes (at random in the simplest derivations), initially the connections only link two isolated nodes. This creates some incremental value per connection added, but not a tremendous amount (zone 1).

As more nodes are connected, a state is reached in which many of the nodes are already connected to other nodes. As a result, instead of linking two isolated nodes, new connections usually link clusters of already connected nodes. This means that incremental connections increase the average size of a network cluster much faster and thus the value of the network increases more rapidly as the number of connections increases (zone 2).

Once this stage of connecting clusters begins, the network quickly becomes one large cluster plus some isolated nodes and small clusters that remain unconnected to the big cluster. At that point, additional connections can no longer create much value, as they can only add on small clusters to the big cluster and often only connect nodes that are already connected through other paths (zone 3).

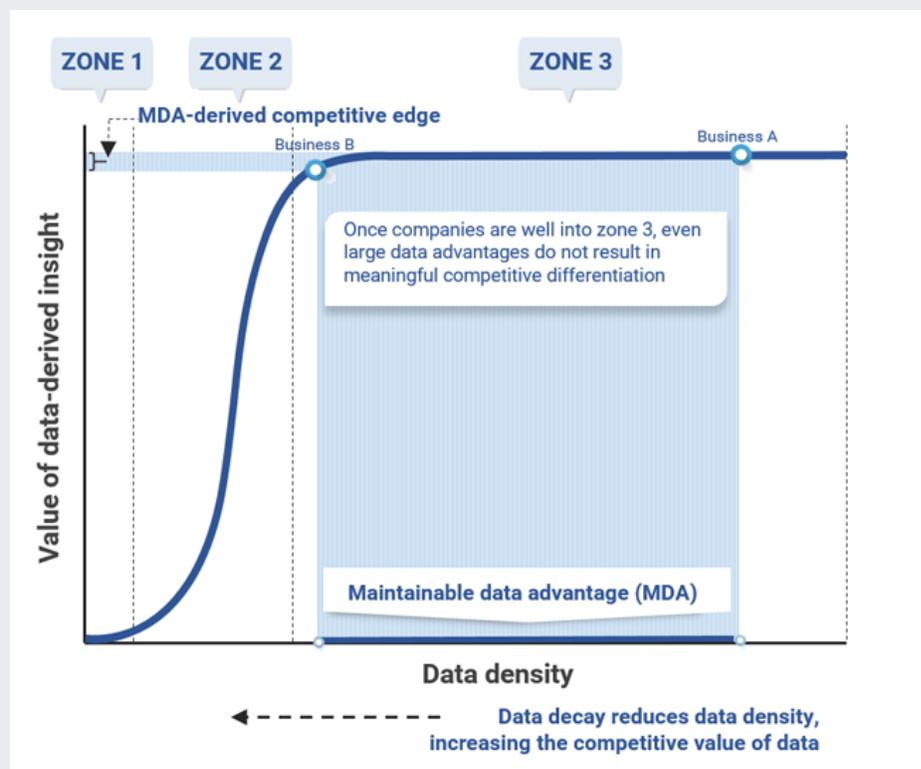
Data density: the rate of collection vs. the rate of decay

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.

As Exhibit 19 shows, in zone 3 – regardless of the size – a maintainable advantage in data density (which we refer to as maintainable data advantage or MDA) has little value for a business since there is typically little competitive differentiation. In zone 2, however, even a small MDA can be very valuable and highly differentiating.

Depending on where a business is positioned along the learning curve, the value of data-derived insight changes and affects whether it is possible to turn a data advantage into a competitive one. To that end, the rate of data decay is critical. 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. If the rate of data decay is high, however, then the business is likely to be limited from progressing past zone 2 or zone 1. The key is that in zone 2 it becomes possible for the business to establish a competitive edge if it can maintain an advantage in data density.

Exhibit 19: In zone 3, there is little competitive differentiation between businesses regardless of the maintainable data advantage (MDA)



Source: Goldman Sachs Global Investment Research

As an example of how data decay works in practice, consider it in the context of navigational maps. Depending on the precise nature of a map's usage, the user's sensitivity to accuracy and to how recently the data underpinning it was collected changes – thus altering 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 on the learning curve (where little or no competitive advantage could be derived from differences in accuracy or how recently the data was collected). However, if we consider the use of maps in the context of a more demanding problem – 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 accumulated data is subject to very high rates of decay, such that reaching zone 3 in the learning curve is very 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 users' 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, which allows 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's service lowers competitors' rates of data collection, reducing the value of their data-derived insights (trapping competitors in zone 1). This further reinforces 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 the searches themselves – particularly recent searches with a short-lived relevancy – became more important to producing the best (most relevant) search results, a clear self-reinforcing dynamic took hold. This is especially true in the case of popular or trend-based searches.

As a result, Google – which pioneered the use of its repository of past searches to improve the applicability of its real-time search results – has been able to translate a data collection advantage into a competitive one 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 large volumes of data with high rates of decay.

As we previously noted, unless the rate of data decay is high for a given problem, all businesses 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 – an MDA – can become a self-sustaining and self-reinforcing advantage. Data-density advantages only translate into competitive advantage in zone 2 and can be maintained only 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 uses of data-smart and data-asset strategies than of successful uses of

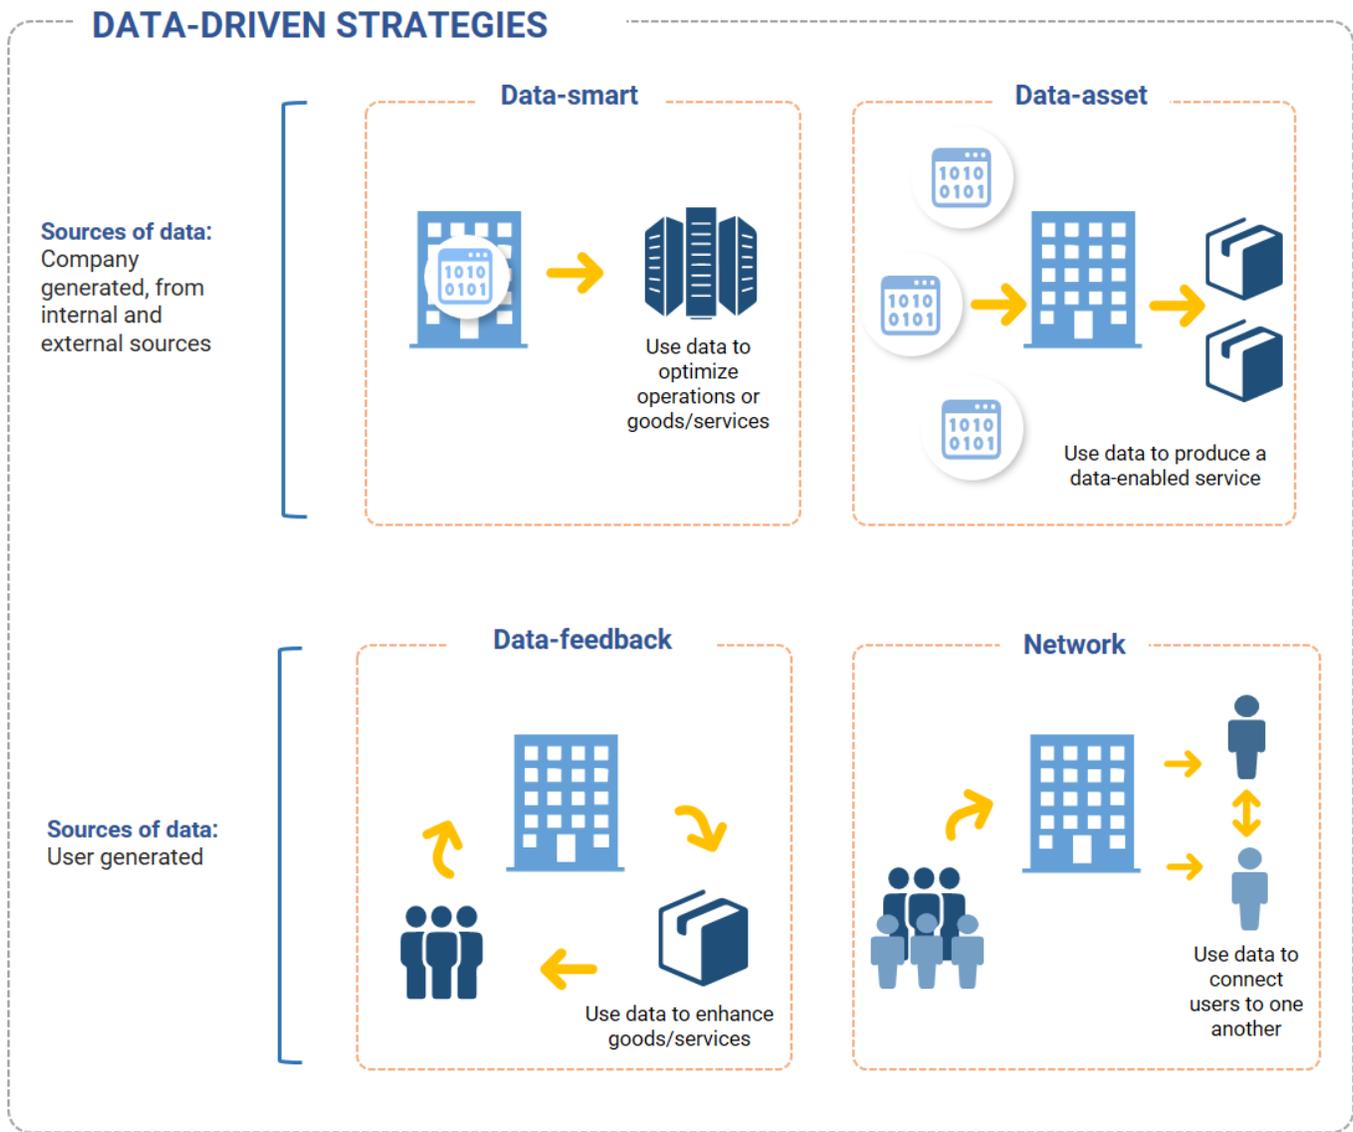
data-feedback or network strategies. In the latter instances, it just isn't that easy to find examples where the runner-up doesn't eventually make it to zone 3.

Data-driven learning strategies

As we noted at the outset of this report, we use the learning curve to analyze four types of data-driven learning strategies: data-smart, data-asset, data-feedback and network.

Exhibit 20 illustrates how each strategy works, both in terms of how data is sourced and used in each case.

Exhibit 20: The four types of data-driven learning strategies



Source: Goldman Sachs Global Investment Research

The data-smart strategy

It's a popular refrain to believe that all companies should become data-smart – since collecting and analyzing one's own operational data may seem like low-hanging fruit. In practice, however, it may not be relevant or possible to pursue this strategy.

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 business from zone 1 to zone 2. Thus, in many cases, because the quantity of data is simply insufficient, being data-smart simply proves impossible.

But, when an individual company is able to generate enough operational data to successfully reach zone 2 or even zone 3, it is likely that the information that is collected is related to highly repetitive tasks (related to logistics, simple customer support or other basic operations, as examples). It is worth noting that 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 unfavorable; said another way, it is worth considering whether the PVD is sufficiently high relative to the investment involved.

Rather than trying to use this kind of 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. In this case, the first analysis can be run using the high initial volume of data, which can be complemented over time through high ongoing usage that allows even small efficiency improvements to accumulate with meaningful results.

Taken together, these factors suggest that feasible data-smart strategies are often likely to culminate in zone 3, which means they will generally be defensive in nature – they may be a cost of entry, for example. What's more, businesses that fail to realize the basic efficiencies associated with a data-smart strategy are likely to be at a significant competitive disadvantage relative to those businesses that have realized those efficiencies.

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

The data-asset strategy

Unlike data-smart strategies, data-asset strategies are predicated on a business's ability to build robust proprietary databases – often using secondary sources of data – that allow it to produce data-driven services.

Constructing such 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 business 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.

As we touched on earlier, second movers do not face the same level of uncertainty that first movers do, which means their investments are subject to a more favorable risk-to-reward trade off. Second movers not only know that valuable data-based insights do exist, but 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.

Copy risk can be difficult to determine, particularly before a company knows how much value a particular data-asset strategy will generate to address a specific problem or question. This only reinforces the need for businesses pursuing data-asset strategies to diversify and to have sufficient capital to experiment again.

Broadly speaking, however, as businesses decide which strategies to invest in, 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-to-reward perspective. Second, 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 made clearer through examples. Consider a "safe" example first, meaning a case involving low copy risk. As in the case of IBM's Watson Health business, interpreting magnetic resonance imaging (MRI) data requires a large initial 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.

Next, consider a less safe example involving a data-asset enabled maintenance service for elevators. This service is driven by data collected from sensor arrays or from past instances of elevator maintenance. On the one hand, if producing the maintenance service requires a complex assessment of the sensor input data, copying the strategy could be difficult – in which case the originating firm would be able to maintain a competitive advantage. On the other hand, if the maintenance service could be approximated through simpler rules, for example by counting hours of service rather than calendar time associated with the service, then the service could be copied at a lower cost – and the associated copy risk would be high.

Economies of scope can easily play a significant role in driving data-asset efficiencies. 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 business pursuing a data-asset strategy to become even better at both assessing the risks and lowering the cost associated with new ventures related to their particular area of expertise.

Data-asset strategies: a comparison to deep-water drilling for oil or new drug development

The risk-to-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 there is significant uncertainty associated with the potential 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-to-reward trade off.

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

There are also 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 (as discussed earlier), since potential competitors (for example, new entrants) face very different incentives and hurdles than the innovators that preceded them.

The data-feedback strategy

The most complex – but perhaps most talked about – data strategy is one that relies on the collection of user data to refine the user experience, which we refer to as the data-feedback strategy. There are two key challenges related to pursuing a data-feedback strategy: the first is determining whether an advantage actually exists, and the second is determining whether it is possible to maintain that advantage.

Historical efforts suggest that finding a true advantage based on customer data isn't easy. Individuals' "discovered behavioral patterns" generally aren't complex or surprising. Amazon offers an illustrative example. The firm's use of its own 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 its product placement and sales strategies are often quite simple, to the point where third-party retailers can mimic Amazon's strategies and now outpace Amazon in terms of unit sales on Amazon's own retail Platform.

To make this example more granular, consider that a business 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. This is another kind of copy risk.

Accordingly, when an advantage can be found (when the PVD is high) the data-accumulation process must be sufficiently difficult that the business pursuing the data-feedback strategy is able to progress up the learning curve (and capture a significant portion of that PVD). As well, competitors must be constrained from doing the same.

The network strategy

Network strategies are similar to data-feedback strategies in that they also leverage user data in ways that reinforce the value of their products or services. The primary difference is that network strategies use this data to connect users to each other (while data-feedback strategies leverage user data to enhance the output that is provided to each user).

For businesses that use network strategies, data density is defined by the number of active users. The key driver of data decay typically has more to do with activity levels than a change in the data – as is the case for the other types of strategies we have identified.

The competitive issues associated with network strategies are similar to those associated with data-feedback strategies. For both types of strategies, progressing out of zone 1 involves significant hurdles related to amassing sufficient user data. After doing so, the business's ability to build a sustainable competitive advantage is dependent on whether competitors are also able to reach zone 3 – where there is little to no differentiation.

Businesses pursuing network strategies must consider: 1) how an active user in the space is defined, and 2) whether being an active user in one network precludes the user from being active in others. If businesses are competing for users' time (as with Netflix or Instagram), 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 in terms of car service, or Airbnb and VRBO for temporary lodging), the market is more likely to have multiple vendors that are in ongoing competition, and the network alone is unlikely to result in a persistent advantage.

Communities of users – and the relevant boundaries – play an important role in driving the economics of network strategies. In some circumstances, networks naturally divide into communities in which there is an advantage in specializing in providing network services to a specific community rather than to the general population. Modern dating applications – like Bumble and e-Harmony – are examples of businesses leveraging network strategies and where success is determined by active users.

For businesses using network strategies, the ability to monitor and regulate membership can become a sustainable advantage. As examples, the ability to offer high-quality drivers, dynamic rental spaces, verified vendors, or other specific community affiliations may represent key competitive strengths. In such cases, the

business may be mixing two types of learning strategies: a network strategy (based on the directory of users) with a data-asset strategy (using reviews, background checks) to police the service. The mix can result in a hard-to-replicate business model.

The four-part test

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

1. Is there sufficient data to analyze?
2. Are the insights gained from the data novel and valuable enough to be of competitive benefit?
3. Is the data-derived strategy difficult to copy without the data?
4. Is the data sufficiently scarce or hard enough to collect that competitors cannot replicate the analysis in normal course?

If each of these questions elicits an affirmative response, building a sustainable competitive edge through data is possible. More often than not, however, this is unlikely to be the case. As a result, data often serves as a cost of entry to a given market and copying a data-based strategy – rather than leading one – is likely to be more efficient over the long term.

Chapter 6: Concluding thoughts

We titled this collection of essays “A survivor’s guide to disruption” because understanding the roots and nature of disentanglement is key to adapting to the pace of disruptive change. From companies’ perspectives, disruption is perceived as a threat. But the shift to the EaaS economy gives firms the opportunity to respond to this threat by organizing themselves around their sources of competitive advantage to establish a sustainable competitive edge.

These dynamics are the result of a natural progression over years of companies reengineering their operations in a self-reinforcing cycle of disentanglement and standardization that only accelerates the pace of change. We expect this process to continue, enabled by the introduction of new technologies, until the most optimal and efficient outcomes are achieved. In other words, continued disruption is highly likely.

The new business environment is in many ways easier for firms to operate in but also harder for them to succeed in. And this easier/harder paradigm can be used to summarize much of what is confusing about the new economy:

- Innovations are in general easier to execute, but are less likely to provide a sustainable advantage.
- The use of external service providers allows firms to have a more narrow, focused approach to business but also creates a web of interdependencies that needs to be managed.
- The path to success is actually more straightforward today: firms must have the best product, but the ability to stay on that path is far harder because someone else is coming!

An economic policy postscript

Thus far, we have focused on the specific implications of the EaaS economy for individual businesses, but there are broader implications that are worth considering. We only outline them here to help engender discussion.

First, jobs and capital are likely to accumulate in different places in the economy. Effective policy and economic analysis is going to have to learn to separate what is happening in the consolidating layers from the fragmenting layers because the trends underlying these layers are likely to be quite different.

For example, jobs are more likely to expand in the fragmenting layers – think Uber drivers (not Uber), while capital accumulates in the consolidating layer – think Amazon. Education and skills-based solutions to improving job creation, while important in the consolidating layers, are likely to be numerically small as skills leverage (the ability to hire and borrow expertise) is increased. In contrast, job creation and wage growth in the fragmenting layers, where the numbers are likely to be larger, are likely to be associated with business formation.

Economic statistics also face challenges. While the efficiencies in the consolidating layers follow the normal notions of productivity, investment in the fragmenting layers is primarily focused on creating consumer surplus, which is often ignored by standard economic statistics. In particular, the kind of micro targeting discussed in the Netflix example is productivity destroying because it increases the average cost of production without accounting for the intensity of “joy” that consumers experience from highly targeted programming.

The new EaaS economy also poses challenges for regulators. The inability to define a modern company by its product raises the need to regulate activities rather than product markets. This would better match the modern structure of the EaaS economy and would allow the broader business environment to continue to re-optimize for economic gains rather than to arbitrage regulations.

The EaaS economy also requires a rethinking of competition policy. This is because most classic anti-competitive concerns are reduced in the new business ecosystem, for three key reasons:

1. 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.
2. 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.
3. The plug-and-play aspect of the EaaS economy lowers switching costs and makes it easier to displace “bad” actors, further limiting the scope for anti-competitive activity.

However, the new Platforms may end up with classic monopoly-type power in activities the way they once did in products. But without clear market definitions, regulators lose the signals they have historically had from the associated concentration ratios.

Also, EaaS may create competition problems in some small markets. As with narrowly used drugs or patented technologies, as examples, it has become easier to set up capital-light, narrow-purpose entities that can exploit existing local markets. This could be called “niche exploitation.”

Consider 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 EaaS economy. Accordingly, this type of company can exploit the fact that it has only narrow market power, which is usually not a target of antitrust 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 EaaS economy, smaller businesses and smaller markets may have the largest inherent antitrust risk as opposed to larger businesses and larger markets.

While the EaaS economy certainly creates significant stresses and strains within the system, not the least of which is greater economic insecurity, one has to believe that faster, easier innovation is – in the end – net positive.

Appendix A: Considering communities

Principal component analysis of Reddit comments on 50 Netflix programs

We leverage public posts on Reddit's discussion forum to gain insight into Netflix's community clustering. First, we identify a universe of 50 television series available to be streamed through Netflix. While this universe is not guaranteed to capture the 50 most viewed/commented on/liked programs, we consider the sample to be a broadly representative cross-section of series that includes both Netflix Original content as well as externally licensed content. For each program that we considered, we required there to have been at least 150 users who have mentioned the program and for the program to have an unambiguous name⁵.

We consider all Reddit comments from January 1, 2010 to May 8, 2019 within the "Television" subReddit, which we use to further reduce spurious program mentions and unrelated comments. Similarly, we require that Reddit users mention fewer than 20 programs overall as this removes accounts programmatically posting updates on platform content, such as automatically generated lists and updated content postings (e.g. "bot accounts"), with a negligible effect on the number of users considered (0.3% of users removed). To create our data set we aggregate all program mentions by unique user and assign each program a score of one if the user mentions it, otherwise the score is set to zero.

We use a principal component analysis (PCA) framework to quantify patterns within users' interests across programs, with observations and variables represented by users and television programs, respectively. We employ PCA to provide an orthogonal representation of the entire data set, where the resulting components are mutually orthogonal and represent linear combinations of all programs. Components are ordered by the amount of total variance captured from the data set, where the first component has the largest contribution.

Because the data (by construction) is comparable across variables, we do not rescale the data sample as the weights incorporate not only similar user interest in programs but also the frequency with which programs are mentioned. Thus Breaking Bad, which is mentioned by the most Reddit users, would be expected to have a large weight in the first component due to the nature of the data and the construction of the space.

To represent users in a three-dimensional space (e.g. the first three principal components), we take the inner product of their associated observation with the weight of each respective component. To view specific cross-sections of users based on program interest, we filter the input data to include only users who have commented on

⁵ The complete list of programs considered: 13 Reasons Why, The Alienist, Altered Carbon, Always Sunny in Philadelphia, American Horror Story, American Vandal, Arrested Development, Ash vs Evil Dead, Bates Motel, Better Call Saul, Black Mirror, Bojack Horseman, Breaking Bad, Broadchurch, Crazy Ex-girlfriend, Daredevil, Dear White People, Disenchantment, Fargo, Gilmore Girls, House of Cards, Luke Cage, Mad Men, Master of None, Mindhunter, Narcos, Orange is the New Black, Orphan Black, Ozark, Parks and Recreation, Parts Unknown, Peaky Blinders, Portlandia, Riverdale, Russian Doll, Sabrina, Shameless, Star Trek Discovery, Stranger Things, The Crown, The Good Place, The Haunting of Hill House, The IT Crowd, The Last Kingdom, The OA, The Office, Twin Peaks, Unbreakable Kimmy, Vampire Diaries, The Walking Dead.

the program(s) of interest. In the case of multiple programs, we employ a logical ORing, e.g. for “users who comment on Russian Doll or Disenchantment” we include users who have commented on only Russian Doll, on only Disenchantment or on both programs.

Finally, we use the following technique to create a metric approximating a program’s uniqueness (meaning, to determine a value that is representative of a niche and narrow audience). We determine the pairwise correlation for each of the 50 selected programs using the previously defined user interest data set. For each program, we use the mean of the highest five correlation values as the score. Scores close to zero represent high uniqueness, whereas programs with similar user trends will have higher values. To validate the stability of the results, we vary the number of values considered in order to determine the mean correlation (baseline of five programs) and find that the scores are largely independent of group size.

The word clouds: a glimpse into community perceptions using Twitter hashtags

We leverage Twitter posts to quantify the similarities across tweets referencing #coach, #katespade and #stuartweitzman. We select all English-language tweets posted between January 1, 2015 and May 22, 2019 that include the hashtag representing the respective entity, e.g. for coach we included all tweets that have the term “#coach”.

To remove spam and similarly spurious tweets, posts must not be “re-tweets” (re-forwarded posts) or include links to websites. For both #katespade and #coach we apply subfilters specific to each to ensure posts are relevant to the brand (e.g. we remove tweets corresponding to a sporting team’s coach). Additionally, for #katespade we restrict the earliest post-date to be August 1, 2018 given significant non-brand related discussions in prior periods.

To compare similarities across brands’ posts, we extract all hashtags in selected posts along with the number of times each hashtag was mentioned. For the #coach, #katespade and #stuartweitzman screen we assign a light blue color to all hashtags appearing in at least two brands’ tweets. In contrast, language that is unique to each hashtag appears in dark blue. The word clouds’ hashtag sizing is related to the frequency of mentions for each brand and each brand’s word cloud is volume normalized such that a uniform font size can be applied across all images.

Bibliography

- Bezos, J. P. (Apr 2018). *2017 Letter to Shareholders*. Retrieved from <https://blog.aboutamazon.com/company-news/2017-letter-to-shareholders>
- Bezos, J. P. (Apr 2019). *2018 Letter to Shareholders*. Retrieved from <https://blog.aboutamazon.com/company-news/2018-letter-to-shareholders>
- Buffett, W. E. (Feb 2018). *2017 Letter to Shareholders*. Retrieved from <http://www.berkshirehathaway.com/letters/2017ltr.pdf>
- Buffett, W. E. (Feb 2019). *2018 Letter to Shareholders*. Retrieved from <http://www.berkshirehathaway.com/letters/2018ltr.pdf>
- Burgstaller, S. (May 2017). *Rethinking Mobility: The 'pay-as-you-go' car: Ride hailing just the start*. 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.
- Love, J. F. (Aug 1995). *McDonald's Behind the Arches*. Bantam Books.
- Netflix Case Study*. (2016). Retrieved Jun 20, 2019, from Amazon Web Services: <https://aws.amazon.com/solutions/case-studies/netflix/>
- Strongin, Steve et al (Apr 2019). *What the market pays for*. Goldman Sachs Global Markets Institute.
- Strongin, Steve et al (Dec 2018). *The Everything-as-a-Service economy*. Goldman Sachs Global Markets Institute.
- 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*.

Disclosure Appendix

Disclosures

This report has been prepared by the Global Markets Institute, the research think tank within 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.

Third party brands used in this presentation 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.

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.

© 2019 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.