

2022

11 laws driving success in tech:

Amazon's 2-pizza rule,
the 80/20 principle, & more



Where is all this data from?

The CB Insights platform has the underlying data included in this report

[Click here to sign up for free](#)



Table of Contents

1. Moore's Law	5
2. Metcalfe's Law	12
3. Law of Mobility	18
4. Gall's Law	25
5. Law of Modularity	32
6. The 2-Pizza Rule	38
7. Conway's Law	44
8. Yule's Law of Complementarity	49
9. The Law of Shitty Clickthroughs	53
10. Zimmermann's Law	59
11. Pareto Principle	64

What separates success from failure? These 11 laws contain some of the most influential ideas that the biggest tech companies use to run their operations, design business models, and build products.

Since the early days of computers, the rises and falls of tech companies have inspired countless theories about what drives success — and predicts failure — in the fast-moving world of startups. While many such theories fall flat, a few have become well-regarded descriptions of how the tech business works.

Some, like Moore's Law, have been extremely prescient. Others, like Conway's Law, provide counterintuitive insights — such as why Microsoft sells Xbox consoles at a loss, or how Facebook became one of the most valuable companies in the world by offering a free service.

In this report, we dive into 11 of the most famous tech laws, with examples from tech leaders like Amazon, Apple, IBM, and Salesforce.

1 **Moore's Law: The self-fulfilling prophecy that ushered in the digital age**

In 1965, Intel co-founder Gordon Moore made a prediction.

Moore had observed that, every year, semiconductor manufacturers were able to double the number of transistors on a single microchip and halve the cost of building them.

When he was asked by Electronics magazine to contribute his perspective on the future of semiconductor manufacturing, Moore said this trend would continue. And thanks in large part to advances in transistor miniaturization – most notably, the advent of the MOSFET transistor type – Moore's prediction panned out.

Today, we call his observation Moore's Law.

HOW MOORE'S LAW PREDICTED CHIP PRODUCTION AT INTEL FOR DECADES

Over the decades following Moore's prediction, the number of transistors on the average Intel computer chip produced grew exponentially, from just over 2,000 in 1971 to 50B in 2021.

While this progress relied on advances in technology and manufacturing to make ever-smaller transistors, it also became a kind of self-fulfilling prophecy in the industry.

For decades, the heads of rival semiconductor companies used Moore's Law to construct their annual production goals, according to Ethan Mollick of the MIT Sloan School of Management.

This approach pushed them to increase the amount of transistors on their chips – mainly out of the fear of being left behind by companies like Intel.

“If you go look at Moore’s Law, Moore’s Law was never one thing. There was transistor architecture, strain, materials, 3D architecture... Every single component of shrinking the transistor changed year over year. [...] Imagine a business where a million people wake up every day, working on making things smaller and better. It was that collective belief system about this [that helped] the inventions keep coming.”

– Jim Keller, senior vice president of Intel's Silicon Engineering Group

CHIPMAKERS ARE BLITZING PAST MOORE'S LAW TO POWER AI

Artificial intelligence (AI) applications and internet of things (IoT) devices require ultra-fast processing speed, and manufacturers are outpacing Moore’s Law to develop suitable chips.

Leaders like Nvidia and AMD have come out with chip processors capable of supporting massive neural network operations and cloud workloads. Intel itself is developing a 100B-transistor AI chip in a bid to challenge Nvidia’s dominance in the sector.

Meanwhile, new competitors have emerged to challenge Intel and incumbents. In 2018, Amazon revealed that it had been stealthily developing its own chips for use in its data centers worldwide. Its latest chip, Graviton3, has 55B transistors – 25B more than its predecessor – and is designed to increase the speed of machine learning, scientific, and cryptographic workloads.

Meanwhile, [Cerebras](#) – one of the more [well-funded AI chip startups](#) – has developed the Wafer Scale Engine 2, which it claims to be “the largest chip ever built” with 2.6T transistors and 850,000 AI cores. Each chip costs around \$2M to purchase.

As AI applications and IoT systems become more ubiquitous, the [processing power of specialized chips will grow](#), too. Google designed its Tensor chip to run machine learning models for tasks like speech recognition and face detection on Pixel phones while consuming less power than the Qualcomm chips it previously used. ARM claims that its Cortex-M55, a chip built for processing AI tasks on the edge, delivers up to 15x better performance in machine learning than preceding chips. With more power, the M55 will be able to run AI tasks like object recognition on the edge – meaning data does not need to be sent to the cloud and back to be processed.

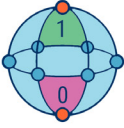
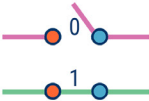
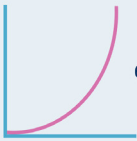

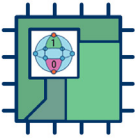



QUANTUM COMPUTING AND THE NEW MOORE'S LAW

If Moore’s law doesn’t hold as much water as it did in past decades, we will soon have to find a different way of processing information if we want to continue to reap the benefits of rapid growth in computing capabilities.

Enter qubits.

In classical computing, a “bit” can have a value of 0 or 1, and whether it is one or the other is represented by the state of a transistor.

Quantum computers, however, don't use traditional transistors or bits. The basic unit of processing in quantum computing is the qubit, a quantum state that can represent multiple values at the same time. Stringing qubits together could allow some computations to be run much faster than otherwise possible, as the information represented by the qubits can be processed simultaneously.

Quantum Computing	Vs. Classical Computing
 <p>Calculates with qubits, which can represent 0 and 1 at the same time</p>	 <p>Calculates with transistors, which can represent either 0 or 1</p>
 <p>Power increases exponentially in proportion to the number of qubits</p>	 <p>Power increases in a 1:1 relationship with the number of transistors</p>
 <p>Quantum computers have high error rates and need to be kept ultracold</p>	 <p>Classical computers have low error rates and can operate at room temp</p>
 <p>Well suited for tasks like optimization problems, data analysis, and simulations</p>	 <p>Most everyday processing is best handled by classical computers</p>

Quantum computing is still years away from being useful in practice and won't replace classical computing in many contexts. But companies like IBM are already providing access to early quantum computing infrastructure – which would be highly expensive for most companies to operate themselves – through the cloud. Google, one of the field's leaders, wants to build a commercial quantum computer by 2029.

Now, as quantum heats up, [a handful of big tech companies are investing heavily in the area](#) to gain an edge.

Thanks to the combined efforts of big tech leaders, as well as researchers around the world, the number of qubits inside quantum computers has increased at a steady rate over the last 2 decades. In 2018, when Google announced it had built a 72-qubit quantum computer, the qubits chart suddenly started to look more and more like the beginning of an exponential curve. Three years later, IBM unveiled a quantum computing chip with 127 qubits. It has since promised to deliver a 1,000-qubit chip by the end of 2023.

In 2018, Hartmut Neven, director of the Quantum Artificial Intelligence lab at Google, made a similar observation to Moore's: Quantum computers were developing at a "doubly exponential" rate when compared with classical computers.

This double exponential effect comes from a qubit's ability to represent both 0 and 1 at the same time: for every qubit added to a quantum computer, the combinations of information it can handle increases exponentially. As such, if qubits themselves are being added exponentially, then the resulting computing power would actually increase at a double exponential rate – at least in ideal conditions.

Double exponential growth could lead to rapid advancements that far outperform what came immediately before. Nevens summed it up as such: “It looks like nothing is happening, nothing is happening, and then whoops, suddenly you’re in a different world.”

Though measuring the true “power” of a quantum computer is a vexing task – there’s much more that plays into a quantum computer’s capabilities than raw qubit number – some experts believe that “Neven’s Law” will eventually become as prophetic as Moore’s Law.

TAKEAWAY

Moore’s eponymous law – what Intel later called “the golden rule for the electronics industry” – presaged and, in some ways, propelled a revolution in both the performance and price of computing power. Even today, when Moore’s Law may finally be reaching its limit, a successor to Moore’s Law – that predicts similar exponential increases in computing power – is emerging to help define the next age of computing.

2 Metcalfe's Law: Why big networks produce colossal winners

In 1980, Robert Metcalfe – one part of the duo behind Ethernet, a technology for connecting computers together – observed that communications networks increased in value in proportion to the number of users.

One telephone in a city was useless. Even a few hundred, dispersed across millions of inhabitants, wasn't very useful. But once you had friends and family who owned telephones – and restaurants and movie theaters and stores all had telephones – it was suddenly very desirable to own one.

Metcalfe was working at 3Com, a computer network company he co-founded, when he came up with the idea. At the time, he was trying to understand why his company's local area network (LAN) starter kits – which allowed 3 PC users to share a printer and a hard disk – weren't selling.

What he found wasn't necessarily that the cost was too high, but that 3 people didn't constitute a sufficiently large network to justify the purchase.

Like early installations of the telephone, the cost of connecting the network initially exceeded the value generated – but there would also be a point of critical mass where the benefits of connecting did exceed the upfront investment.

HOW FACEBOOK LEVERAGED METCALFE'S LAW TO BECOME A SOCIAL MEDIA GIANT

When Facebook first launched at Harvard University in 2004, it was like a singular telephone: pretty useless. But it wasn't useless for long. Within a month, 50% of the student body was on Facebook.

This was partly a case of pent-up demand: every year, Harvard released a physical “face book” that included every student and faculty member on campus, designed to help everyone at the school get to know one another. “TheFacebook,” as the social network was then known, was an attempt to digitize this already existing physical object and make frictionless the process of snooping on other people’s photos.

With every new Harvard student that joined the platform, it became more valuable for other students to join too. Each new student promised new people to look at, learn about, and “poke.”



SOURCE: GETTY

As Facebook grew, the company added features that were designed to tap into the “more is better” mechanics of Metcalfe’s Law. Photos, groups, likes, comments – they all leveraged this concept to bring users back into Facebook again and again. And the more people there were on Facebook, the more photos were uploaded and tagged, the more groups were created, and so on.

The non-linear effects of Metcalfe's Law would also contribute to the long-term viability and success of Facebook's advertising business. The more people on the platform, the more data and connections, the greater the revenue opportunity in targeted advertising.

Metcalfe himself acknowledged the impact of this law on Facebook's growth, observing that if Facebook's revenue was used as a proxy for its value, it was impossible to deny that the network's value had risen exponentially compared to the steady linear growth of its user base.

WHEN NUMBERS BOOST USEFULNESS: WHY PEOPLE CONTRIBUTE TO OPEN-SOURCE PROJECTS FOR FREE

The past 2 decades have been awash with stories of billionaire tech founders and massive startup valuations. Growing steadily alongside their for-profit competitors were free, open-source software like WordPress and Wikipedia, which amassed cult-like followings, gained millions of customers, and became quality leaders and innovators in their industries. It's safe to say that anyone who's Googled anything in the past decade has visited a website powered by WordPress and learned something new from Wikipedia.

Lots of software are free and open-source, yet they haven't achieved a fraction of the scale of WordPress and Wikipedia. Steven Weber, a professor at Berkeley, cracked the puzzle in 2004: To succeed, open-source software doesn't just have to be shareable. It also has to be anti-rival, which means that its value should increase as more people use the software. Once again, network effects are the differentiator.

The more people use WordPress, the more use cases they discover, which give rise to new features and tools. This massive user base also incentivizes an ecosystem of individual designers and developers, as well as small and growing businesses, to create website solutions like design themes and plugins. Users can choose from more than 9,000 themes and 59,000 plugins, some free and some for a fee.

Likewise, the more people use Wikipedia, the more information gets added, flagged, and edited. Wikipedia relies on over 200,000 volunteers to create records, link to sources, and check entries for unsupported, unverifiable, and inaccurate statements.

These efforts pay off by improving usefulness and quality. In fact, people find so much value in these anti-rival goods that they donate money to keep WordPress and Wikipedia running.

DEATH SPIRAL: WHAT HAPPENS WHEN NETWORKS DON'T REACH CRITICAL MASS

According to Metcalfe, a network's critical mass is a function of the cost of a new connection (e.g., the cost of acquiring a user), the number of users, and the value of each connection.

This approach describes a few mechanics key to network effects. The lower your cost per connection, for example, the lower the number of users that you'll need to hit critical mass — and the same goes for a higher value per connection.

Andrew Chen, general partner at Andreessen Horowitz, has written not only about how the mechanics of Metcalfe’s law are key to growth at companies like Facebook, Twitter, and Snap, but also how the law can work against these networks if user retention isn’t high enough.

In one scenario, each new user that comes to a platform makes the experience better for everyone else. Growth is self-perpetuating because not only are users encouraged to stick around, but there’s an increasing value proposition to entice new users at a steady clip.

In another, not-so-good scenario, new users may still be joining and helping generate the benefits of network effects (for a time), but the network’s retention isn’t great.

If a network starts losing too many users, it can get stuck in what he calls a “social network death spiral.” Because while the forces described by Metcalfe’s Law can help startups gain lots of users quickly, they can also cause those startups to lose users at the same pace. In other words, “as you lose users, the value of your network decays exponentially.”

What’s lacking in this reverse-network effect scenario is the value offered by joining the network. If that value is too low, retention will be low, and it becomes difficult to reach critical mass. Meanwhile, if the value is high, your retention will likely be high as well.

Chen’s death spiral provides a way of understanding what happens when social networks take off for a time but ultimately don’t give each new user sufficient value to stick around — a fate that has befallen prospective social networks from Ello to Path to Peach. These apps often launch with a bang, attract initial interest, and then rapidly lose their user base as the promise of the network fails to materialize.

TAKEAWAY

The telephone, the fax machine, and early LAN networks had what venture capitalists today call “network effects.” The more people that were on these networks, the more valuable it was for new users to join.

This is one of the most important dynamics for businesses in the 21st century, especially in tech. Many of the most valuable businesses of the last several decades, from Microsoft to Facebook to Airbnb to Uber, have succeeded in part thanks to the power of network effects.

While building a business through network effects was valuable before the internet, the internet has made it possible for companies to grow their user base at exponentially faster speeds, often building their entire business models on this growth.

3 **Law of Mobility: The value of making products available anywhere, any time**

In 2003, Russ McGuire, Sprint's director of strategic planning, observed that wireless technology would be the catalyst of another revolution like the microprocessor and internet revolutions before it.

He had noticed that, while the recently released BlackBerry 8510's interface was nowhere near smooth, the convenience of accessing email anytime, anywhere made the device ubiquitous in executive circles. Wireless devices were even starting to transform non-tech businesses – rental car services, for example, brought mobile barcode scanners and credit card readers to parking lots to speed up the process of car returns.

Based on these observations, and drawing on both Moore's Law and Metcalfe's Law, McGuire formulated the Law of Mobility, which states that "the value of any product increases with mobility."

He first wrote about this law in 2005, adding, "If a product is available for my use an increasing percent of my time, without a significant increase in cost (in terms of product cost, operating cost, and convenience), then it will be more valuable to me. As the cost of building mobility into products approaches zero, then virtually every product will become mobile."

McGuire further identified 2 value enablers of mobility: contextual relevance and convergence.

The former meant using information about a user's individual conditions and needs to make the product more valuable. Today, we simply call this "personalization."

Convergence, meanwhile, meant bringing together formerly separate objects into one unit. In other words, mobility doesn't simply mean making a mobile, digital version of a physical or on-premises product or service. If that version can't adapt to the user's changing circumstances and fulfill its function without requiring multiple gadgets, then it hasn't succeeded in delivering the value of mobility.

NETFLIX AND THE OPPORTUNITIES – AND COSTS – OF MOBILITY

From the very start, Netflix co-founder Reed Hastings knew that the company, which launched a DVD rental-by-mail service in 1998, would evolve to offer video streaming through the internet. But he also knew the time was not right.

Many pieces of the puzzle had not yet fallen in place: bandwidth to stream videos without an hours-long wait, buy-in from movie studios in the viability of streaming (which affected their willingness to grant streaming rights to their films), suitable mobile devices, and consumer demand.

Netflix eventually introduced video streaming in 2007 – in part after observing the explosive popularity of YouTube videos despite their low resolution – and the new model took off. By 2010, it had made streaming available on mobile devices from Roku to Xbox to the iPhone. By 2012, Netflix's "subscriber base rivaled that of Comcast, the largest U.S. cable operator," according to author and journalist Gina Keating.

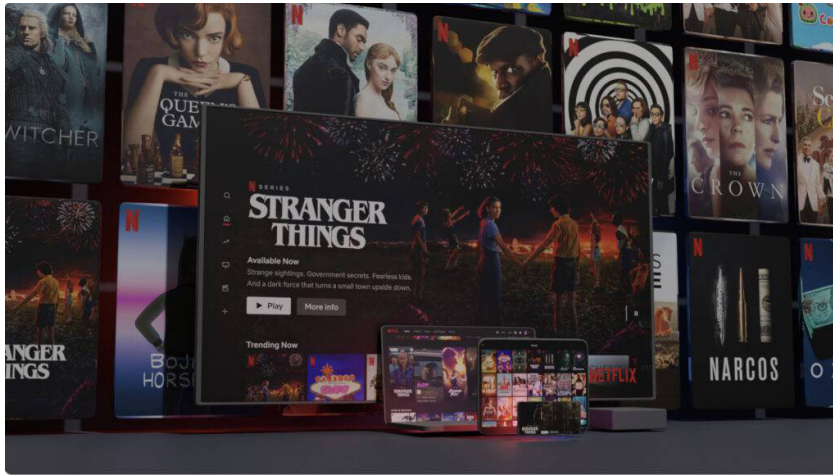
Today, customers no longer need a DVD, a DVD player, and a TV set to watch a film. All they need is a single device, be it a mobile phone, laptop, or smart TV – and a streaming service like Netflix. In the years since Netflix’s pivot to streaming, the video streaming industry has ballooned, while DVD rental is steadily declining.

Netflix now has 222M paid subscribers globally, making it the world’s largest streaming service. Its net income grew from \$49M in 2006 – the year before it introduced streaming – to almost \$5.1B in 2021.

Context – a key ingredient in the law of mobility – has been a major contributor to Netflix’s streaming success.

Netflix first introduced a recommendation system as early as 2000, when it began feeding user reviews into its Cinematch algorithm to recommend similar movies that users might like. Netflix continued to improve its algorithm over the years to more accurately predict how much a customer would like a given film.

Once it began offering streaming services in 2007, Netflix started to feed its algorithm a wealth of data on the context of every film view – from how each user navigated the homepage to whether or not they finished a movie. These actions, along with many other factors, now fuel Netflix’s recommendation engine, which today drives more than 80% of users’ discovery of new TV series on the platform.



SOURCE: NETFLIX

When Netflix began producing original content in 2013, it used its insights into its subscribers' preferences to predict how well shows would perform. In a sign of this model's success, "House of Cards," Netflix's first original series, became the platform's most-watched title in a matter of weeks.

However, the cost of building mobility into Netflix's products is still far from approaching zero.

As a pioneer of mobile streaming, Netflix spent more than a decade and millions of dollars figuring out digital delivery, developing its technology, and battling incumbents and regulatory roadblocks.

The streaming giant also relies on Amazon's AWS cloud service — as opposed to building and owning the requisite data centers itself — to host its streaming services and recommendation engine. As of 2020, Netflix planned to spend around \$1B on its streaming and cloud computing requirements through 2023. It's estimated that Netflix is AWS' top-spending customer.

Netflix also pays hefty sums to maintain its competitive positioning. It has recently been losing titles from its catalogue, either because studios have pulled them to host them on their own streaming services or due to competition from the likes of Amazon Prime Video and Hulu. Netflix produces original content to mitigate these losses, but production is expensive – it is estimated to have spent \$13.6B on original content in 2021, a figure that could reach \$18.9B by 2025.

Nevertheless, Netflix deserves credit for making video streaming on mobile devices a mainstream pastime. The years and resources it spent on innovation paved the way for other streaming services, lowering the cost of mobility for those competitors.

HOW CHALLENGER BANKS AIM TO CONNECT AND SIMPLIFY EVERY ASPECT OF A CUSTOMER'S FINANCIAL LIFE

Challenger banks epitomize the belief that with mobility, a product or service becomes immensely valuable.

In the late 2000s, banking customers still relied heavily on bank branches and were just starting to embrace online banking. But challenger banks bet that online – particularly, mobile – would be the next channel for retail banking distribution. That has proven prescient.

Challenger banks – tech companies that leverage software to digitize and streamline retail banking – tend to be branchless, living entirely online. They offer a mobile-first experience for every aspect of banking, from savings and money transfers to payments, asset management, and loans. Some also offer an ecosystem of financial products like insurance and investments, all through a single mobile platform.

For instance:

- UK-based [Monzo](#) is creating a financial hub for bills management and incentivizing customers to switch to more eco-friendly energy providers so they can save on energy costs. It also offers a buy now, pay later service, tax-free savings, and a feature for assigning savings to different budgets.
- UK-based [Revolut](#) offers investment and insurance products, a “shopper” feature that finds discounts online, a single-use card to secure online purchases, and travel-related financial services.
- Brazil-based [Nubank](#) has integrated major e-commerce platforms into its app, allowing customers to shop without ever leaving the bank’s app.

These ecosystems reflect many fintechs’ ambitions to become [super apps](#) that unite all the financial products and services a consumer could need.

While brick-and-mortar banks could offer the same services, customers would not be able to use them anytime, anywhere. Mobility increases the likelihood of customer engagement in these products and services, especially among younger consumers. EY found that among global consumers who use challengers for their primary financial relationships (PFRs), 37% are ages 18 to 34.

Convenience is one of the main draws of challenger banks. Customers don’t need to visit a physical branch to open an account. Loan applications can be processed more quickly, as many use readily available customer data — such as spending habits, bills payments, and savings — to determine credit score. Customers can become more proactive in saving money by using app features like separate wallets for certain goals (e.g., a travel budget account).

Mobility also enables banks to offer personalized products – such as customized benefits and tailored product features – to consumers who have come to expect this degree of individualization in digital interactions.

Despite the challenges inherent to online banking – including security issues and service outages – more consumers are adopting it worldwide, accessing these services from both challenger banks and incumbents that have pivoted to offer digital banking to keep up with shifting consumer demand.

The market size of challenger banks is expected to grow from \$47B in 2021 to just over \$2T in 2030. Insider Intelligence predicts that by 2025, the number of US residents with digital-only bank accounts will nearly double to 53.7M, up from 29.8M in late 2021.

TAKEAWAY

McGuire recognized that mobility would usher in new ways of working and change customers' expectations for interacting with businesses.

Beyond the evolution of gadgets, the mobility revolution has extended to services and legacy businesses. Nearly every sector – from retail to financial services to healthcare delivery – has to some degree gone digital and, therefore, mobile.

There's a mobile app (or a SaaS product) for virtually every personal and business need and desire. Mobile app development is now an expected part of every software product's roadmap – and failure to fulfill that can make customers defect in droves to the nearest competitor.

4 Gall's Law: Why the best products are built from simple systems

In 1975, the American author and pediatrician John Gall made an observation about systems that would go on to become hugely influential in computing.

"A complex system that works is invariably found to have evolved from a simple system that worked," he wrote. "A complex system designed from scratch never works and cannot be patched up to make it work. You have to start over with a working simple system."

This idea, originating in his book "General Systemantics," would eventually become known as Gall's Law. Since then, it has become one of the most important and foundational ideas in tech, especially when it comes to the design of new products.

HOW GALL'S LAW FOSTERS PRODUCTS THAT PEOPLE ACTUALLY WANT TO USE

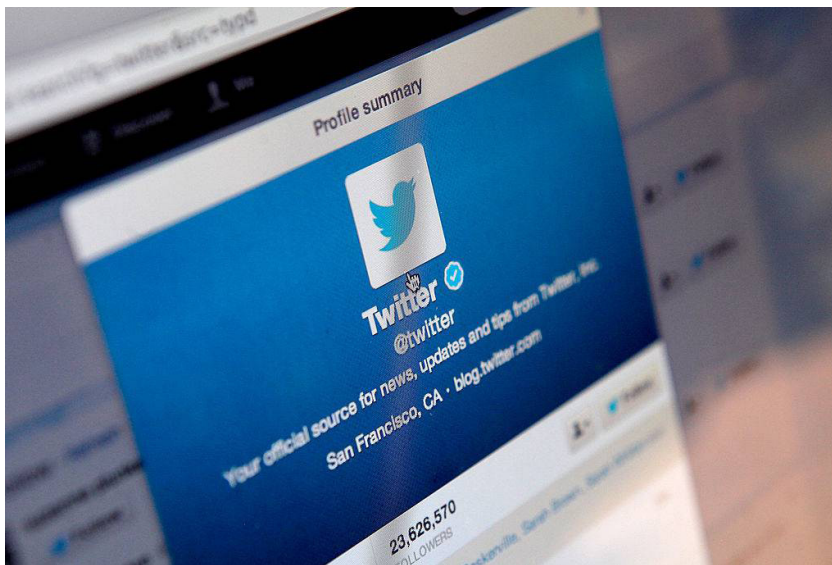
One concept that can help illuminate the power of Gall's Law is the idea of emergence: the notion that complex systems possess properties that the individual constituents of the system do not.

Powerful technologies can be built when users can manipulate and recombine individual pieces in order to create something of greater value.

Twitter, for example, was founded without reply or retweet functions. All users could do was post 140-character messages to their feeds.

Then Chris Messina proposed the idea of using a hashtag to collect all tweets relating to a single topic, and not long after, hashtags were added into the Twitter platform itself.

Replies and retweets followed later as the Twitter team realized that their users were trying to approximate these actions within their constrained set of possible behaviors.



SOURCE: GETTY IMAGES

Twitter could have predicted that users of a public broadcasting app would want some way to share content with their followers, reply to other people’s messages, and even organize information along topical lines.

But cultivating an openness to what users actually did, rather than assuming they knew exactly what to build, was key to how Twitter developed into the product it would become. As co-founder Evan Williams put it, it wasn’t even clear “what [Twitter] was” in those early days, and the product took several turns before settling on what it is now.

This kind of emergence-friendly development process is more or less conventional wisdom in Silicon Valley today, with the creation of a “minimum viable product” (MVP) being standard practice within product businesses.

The thinking behind the MVP is simple and echoes the development processes followed both by Amazon in inventing AWS and by the team at Twitter: build a version of your product with just enough functionality that it can engage your core users. Roll that initial version out and start gathering data and feedback. Then, use that data and feedback to chart the path forward. When done right, the result is what Gall’s Law predicts: a complex system that has arisen from the organic growth of a much simpler system.

HOW AMAZON FOLLOWED GALL’S LAW TO DOMINATE CLOUD SERVICES

The development of Amazon Web Services (AWS) from an internal data storage service to a crucial backbone of the internet encapsulates the power of Gall’s Law.

AWS emerged from a relatively simple desire: Amazon needed a more modular system for its internal engineering work.

Every time a team at Amazon wanted to build out a new web feature, whether for internal use or for a retail partner, Amazon developers would end up building everything – from databases to computing power to storage – from scratch.

That's when the idea emerged to build out each one of these functions as a modularized service. For one, the team had become highly competent at working on these functions. For another, the Amazon team was operating at a unique scale, being one of the biggest early internet companies.

It soon became clear that Amazon had an equally unique opportunity to make databases, computing, and storage easier for other companies and other developers to use.

"We tried to imagine a student in a dorm room who would have at his or her disposal the same infrastructure as the largest companies in the world," Andy Jassy, then-head of AWS, told Brad Stone, author of "The Everything Store." "We thought it was a great playing-field leveler for startups and smaller companies to have the same cost structure as big companies."



SOURCE: FLICKR

While AWS today appears to be an intimidatingly complex system, it emerged from the interaction of a very simple set of modules designed for Amazon to use internally.

The concept was heavily inspired by the 2003 book “Creation,” written by game designer Steve Grand. Grand’s design approach had centered around building simple creatures, then sitting back and “watching surprising behaviors emerge.” For Amazon’s then-CEO Jeff Bezos, this idea proved perfectly applicable to computing.

On the influence of “Creation” across Amazon’s leadership, Brad Stone writes,

“If Amazon wanted to stimulate creativity among its developers, it shouldn’t try to guess what kind of services they might want; such guesses would be based on patterns of the past. Instead, it should be creating primitives – the building blocks of computing – and then getting out of the way. In other words, it needed to break its infrastructure down into the smallest, simplest atomic components and allow developers to freely access them with as much flexibility as possible.”

Today, Amazon Web Services drives over half of Amazon’s operating profit and about 16% of its overall revenue.

FEATURE BLOAT: WHAT HAPPENS WHEN COMPANIES FLOUT GALL'S LAW

Tech history is full of products that failed because they went the opposite way of Twitter and Amazon: they built over-engineered products, bloated with features, that lacked enough simple utility for users.

One of the classic examples of feature bloat is ICQ, the once-popular instant messaging client. In a 2007 blog post, Robert Scoble attributed the service's sudden decline mainly to over-engineering: "It got too cluttered and stopped being developed."

At one point, ICQ had been the king of instant messaging. In 1998, AOL bought the company behind ICQ for \$287M upfront and \$120M in performance-related payments. The service had more than 100M accounts registered around the time it peaked in 2001.

For Scoble, that's where the decline, and the clutter, started. ICQ started branching out from its core utility by adding features around shopping, music, games, and even careers, resulting in a busy interface that felt removed from the purpose of the product.

At the same time, new products like Facebook were emerging that allowed people to more organically connect with their friends outside the constraints of a messenger.

"Feature bloat is how most consumer web and desktop products suffocate themselves," Dropbox co-founder and CEO Drew Houston has said. "ICQ became so comically bloated that they released an 'ICQ Lite' version, but by then they were already on the decline."

Releasing a “lite” version of a product might make sense if you’re differentiating it from a “premium” version that serves an audience with more complex needs, but ICQ and ICQ Lite were the same basic product for the same basic target user.

The release of a “lite” version, in this case, could be seen as a way to actually offer a superior product – an attempt to cut the feature bloat that had crept in over the years.

TAKEAWAY

Today, with free trials, SaaS, and subscription business models in vogue, it’s more sensible than ever to produce a minimum viable product (MVP) and iterate on it based on user feedback.

While companies should avoid feature bloat, aiming for a gradual evolution from simple functionality is a way to build a growing product that actually reflects users’ needs.

And according to Gall’s Law, products that start small – as simple, effective systems – stand the best chance of becoming powerful businesses built on high-functioning, complex systems.

5 Law of Modularity: Why building blocks are essential to modern tech design

“The entire concept of computers has changed,” wrote IBM in a brochure it released in 1964. It wasn’t exaggerating. The company was launching System/360, a new modular computer architecture that let customers select the speed, memory size, and processing features they needed with the option to upgrade and expand them in the future.

Until then, IBM computers had been incompatible with one another. For businesses — the primary buyers of computers at the time — this meant having to buy a completely new machine if their data processing needs or business use cases changed. One couldn’t simply change out parts (like keyboards), connect to any printer, or use the same software or operating system across different computer models. Installing a more powerful processor meant extensive reprogramming.

The shift to the modular structure embodied by System/360 was largely driven by one IBM engineer, Bob Overton Evans. He persuaded the company’s chairman, Thomas Watson Jr., to stop building incompatible machines and instead develop a range of modular computers that would feature common building blocks.

Modularization has since transformed both hardware and software. Adenekan Dedek, a technology and information systems professor at Northeastern University, encapsulates this impact in what he calls Evans’ Law of Modularity: “The inflexibilities, incompatibilities, and rigidities of complex and/or monolithically structured technologies could be simplified by the modularization of the technology structures (and processes).”

The impact of Evans' advocacy is most obvious in computer hardware today. Even in closed systems like Apple's Mac products, a user can connect external hard drives, keyboards, monitors, and cameras, as well as link to other Apple devices.

But modularization has also spurred a revolution in software.

HOW SALESFORCE KICKSTARTED THE MODULARIZATION OF SOFTWARE

In February 2000, at a Siebel Systems conference in San Francisco, 25 (hired) protesters in red T-shirts called for the end of software and gave out invitations to Salesforce's launch party later that day. To enter the party, one had to bring a piece of old software to toss into a trash bin.

Customer relationship management (CRM) platform Salesforce's beef with existing software was that it was bound on-premises, within devices. And once you installed it, you were generally stuck with whatever features it came with. Salesforce founder Marc Benioff believed software should be made available 24/7 via the cloud.



SOURCE: MARC BENIOFF VIA TWITTER

Salesforce went on to become one of the first-ever completely cloud-based products, forcing competing CRM giants to follow suit. Today, it is used by more than 150,000 businesses. The company reported a revenue of \$26.5B in fiscal year 2022.

One reason Salesforce is so popular across many industries is that you can customize the system to suit different use cases. The Salesforce Customer 360 platform has cloud modules for sales, marketing, community, service, commerce, and analytics. On top of that, hundreds of apps can be connected to the platform through APIs. Third-party developers can build custom apps and extensions on top of the platform, too.

HOW APIS ENABLE MODULARITY AND GENERATE MASSIVE VALUE

In the early 2000s, three companies – Salesforce, eBay, and Amazon – demonstrated how making tech more compatible can boost operational efficiency, increase user adoption, open up access to more products and services, and improve commercial performance. They did this through the use of application programming interfaces (APIs) – code that lets 2 different systems work with each other.

APIs were a part of Salesforce’s business model from the beginning. Two weeks before its notorious protest, the company launched its web API, which let customers share data across different web applications. Many consider this to have been the first modern API.

In the same year, eBay launched an API for developers to create systems that helped sellers manage their growing businesses. And in 2002, Amazon launched an API that enabled third-party sites to display and search products from the e-commerce website.

“APIs were a convenient way for companies with a lot of resources to share those resources with developers and expand their footprint, spark new ideas and accelerate development.”

— *Matt Hawkins, co-founder & CEO of Hoss*

Modularity through APIs unlocks new use cases for both tech and traditional companies. E-commerce sites show you what products are available thanks to APIs connecting the site to inventory software. In countries with open banking APIs, third-party apps can access financial information — like bookkeeping software that reflects a company’s bank account balance — in real time.

Today, modularity isn’t just an attribute of tech products — it’s a product in itself. One of the earliest examples is the phone-calling API that Twilio started selling in 2008, which enabled apps to make cloud-hosted phone calls.

Many tech products today rely on communication APIs to improve usability and customer service, from ride-hailing apps that let drivers and passengers call each other to banks that use SMS and video for identity verification.

THE INTERNET OF THINGS: A NEW FRONTIER FOR FRAGMENTATION

Modularity could pave the way for the future of the smart home.

Until now, internet of things (IoT) architectures have tended to support a specific use case (for example, a smart home component or a device for health monitoring), but this can make it difficult to efficiently stitch together different IoT devices and systems into a network.

For example, a homeowner shouldn't have to buy all of their appliances and electronics from a single vendor to have a smart home. Instead, one could put together a Nest thermostat, a Philips lamp, an Amazon Echo speaker, and a Xiaomi robot vacuum, and then manage them all through a single interface, storing data from these devices in the same location in the cloud.

In the same spirit as its 2002 API launch, Amazon has been an early mover here, seeking to build interoperability among intelligent voice assistants (IVAs). If it succeeds, its own IVA, Alexa, would work with any smart device. Google and Apple have largely resisted the initiative. However, the 2 have teamed up with Amazon and the Connectivity Standards Alliance on Matter, a connectivity standard that promotes interoperability among smart home devices and IoT platforms.

Businesses have more complex IoT use cases, and that amplifies the need for modularity. Consider factories that require a single analytics platform for all of their smart sensors and machines, manufacturers that must put together an automated assembly line using varying robotics and smart technologies, and fleets of autonomous trucks that need to maintain connectivity when they enter another country with different network service providers.

TAKEAWAY

Across both hardware and software, modular design can drive a step-change in the value of tech products and services – both for the end-user, and for the business itself.

Flexible, customizable, and interoperable systems can also invite outside engagement, whether by creating space for third-party developers to build on top of a platform or through APIs that create new value via interconnected services.

6 The 2-Pizza Rule: Why small teams lead to big success

Jeff Bezos didn't just leverage the basic insight of Gall's Law to develop what has become Amazon's fastest-growing internal unit – he is also responsible for popularizing one of his own tech rules.

In early 2002, Bezos decided that to reduce communication overhead and improve productivity, all of Amazon would be re-organized into so-called “2-pizza teams” – squads small enough that just 2 pizzas would be enough to fully feed them when working late in the office.

While this has also been interpreted as a rule for meetings – either capping the maximum size of meetings at Amazon or capping the size of meetings that Bezos himself will attend – the original formulation was an edict about team size.

WHY 2-PIZZA TEAMS GAVE AMAZON A STRATEGIC ADVANTAGE

According to analyst Benedict Evans, the 2-pizza rule has been a crucial structural advantage for Amazon.

The idea stemmed from Bezos' attitude toward both centralization and communication. Bezos wanted Amazon to remain nimble even as it grew, and that started with encouraging independent decision-making rather than an over-reliance on hierarchy.

And Bezos hated the overhead created by excess communication. In “The Everything Store,” Brad Stone quotes Bezos as saying that “communication is a sign of dysfunction. It means people aren't working together in a close, organic way. We should be trying to figure out a way for teams to communicate less with each other, not more.”



SOURCE: GETTY IMAGES

The main advantage this conferred upon Amazon was simple: having smaller and independent teams meant the company could spin up new teams much faster, giving it the power to scale more cheaply, explore new ideas easily, and ultimately ship more products to customers.

While many of the new initiatives spearheaded by these small teams have failed, some – like Prime, Kindle, and AWS – have become core businesses for Amazon.

For Bezos, it comes back to the idea of building a corporate structure that can generate the maximum amount of innovation for customers. Big, centralized teams maintain companies. Small, autonomous teams find new ideas.

“We have the good fortune of a large, inventive team and a patient, pioneering, customer-obsessed culture – great innovations, large and small, are happening every day on behalf of customers, and at all levels throughout the company,” he wrote in his [2013 shareholder letter](#). “[Decentralized] distribution of invention throughout the company – not limited to the company’s senior leaders – is the only way to get robust, high-throughput innovation.”

In other words, as Benedict Evans writes, “you don’t (in theory) need to fly to Seattle and schedule a bunch of meetings to get people to implement support for launching make-up in Italy, or persuade anyone to add things to their roadmap.”

There is a machine, and then there is a machine to build the machine, says Evans. And Amazon’s 2-pizza rule ensures that the company, aside from being a ubiquitous tech leviathan, can operate a machine to build more Amazons.

NOT EVERYONE IS SATISFIED WITH 2 PIZZAS

According to organizational psychologist J. Richard Hackman, the more interconnections you have between people, the slower your decision-making and the higher the management costs for the organization.

As you add people to an organization, the number of communication links increases exponentially.

“The cost of coordinating, communicating, and relating with each other snowballs to such a degree that it lowers individual and team productivity,” writes blogger Janet Choi. Two-pizza teams solve this inherent scaling problem by artificially capping the number of links.

Despite the organizational clarity derived from the 2-pizza team, not everyone at Amazon and in the wider tech world is a fan of the rule.

Some former Amazon employees and others believe the strategy was counterproductive, and some think that building products using 2-pizza teams creates a disconnected user experience.

According to Brad Stone, the concept of 2-pizza teams was unevenly applied throughout Amazon. It took root most of all in engineering, while the idea barely touched the finance and legal departments.

Additionally, each team had to set up its own “fitness function” – some kind of quantitative, linear equation that could be used to judge whether that team succeeded or failed in its mission. For a marketing team, that might be the average email blast open rate multiplied by ensuing order value.

Yet, Stone writes, making some teams define this function for themselves was like “asking a condemned man to decide how he’d like to be executed.” For others, it was merely ineffective.

Kim Rachmeler, former VP at Amazon, said, “Being a two-pizza team was not exactly liberating. [...] It was actually kind of a pain in the ass. It did not help you get your job done and consequently the vast majority of engineers and teams flipped the bit on it.”

Small, disconnected teams run the risk of producing disjointed products that don’t contribute to a seamless experience, according to product management consultant Matt LeMay. For LeMay, the most important aspect of a product for a customer is not the individual features of a product, but how those features come together to deliver a cohesive experience.

Given the vast number of products that Amazon has put out over the years, and how cluttered many of its website’s menus can seem, it’s fair to say that Amazon’s drive for 2-pizza teams has not been without its drawbacks.

In fact, it’s no longer the norm. In their book “Working Backwards,” former Amazon executives Colin Bryar and Bill Carr reveal that the company eventually figured out the biggest predictor of a team’s success — and it wasn’t size. Rather, it was having a leader with the relevant skills, authority, and experience to manage a team focused on a single goal.

As a result, Amazon replaced the 2-pizza model with single-threaded leader (STL) teams, which focused on one project at a time. Each time was headed up by a leader who was solely committed to that one focus area, rather than dividing their time between projects.

The adoption of STL teams enabled Amazon to realize ideas that had languished in the vagueness of the 2-pizza model. Fulfilled by Amazon (FBA), a service that gives businesses access to the company’s logistics network, finally came into life through an STL — and now provides the company a steady stream of income through fulfillment and storage fees paid by sellers.

TAKEAWAY

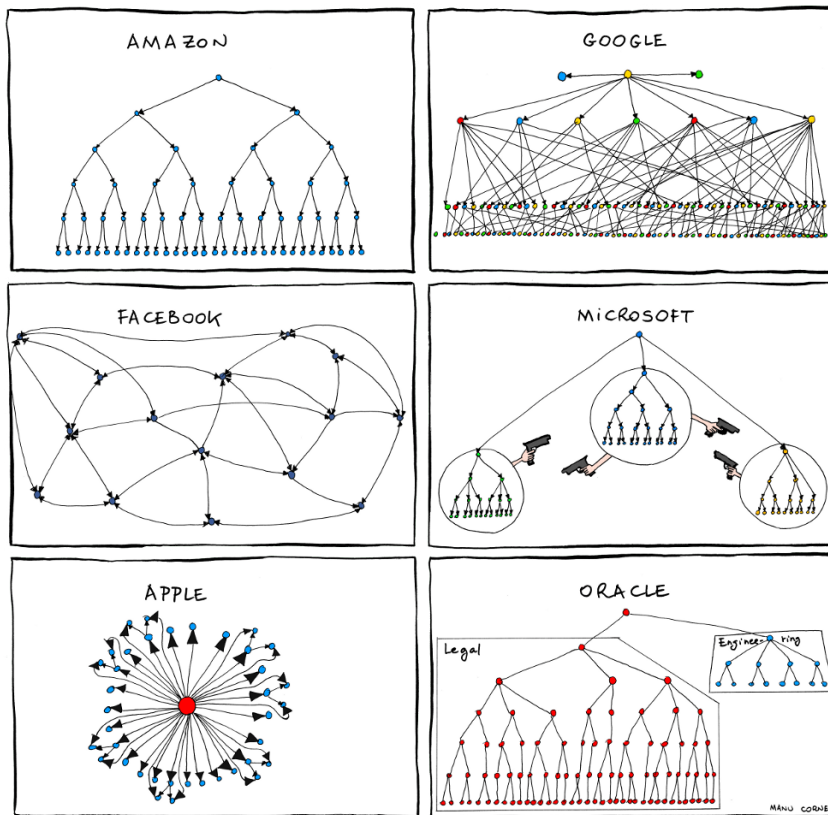
Jeff Bezos wanted small, autonomous teams that could be “independently set loose on Amazon’s biggest problems,” as Brad Stone writes. These teams wouldn’t have to waste cycles communicating with other teams, and they would each have all the resources and people necessary to launch new products. The result, Bezos thought, would be more creative offerings and faster results for customers.

Though not without its detractors, Bezos’ idea that over-communication between teams risks stoking inefficiencies has been implemented at other companies, including Google.

7 Conway's Law: Why corporate structure is vital to product development

In 1967, the computer scientist Melvin Conway made a key observation about organizational structure.

The way that a team communicated and the design of that team's products, Conway argued, were mirrored – one always reflected the other.



A DROLL ILLUSTRATION OF CONWAY'S LAW. SOURCE: MANU CORNET, BONKERSWORLD.NET

In a simple explanation of Conway's thesis, there are two pieces of software: software A and software B. If the developers of these two pieces of software don't communicate, there's no easy way for the software to integrate.

When communication between the developers occurs often and openly, on the other hand, the odds of a seamless experience are far greater.

HOW APPLE PRODUCES AN END-TO-END CUSTOMER EXPERIENCE

At Apple, teams are organized according to what's called the Unitary Organizational Form. The basic idea — rooted in Conway's Law — is that the company should be organized around functional expertise rather than products.

That means that instead of dedicated teams for products like the iPhone, the Mac, or the iPad, Apple has teams that work on design, teams that work on engineering, teams that work on marketing, and so on.

This structure encourages coordination between teams and helps Apple deliver a unified experience across products. No product is ever released that departs from the predominant Apple design, engineering, or operational paradigm. Even its credit card has an unmistakable Apple "feel."

After the iPad first launched, Steve Jobs said this “post-PC device” needed to be “even easier to use than a PC” and “even more intuitive than a PC, [...] where the software and the hardware and the applications need to intertwine in an even more seamless way than they do on a PC. [...] We think we have the right architecture not just in silicon but in our organization to build these kinds of products.”

Reorganizing Apple along functional rather than divisional lines was one of the first things Jobs did upon returning to Apple in 1997, and his successor at Apple, Tim Cook, still attributes the success of the company’s products to this move.

“We’ve found a way to make our products such that the experience is jaw-dropping,” Cook told Businessweek.

WHY GITHUB IS STRUCTURED LIKE AN OPEN-SOURCE PROJECT

GitHub provides an example of a company that obeys Conway’s Law while doing so in a remarkably different way from Apple.

Instead of integrating in order to promote an end-to-end experience, GitHub is intentionally structured like one of the open-source projects the service hosts: decentralized, autonomous, and asynchronous.

That structure reflects the kind of product GitHub has built – one designed for developers more than managers – as well as how that product works.

The GitHub tool is built for asynchronous collaboration: new code can be submitted anytime from anywhere, then reviewed at the responsible party's leisure. Developers across the globe can use GitHub to collaborate on a project without having to deal with overlapping codebase changes or inconsistencies between their work.

The company itself is also built for asynchronous collaboration, with many of its basic organizational tenets ripped directly from the processes of open source development.

There are no codified standards about what time to come into the office, and most work is completely self-directed. "If you're interested in working on something, then work on it," wrote Zach Holman, one of the first engineers at GitHub.

Members of the team are spread out across the world, there are no daily stand-up meetings, and most of the communication you have with your colleagues happens asynchronously, over chat or email.

Collaborating asynchronously to the extreme is a way for the team at GitHub to stay close to one of the biggest problems they want to solve with the company: the difficulty of collaborating asynchronously when working on software projects. One of the ways they take on this problem is through "open, easy-to-use platforms" – precisely what GitHub itself is trying to build.

Apple, in other words, uses an integrated organization in order to build products that give the customer a seamless, end-to-end experience. GitHub uses an organization structured like an open-source project because its goal is to give its user base of developers a collaboration platform that allows distributed, decentralized teams to build great products.

TAKEAWAY

Conway's Law helps explain not just how companies operate – and how their structures enable or hinder business activity – but also how they are managed.

As computer scientist Fred Brooks has pointed out, for an organization that provides some good or service, the structure that it naturally takes on as it grows is unlikely to be the ideal system for delivering that offering. Remaining flexible is essential to the organization's structure.

The people that run companies, in other words, must consider organizational design on a similar level as operations, R&D, and products. Much of what we attribute to the latter is rooted in the former.

8 Yule's Law of Complementarity: When a loss-making product is good for business

When British statistician George Udny Yule formulated the Yule distribution in 1925, he was making a mathematical observation about the growth of species per genus in some living organisms. But for Adenekan Dedeke (the same academic who described Evans' Law of Modularity), Yule's equation basically describes the relationship between two attributes, and that can apply to tech products too.

In the context of the tech industry, Dedeke calls this Yule's Law of Complementarity: "If two attributes or products are complements, the value/demand of one of the complements will be inversely related to the price of the other complement."

Simply put, a drop in the price of one complement drives up demand for the other. A classic example is how printer sales increase the demand for ink cartridges. When manufacturers make printers more affordable, demand increases; in turn, more printer usage drives up demand for ink cartridges and refills.

There's a parallel concept in business: cross-subsidization, where revenues from one product or service support a (usually complementary) product or service. For example, banks eliminate signup fees to gain new credit card account holders, and airlines offer low-priced travel insurance to increase seat sales.

Or, as Bharat Anand, a professor at Harvard Business School, explains, "Two products are complements if a user's value from consuming both is greater than the sum of her values from consuming each alone." Free songs and music videos on YouTube are the perfect complements to expensive concert tickets.

In tech, Yule's Law of Complementarity often plays out in software driving demand for hardware – or vice versa – where one tends to be profitable, and the other barely breaks even or operates at a loss.

WHY SOFTWARE IS THE BATTLEGROUND FOR HARDWARE MAKERS

In May 2021, during testimony in the antitrust trial between Epic Games and Apple, Microsoft's Lori Wright, the head of business development at Xbox, revealed that the company has never made a profit on the sale of any Xbox gaming console.

But it doesn't really matter because the revenue from game sales and subscription services more than makes up for the loss. In fact, Xbox content and services made up 77% of Microsoft's gaming revenue in 2021 – a year that saw the Xbox division's highest revenues ever – driven largely by first-party games and Xbox Game Pass subscriptions.



SOURCE: XBOX

A similar story plays out at Sony: game consoles, in general, are low-margin products, but they introduce consumers to a whole ecosystem of games, accessories, related content, and services.

The practice of relying on the revenues of peripherals and games is so commonplace that when Nintendo announced the Switch (then codenamed NX), company president Tatsumi Kimishima felt compelled to announce it wouldn't be selling the device at a loss. The Switch console and ecosystem are turning out to be one of Nintendo's most successful: since its 2017 launch through Q1'22, Nintendo has sold 100M+ hardware units and 800M+ software units.

Software can equally serve as a driver for hardware sales. During the early years of Apple's iPod, for instance, it was the hardware that subsidized music sales, even though the low price of the latter drove up demand for the former.

Over time, analysts have observed how iTunes served as a loss leader for the iPod. In 2007, the LA Times quoted research analyst Susan Kevorkian as saying, "iTunes was developed to promote iPod hardware sales. The introduction of services that offer digital music to the installed base of iPod users will help drive more iPod sales." This was six years after the launch of the first-ever iPod and four years after the launch of the iTunes Music Store and its 99-cents-per-song pricing model.

For every 99-cent song iTunes sold, a paltry 9.1 cents went to the artist out of a total 70 cents that went to the record label. The 29 cents left over for Apple went more toward the maintenance of iTunes rather than to profit. That's a striking contrast to the gross profit margin of 50% or more on some iPod models, according to several analyses.

But iTunes boosted the iPod's appeal beyond just being a slick piece of hardware. It made it quick and simple to download songs to an MP3 player – literally one click.

“Steve Jobs had apparently learned, painfully from the earlier Mac experience, the importance of having software for your hardware: He had spent months negotiating with the major recording studios to ensure that on the day the iPod was launched, it had a song library to ensure its value. The price was right, too: 99 cents for a legal download. And the system was open: iTunes software could be installed on a PC, making the iPod compatible with the largest computer platform. Apple’s ability to produce a great product was not the only game changer during the last decade. Its ability to manage complements was.”

– *Bharat Anand, The Content Trap: A Strategist's Guide To Digital Change*

TAKEAWAY

When a tech company wants to deliver an end-to-end experience for consumers, it's important to identify the component that will drive the most demand for related products and services. More than the question of sales volume or profit for this single component, the company should consider how much value it brings in terms of driving adoption of the entire ecosystem.

9 The Law of Shitty Clickthroughs: Why innovative marketing is better than expensive marketing

Coined by Andreessen Horowitz's Andrew Chen, the Law of Shitty Clickthroughs states that new marketing channels, no matter how useful at first, gradually lose their effectiveness over time.

Emails, waiting lists, referrals, search engine marketing — all obey the Law of Shitty Clickthroughs. Chen gives 3 main reasons for this trend:

- **Novelty:** The first time you use a new marketing technique on a customer, you can get them to respond out of curiosity. Humans are attracted to novelty, but they also quickly recognize patterns.
- **Fast followers:** Once news of a new channel's effectiveness gets out, competitors are fast to provide similar services, and customers start to get fatigued.
- **Scale:** While early adopters may respond positively to novel marketing efforts, the mass market will be more hesitant.

For example, Chen says, when banner ads first debuted on the website HotWired in 1994, their clickthrough rate (CTR) was 78%. They were novel, and no one had ever seen them before. But soon, every website had banner ads. By Q1'20, the CTR for ads in Facebook's News Feed was just over 1%, according to AdStage.

WHEN D2C USER ACQUISITION MET THE LAW OF SHITTY CLICKTHROUGHS

With the advent of digital marketing, it became relatively cheap to acquire new customers. Brands no longer needed to spend on billboards, print, and broadcast ads, nor did they need inventories or expensive shiny catalogs to attract customers.

Facebook and Google, with their massive and broad user bases, offered brands a particularly cheap channel for mass customer acquisition – and few industries took to it as well as the burgeoning direct-to-consumer (D2C) market.

For companies like [Warby Parker](#) and [Bombas](#), digital ads were pivotal to their rapid growth. Companies could reportedly acquire customers for as little as \$10 a piece. In 2015, D2C luggage startup Away was making \$5 for every \$1 that it spent on Facebook, according to co-founder and former CEO Steph Korey.



SOURCE: GETTY

A key driver of this success was the fact that these digitally native brands, often led by founders or co-founders with technical backgrounds, were effective at leveraging the powerful targeting options available on platforms like Facebook.

The contact lens startup [Hubble](#), for example, found many of its first customers through Facebook's Lead Ads, which allows users to send their already uploaded email addresses to a company. Running Lead Ads became Hubble's most powerful early sales generator.

Doing this kind of advertising on a traditional website would have required prompting users for their email address and having them manually type it out. On Facebook, it was all done in a single click.

The problem for Hubble was that Lead Ads eventually started to become less effective. So did all the other ad types that Hubble tried using. “No matter what new ads they put in an ad set, the growth rate of sales declined and the cost per acquisition went up,” Burt Helm wrote in *The New York Times*.

This gradual decrease in effectiveness – with growth faltering and costs increasing – hasn’t been isolated to Hubble. Customer acquisition costs on Facebook and Google have increased across the board by as much as 3x in recent years, according to Digiday.

For some companies, the solution has been to get off digital ads altogether. After the beverage brand [Iris Nova](#) saw its advertising costs on Facebook increase, it decided to cut its spend entirely and focus on events and physical expansion.

Other brands are increasingly joining forces to acquire customers together, with some teaming up to offer joint product lines, events, and giveaways.

In September 2019, for example, beauty brand [Glossier](#) partnered with dog-walking startup [Bark](#) to launch a line of Glossier-branded toys for dogs. This approach appeals to both companies because the brands are not direct competitors, the risk of poaching customers from each other is limited, and both brands get to target a demographic that overlaps with their target audience.



SOURCE: GETTY

However, collaborations are still not a big part of most D2C brands' customer acquisition strategies, as it can be tough to find complementary product fits that don't pose a competition risk.

HOW AXIOS BUCKS THE LAW OF SHITTY CLICKTHROUGHS BY FOCUSING ON ENGAGEMENT

While email marketing is just as subject to the Law of Shitty Clickthroughs as banner ads, the news startup [Axios](#) looks to keep its various email newsletters effective through its focus on customers with high lifetime value.

A big mistake media companies tend to make with email is to treat top-line open rates as a sacred metric, according to Rameez Tase, former VP of growth for Axios.

If some readers are opening up a newsletter 100% of the time, and another group of readers opens it 0% of the time, that doesn't mean that there's a singular 50% open rate, he argues. Instead, there are 2 cohorts: one that is highly engaged and one that isn't.

For Axios, focusing on how to improve engagement with each specific cohort — instead of obsessing over a top-line metric like open rate — is key to building a sustainable business that can overcome the natural decay of the Law of Shitty Clickthroughs.

What Axios has found is that members of the highly engaged cohort (those with high newsletter open rates) tend to be similar people in terms of profile. They're of above-average income, they're professionals, and they tend to be decision-makers in media and government. This is Axios' ideal customer profile in terms of customer lifetime value.

The more readers like this that Axios gets, the more opportunities it has to sell its premium membership products, acquire big sponsorships, and expand its business.

To grow this high-value segment, the company encourages its most engaged subscribers to share the newsletter with others. For example, Axios uses social retargeting, ambassador reward plans, and surveys — whatever will generate the “motivation [...] to share this product with [readers'] personal or professional networks.”

But it's not just about sharing — it's about the kinds of people that readers in this cohort tend to share Axios' newsletters with. Getting high-value subscribers to share the newsletter, Axios has found, tends to bring in other subscribers like them. “An 80 percent open rate user begets another 80 percent open rate user,” Tase says.

TAKEAWAY

In an effort to outrun the Law of Shitty Clickthroughs, marketing channels and advertising platforms are always looking to develop new kinds of targeting methods and formats.

At Facebook, the race has generated new features ranging from geotargeted ads to Custom Audiences to Lead Ads to split testing. Each new type of feature Facebook puts out promises greater specificity, better results, and less competition.

But what may be more important in determining marketing success is how companies define their markers of success. Campaigns built around a deep understanding of a company's most valuable and engaged customers, rather than purely measuring success by elevating clicks and open rates, could be a way to sidestep the Law of Shitty Clickthroughs and ensure the most reward for each marketing dollar spent.

10 Zimmermann's Law: How free products can build rich businesses

In 2013, email encryption pioneer Phil Zimmermann stated that the natural trajectory of technology today was to move in the direction of making surveillance easier, with computers' estimated ability to track users doubling – in a wry nod to Moore's Law – every 2 years.

Data is frequently said to be the most valuable commodity in the world, with some of the most successful companies in recent history capitalizing on it – often by offering free services in return for user data.

HOW DATA COLLECTION MADE FACEBOOK ONE OF THE WORLD'S MOST VALUABLE COMPANIES

Google was the first company that found massive success providing a free service in exchange for valuable user data, but Facebook took that formula and applied it to keep tabs on user behavior on its platform.

Facebook's success did not lie in being a first mover, but in learning from the mistakes its competitors were making.

Years before Facebook was launched, social networking began with SixDegrees.com. The website's founder described it as “placing your Rolodex in a central location.”

SixDegrees was followed by companies like LiveJournal, Friendster, and MySpace. At its peak in 2006, MySpace had 100M users. Facebook that year recorded only 12M users.

However, within a decade of its launch, Facebook grew to have over 1.2B monthly active users. Today, it boasts over 2.9B.

Facebook does not charge a single penny for its services. Users pay in kind through their usage data – their likes, dislikes, check-ins – all of which helps Facebook profile its users to fuel its targeted advertising business.

Over the years, Facebook’s data collection ability has steadily grown in power and resolution – as Zimmermann may have predicted – fueling the growth of the company’s ad platform and its increasing revenue.

The company’s mounting ability to extract data and understand its users has gone hand in hand with the increasing data collection potential of smartphones. The company’s founder and chief executive, Mark Zuckerberg, declared in 2014 that Facebook was “a mobile company now.” By Q2’19, nearly 94% of Facebook’s advertising revenue came from mobile ads.

Facebook achieved the counterintuitive feat of becoming one of the most valuable companies in the world by offering a free service.

However, this model has come under fire in recent years, especially in the wake of the Cambridge Analytica scandal in 2018. This was a costly incident for Facebook: The FTC fined Facebook \$5B for mishandling the data of 87M users.

Since then, data privacy laws like the EU’s General Data Protection Regulation (GDPR) have restricted the collection, use, sharing, and sale of consumer data. While tech businesses like Facebook don’t seem to be abandoning their business model anytime soon, they’ve realized that exploiting Zimmerman’s Law without trying to protect their users can cost them steep fines, damage their reputation, and drive customers away.

SNOOPING TOASTERS OFFER COMPANIES A NEW STREAM OF REVENUE

Smart home devices like speakers and security devices have become a part of many people's homes, with sales of these devices on the rise.

Over the last decade, we've seen physical, internet-connected technology crop up in virtually every corner of the household. There are hubs (Amazon Echo, Google Home), thermostats (Nest), video doorbells (Ring), security cameras (Nest, Netgear), scales (Withings), and more.

These devices can track and store sensitive data, such as how many people live in a household or your buying history. This data improves the system's responses to user commands and questions, in addition to helping associated services make relevant suggestions.

For example, if a user asks an Amazon Alexa about the symptoms associated with pregnancy, that data can be fed into other Amazon systems to show them suggestions of diapers, baby formula, and other baby products when shopping on Amazon.

These data-collecting devices offer convenience and personalization to users, but at the cost of reduced privacy — a trade-off that many consumers seem willing to make.

While Amazon does not reveal detailed sales numbers for Alexa smart home devices, it announced in 2020 that there were "hundreds of millions" of Alexa-enabled devices worldwide. This included third-party devices like Facebook Portal that had the Alexa smart assistant built in.

Many of these products explicitly or implicitly surveil their users. Security cameras and baby monitors do so by definition. Amazon's and Google's smart speakers need to always be listening for the command that triggers them to start working. But even products like smart fridges and smart TVs have been known to record their users or collect their watching habits and sell the information to third parties.

None of this should be a surprise, according to Zimmermann's Law. It is the natural trajectory of business to seek out new ways to drive revenue from products like microwaves, televisions, refrigerators, and speakers. And now that microwaves and TVs can effectively operate as mini-computers, it feels inevitable that manufacturers would look to collect potentially valuable data – whether for resale, for product optimization, or to bring down the sticker price of the device.

The topic of surveillance in regard to the IoT space is a controversial one. A big fear is security.

There have been several high-profile incidents of security cameras and baby monitors with lax security being hacked. New regulations and better security standards may mitigate some of these privacy risks while retaining the benefits, but it remains to be seen if companies will face a backlash from consumers before then.

One thing that seems sure is that the prevalence of these connected devices – and their data-collecting reach – will only increase.

TAKEAWAY

As technology improves, it's likely that data collection will increase as well. This may help lower consumer prices, enable more personalized services, and boost convenience, but it also raises fears around privacy that could give consumers pause.

Whether or not these kinds of benefits can balance out the negative possibilities will likely continue to define the conversation around tech and surveillance — especially if Zimmermann's Law holds to be accurate in its predictions about the natural trajectory of technology.

11 Pareto Principle: Why startups can raise capital even though most will eventually fail

In the late 19th century, Italian economist Vilfredo Pareto recorded his observation that just 20% of the population in the country owned about 80% of all the land.

After doing a number of surveys abroad, Pareto was surprised to find that the same was true in a number of other countries.

What Pareto probably didn't predict is the wide extent of the contexts to which this rough 80/20 distribution can be applied.

Since Pareto coined his rule of thumb, the 80/20 split has been used to provide insights into everything from optimization efforts in engineering control theory to the overwhelming influence of stars in baseball (approximately 15% of the players in the MLB are responsible for 85% of the total wins).

In 2002, Microsoft found that the Pareto Principle even applied to quality assurance and software development: fixing the top 20% of reported bugs handled 80% of the crashes and errors in a given system.

HOW THE PARETO PRINCIPLE GUIDES UNION SQUARE VENTURES' STRATEGY

For Union Square Ventures (USV) co-founder Fred Wilson, his firm's entire allocation strategy essentially hinges on the Pareto Principle.

"Every really good venture fund I have been involved in or have witnessed has had one or more investments that paid off so large that one deal single handedly returned the entire fund."

– Fred Wilson, co-founder of USV

At USV, each fund holds about \$200M and invests in around 20 to 25 companies. To get the 3x return that Wilson says venture firms should be shooting for, USV's focus is on making just 2 or 3 of each fund's companies into what it calls "high impact" – but the ultimate goal is that at least 1 or 2 of them exit "at a billion dollars or more."

Since 2011, USV has had [13 billion-dollar exits from its portfolio](#).

There are a few different tactics that USV uses to try and encourage these kinds of home runs.

First, cut your losses. "In our 2004 fund, we invested a total of \$50M out of \$120M of total investment in our nine losers. That wasn't so good," Wilson has written. "We could have, and should have, recognized our bad investments earlier and cut them off."

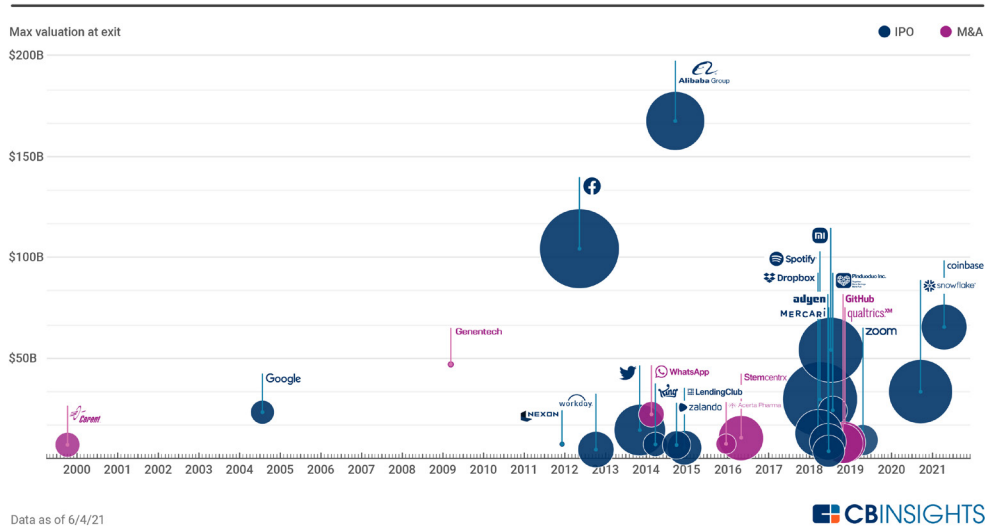
Second, as you start to discern the winners of a fund, funnel the majority of your follow-on money and operational expertise to them: "A few really good companies can carry a fund to the moon. You must make sure you can get a disproportionate amount of your time and money invested in those great investments."

This doesn't mean, however, that the optimal venture strategy is to simply swing for the fences every time. For Wilson, the real goal is to identify a wide swath of "great investments" and nurture them over a period of several years, hoping that a half-dozen or so turn into companies with "home run potential." After about 10 years, a few could lead to billion-dollar exits.

"There are hitters in baseball, the best hitters in fact, that hit balls out of the park when they are just trying to make good contact," Wilson writes. "That's how you have to do it in the venture business."

25 OF THE BEST VC-BACKED EXITS OF ALL TIME

Based on ratio of exit valuation to total pre-exit funding. Bubble size represents total funding at exit.



HOW ONE INVESTMENT MADE THE FUND FOR SEQUOIA

Sequoia's investment in WhatsApp is a great illustration of the Pareto Principle in action.

WhatsApp, which was founded by Brian Acton and Jan Koum in 2009, was not the kind of company many investors would have necessarily expected to become a fund-returning company.

The startups' founders were iconoclasts who rejected the idea of advertising or making money through degrading the user experience. Nor were they eager to work with venture capitalists: Koum and Acton insisted on working with, at maximum, one venture capital firm.

Fortunately for the venture firm, the founders chose Sequoia.

Sequoia would go on to put about [\\$60M in the company](#) through 2 rounds of funding to earn an 18% stake in the business.

Then, in 2014, Facebook acquired WhatsApp in a \$22B deal, which remains one of the largest private acquisitions of a VC-backed company ever.

Sequoia's entire Venture XI Fund, which raised \$387M from about 40 limited partners in 2003, returned \$3.6B in gains over the 11 years it had been open.

By the time WhatsApp was sold, Sequoia's stake in the company was worth around \$3B – equivalent to just over 80% of the Venture XI Fund's returns, and a 50x return on its investment in the company.

TAKEAWAY

Venture capital returns – where about [80% of the returns do tend to come from 20% of the companies](#) – present one of the most prominent examples of the Pareto Principle in action in the tech world.

In VC, most investments fail, but entire funds are often made through just a few home run deals.

If you aren't already a client, [sign up for a free trial](#) to learn more about our platform.