

---

## Subscriber's Daily Update

Tuesday, June 30, 2020

India Bans Chinese Apps, The App Store Firewall, Reddit and The Donald

---

# Apple and Facebook

Monday, June 29, 2020

When it comes to the big four consumer tech companies — Microsoft's decision to **close its retail stores** was the culmination of a step-back from the consumer space **five years in the making** — Google and Amazon have always had moats that were easier to understand. Google has a huge advantage in data and infrastructure, augmented by its control of consumer touch points (by owning Android outright, and paying heavily for default placement on other platforms); Amazon has a huge advantage in infrastructure and data, augmented by its control of the number one consumer touch point for shopping (the Amazon search box). Building a competitor for either feels daunting at best, impossible at worst.

## Apple's Moat

Apple, though, faced questions about the sustainability of its business for years, even after the runaway success of the iPhone; one of the topics that helped Stratechery gain traction in its first year was arguing that the iPhone was actually not about to be disrupted as so many — including **Professor Clayton Christensen himself** — were sure was going to happen. I explained why Apple's approach was sustainable in 2014's **Best**:

Moreover, integrated solutions will just about always be superior when it comes to the user experience: if you make the whole thing, you can ensure everything works well together, avoiding the inevitable rough spots and lack of optimization that comes with standards and interconnects. The key, though, is that this integration and experience be valued by the user. That is why — and this was the crux of my criticism of Christensen's development of the theory — the user experience angle only matters when the buyer of a product is also the user. Users care about the user experience (surprise), but an isolated buyer — as is the case for most business-to-business products, and all of Christensen's examples — does not. I believe this was the root of Christensen's blind spot about Apple, which persists. From **an interview with Henry Blodget** a month ago:

*You can predict with perfect certainty that if Apple is having that extraordinary experience, the people with modularity are striving. You can predict that they are motivated to figure out how to emulate what they are offering, but with modularity. And so ultimately, unless there is no ceiling, at some point Apple hits the ceiling. So their options are hopefully they can come up with another product category or something that is proprietary because they really are good at developing products that are proprietary.*

*Most companies have that insight into closed operating systems once, they hit the ceiling, and then they crash.*

That's the thing though: the quality of a user experience *has no ceiling*. As nearly every other consumer industry has shown, as long as there is a clear delineation between the top-of-the-line and everything else, some segment of the user base will pay a premium for the best. That's the key to Apple's future: they don't need completely new products every other year (or half-decade); they just need to keep creating the best stuff in their categories. Easy, right?

It's not easy, of course, and yet when it comes to hardware in particular, Apple's lead is greater than it ever was, thanks in large part to its **superior systems-on-a-chip**; the headline news from WWDC was that Apple is extending that particular advantage **to the Mac**.

Of course having the best device is not enough either — this was the other reason why Apple was, according to many, eternally doomed; here is Blodget again in **Business Insider**, in 2013:<sup>1</sup>

If smartphones and tablets were not a platform — if the only thing that mattered to the value of the product and a customer's purchase decision was the gadget itself — then Apple's loss of market share would not make a difference. Apple zealots would be correct when they smugly assert that what matters is Apple's "profit share" not "market share."

But smartphones and tablets are a platform. Third-party companies are building apps and services to run on smartphone and tablet platforms. These apps and services, in turn, are making the platforms more valuable. Consumers are standardizing their lives around the apps and services that run on smartphone and tablet platforms. Because of these "network effects," in platform markets, dominant market share is huge competitive advantage. In platform markets, as the often-hated but always insanely powerful Microsoft demonstrated for decades in the PC market, the vast majority of the power and profits eventually accrue to the market-share leader.

In fact, it turned out that Apple's prioritization of the user experience wasn't simply a moat, but also a point of leverage with developers, who need Apple much more than Apple needs them. Thus the second big story over the last two weeks, which has been Apple's **App Store policies**: it appears that Apple has significantly tightened its unwritten rules over the last year as the company seeks to increase its Services revenue. Developers, from **the smallest** to **the largest**, have no choice but to accede to the iPhone maker's demands because Apple has combined the loyalty of the most valuable users with App Review, an unavoidable gatekeeper in terms of getting apps onto those users' iPhones.

## The Bill Gates Line

Facebook, meanwhile, is often thought of as being the opposite of Apple: Apple sells products, and Facebook sells advertising. Apple minimizes data collection, and Facebook maximizes it. Apple is a platform, and Facebook is **just an app**.

What both share, though, is a sort of eternal skepticism from Silicon Valley in particular. Facebook, for its part, has been doubted from its formation to its decision to decline Yahoo's acquisition offer to its

post-IPO stock price slump; every hot new social app, from Twitter to Snapchat to TikTok, is framed as the service that will finally doom Facebook to being the next MySpace.

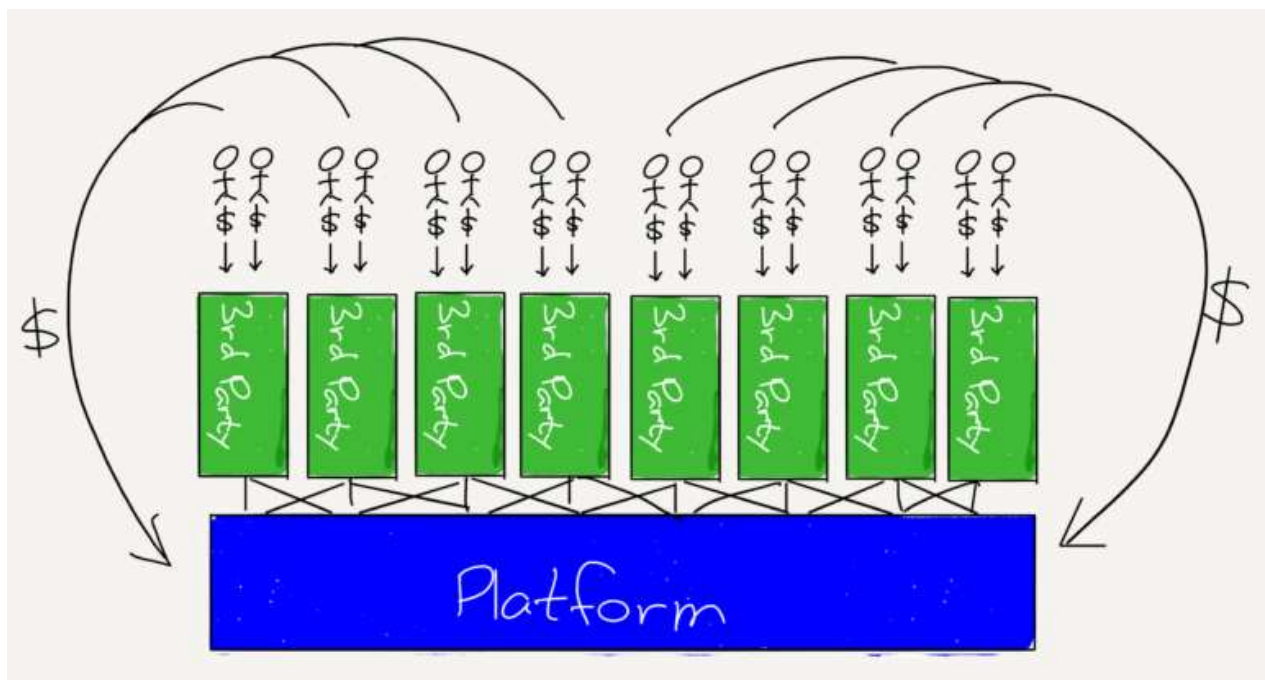
The truth, though, is that in many respects Facebook is more of a platform than Apple is. In 2018 I wrote about [The Bill Gates Line](#), which was actually coined as criticism of Facebook:

Over the last few weeks I have been exploring what differences there are between platforms and aggregators, and was reminded of this anecdote from Chamath Palihapitiya in [an interview with Semil Shah](#):

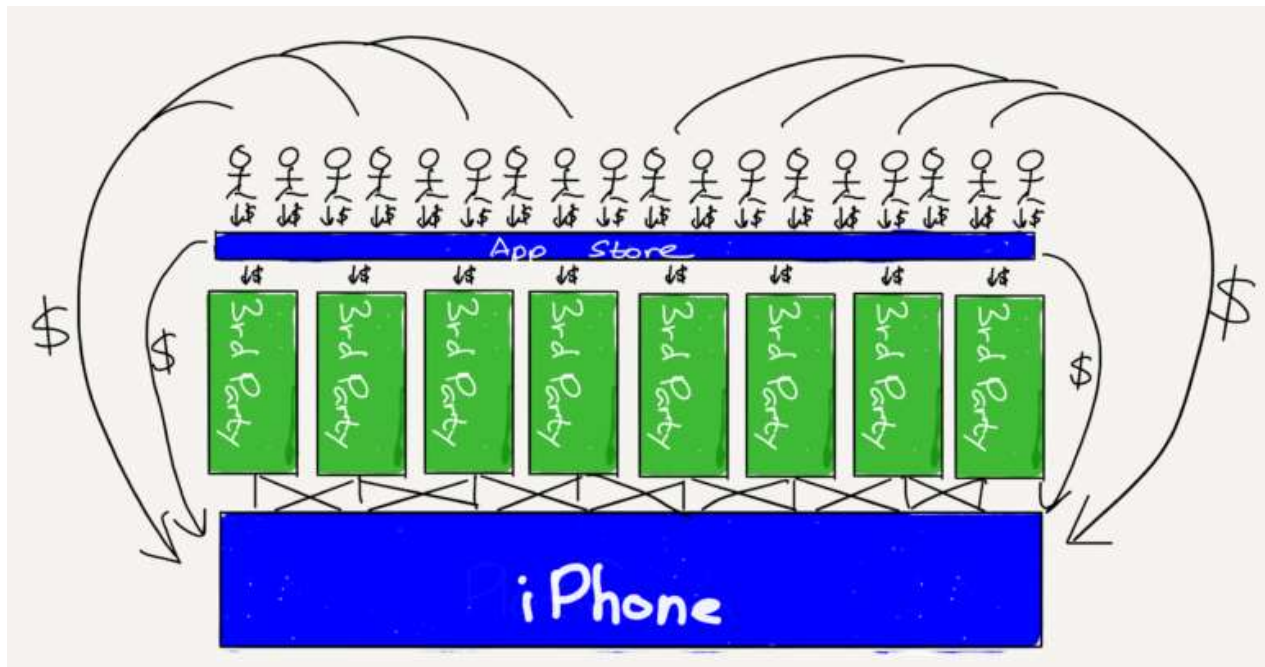
*Semil Shah: Do you see any similarities from your time at Facebook with Facebook platform and connect, and how Uber may supercharge their platform?*

*Chamath Palihapitiya: Neither of them are platforms. They're both kind of like these comical endeavors that do you as an Nth priority. I was in charge of Facebook Platform. We trumpeted it out like it was some hot shit big deal. And I remember when we raised money from Bill Gates, 3 or 4 months after — like our funding history was \$5M, \$83 M, \$500M, and then \$15B. When that 15B happened a few months after Facebook Platform and Gates said something along the lines of, "That's a crock of shit. This isn't a platform. A platform is when the economic value of everybody that uses it, exceeds the value of the company that creates it. Then it's a platform."*

By this measure Windows was indeed the ultimate platform — the company used to brag about only capturing a minority of the total value of the Windows ecosystem — and the operating system's clear successors are Amazon Web Services and Microsoft's own Azure Cloud Services. In all three cases there are strong and durable businesses to be built on top.



Once a platform dips under the Bill Gates Line, though, the long-term potential of a business built on a “platform” starts to decline. Apple’s App Store, for example, has all of the trappings of a platform, but Apple quite clearly captures the vast majority of the overall ecosystem, both because of the profitability of the iPhone and also because of [its control of App Store economics](#); the paucity of strong and durable businesses on the App Store is a natural outgrowth of that.



Note that Apple's ability to control the economics of its developers comes from intermediating the relationship of those developers with customers.

What is missing in this story is how exactly all of those developers made money on the App Store. Yes, without question, a big part of it was the iPhone's explosive growth and how the App Store made it easy and safe to install apps. Another very big part of it, though, was Facebook.

## Facebook's Platform

Facebook's early stumbles on mobile are well-documented: **the company bet on web-based apps** that just didn't work very well; the company completely rewrote its iOS app even as it was going public, which meant it had a stagnating app at the exact time mobile was exploding, threatening the desktop advertising product and platform that were the basis of the company's S-1.

The re-write turned out to be not just a company-saving move — the native mobile app had the exact same user-facing features as the web-centric one, with the rather important detail that it actually worked — but in fact an industry-transformational one: one of the first new products enabled by the company's new app were app install ads. **From TechCrunch in 2012:**

Facebook is making a big bet on the app economy, and wants to be the top source of discovery outside of the app stores. The mobile app install ads let developers buy tiles that promote their apps in the Facebook mobile news feed. When tapped, these instantly open the Apple App Store or Google Play market where users can download apps.

The ads are working already. One client TinyCo saw a 50% higher click through rate and higher conversion rates compared to other install channels. Facebook's ads also brought in more engaged users. Ads tech startup Nanigans clients attained 8-10X more reach than traditional mobile ad buys when it purchased Facebook mobile app install ads. AdParlor racked up a consistent 1-2% click through rate.

Facebook's App Install product quickly became the most important channel for acquiring users, particularly for games that monetized with Apple's in-app purchase API: the combination of Facebook data with developer's sophisticated understanding of expected value per app install led to an explosion in App Store revenue. And yet, even this was seen as a reason to doubt Facebook; in 2015 I wrote about **a prominent venture capitalist's Facebook skepticism**:

Much of the criticism of app install ads rests on obsolete assumptions that view apps as fun baubles instead of the dominant interaction layer between companies and consumers. If you start with the premise that apps are more important than web pages or any other form of interaction when it comes to connecting with consumers, being the dominant channel for app installs seems downright safe.

Surely, though, at least some portion of Facebook's revenue must come from app-install ads for games, no? Absolutely! But even that is less of a danger than critics think for three reasons:

- It's a mistake to assume that just because venture-backed companies have a tendency to be profligate with their money that app install ad spending is little more than "spray-and-pray". In fact, app install ads are not just direct marketing, which is much easier to track, Facebook app install ads in particular are one of the most data rich ad formats around...
- That said, were I a venture capitalist like Gurley, I would be gun-shy about mobile gaming; there are a whole host of examples of one-shot wonders that attract funding or even IPO on a hit game only to struggle to recreate their initial success. I think it's a really hard area to invest in. Facebook, though, doesn't need to care if a particular gaming company succeeds or fails, because they aren't exposed to any one gaming company: they have exposure to the industry as a whole. And on that note:
- Mobile gaming revenue is not a flash-in-the-pan. **According to Newzoo**, a video game research firm, global mobile game revenues is expected to surpass console revenue this year. Interestingly — and perhaps this makes the space a bit of a blind-spot for U.S.-based observers — console spending will still dominate in North America (\$11.1 billion to \$7.2 billion); it's in the rest of the world — particularly in Asia — where mobile is absolutely dominant.

...Ultimately, if there is a bubble, everyone will suffer, including Facebook. But big picture I continue to consider the company the most underrated in the Valley (which is kind of amazing). Facebook has barely scratched the surface of their monetization capabilities, brand advertising has yet to migrate from TV, Instagram still isn't monetized at all, and to cap it off Facebook usage is still increasing. This is a juggernaut that, if there is a downturn, is more likely to be the exception to industry doldrums than the rule.

While this excerpt is about mobile games, the analysis applies to a whole host of industries that have grown up on Facebook, **like direct-to-consumer e-commerce companies**:

- Facebook's targeting is a particularly potent combination for companies that convert on-line. Indeed, the biggest mistake I've made in evaluating the company is over-estimating the potential of brand advertising and under-appreciating just how large the direct response opportunity was.
- Relatedly, the direct response opportunity is so large because Facebook has created the conditions for new direct response-based companies to be created. For apps, this opportunity was created in conjunction with Apple and Google (Android); for e-commerce this opportunity was created **in conjunction with Shopify**. What both examples have in common is that Facebook's advertising was a critical factor in creating new businesses that were unique to the Internet.

- It's not just mobile gaming that is more than a flash-in-the-pain: **the transformative impact of the Internet is only starting to be felt**, which is to say that the long run will be less about traditional companies adopting digital than it will be about their entire way of doing business being rendered obsolete.

One of my favorite examples are CPG companies.

## Facebook's Anti-Fragility

In 2016's **TV Advertising's Surprising Strength — And Inevitable Fall** I observed:

The very institution of television advertising is intertwined with the kinds of advertisers that use it the most, the products they sell, and the way they are bought-and-sold. And what should be terrifying to television executives is that all of those pieces that make television advertising the gold mine that it has been are under the exact same threat that TV watching itself is: the threat of the Internet...

CPG is the perfect example: building a "house of brands" allows a company like Procter & Gamble to target demographic groups even as they leverage scale to invest in R&D, bring down the cost of products, and most importantly, dominate the distribution channel (i.e. retail shelf space). Said retailers, meanwhile, are huge in their own right, not only so they can match their massive suppliers at the bargaining table but also so they can scale logistics, inventory management, store development, etc. Automobile companies, meanwhile, are not unlike CPG companies: they operate a "house of brands" to serve different demographics while benefitting from scale in production and distribution; the primary difference is that they make money through one large purchase instead of over many smaller purchases over time.

Similar principles apply to the other companies on this list: all are looking to reach as many consumers as possible with blunt targeting at best, all benefit from scale, and all are looking to earn significant lifetime value from consumers. And, along those lines, all can afford the expense of TV. In fact, the top 200 advertisers in the U.S. love TV so much that they make up 80% of television advertising, despite accounting for only 51% of total advertising (and 41% of digital).

This is a very different picture from Facebook, where **as of Q1 2019** the top 100 advertisers made up less than 20% of the company's ad revenue; most of the \$69.7 billion the company brought in last year came from its long tail of 8 million advertisers.

This focus on the long-tail, which is only possible because of Facebook's fully automated ad-buying system, has turned out to be a tremendous asset during the coronavirus slow-down. I explained after **Facebook's Q1 2020 earnings**:

That first bit gets at the other thing the Wall Street Journal article got wrong: it is not simply that direct response stayed strong while brand advertising declined, but rather that Facebook actually received more direct response advertising because brand advertising declined...

Notice, though, what happens in a situation like the coronavirus crisis, where a segment of advertisers competing for limited inventory stop buying ads: the mobile gaming company doesn't reduce their budget — to do so would be to kill the company! — but in fact ends up getting more efficient spend.

Suddenly the clearing price for the auction to show those app install ads is \$0.75 per app install; now the mobile gaming company is getting 26,667 app installs for its \$20,000 spend, which results in an expected profit of \$6,667.

This does, obviously, entail downside for Facebook — that extra ~\$6,000 in profit is out of Facebook's pocket — but at the same time the loss is capped because not only is the mobile gaming company not reducing its spend, it is in fact incentivized to increase its spend given the reduced competition and thus increased profitability for its ads. And, of course, as that opportunity is seized on by more and more companies, Facebook's profits, which in the end are gated by the amount of inventory it has, not only return to normal but arguably have more upside, given that usage of the platform is increasing.

This explains why the news about large CPG companies boycotting Facebook is, from a financial perspective, simply not a big deal. Unilever's \$11.8 million in U.S. ad spend, to take one example, is replaced with the same automated efficiency that Facebook's timeline ensures you never run out of content. Moreover, while Facebook loses some top-line revenue — in an auction-based system, less demand corresponds to lower prices — the companies that are the most likely to take advantage of those lower prices are those that would not exist without Facebook, like the direct-to-consumer companies trying to steal customers from massive conglomerates like Unilever.

In this way Facebook has a degree of anti-fragility that even Google lacks: so much of its business comes from the long tail of Internet-native companies that are built around Facebook from first principles, that any disruption to traditional advertisers — like the coronavirus crisis or the current boycotts — actually serves to strengthen the Facebook ecosystem at the expense of the TV-centric ecosystem of which these CPG companies are a part.

## The Apple Vulnerability

Facebook, though, has a vulnerability: Apple. From [AdAge](#):

Apple announced new privacy changes to its upcoming iOS 14 software that will significantly hinder how media buyers and brands target, measure and find consumers. One change will make it harder for apps to track iOS users across different apps and websites. Another will make attribution — determining which tactics contribute to sales or conversions — harder for marketers.

The changes, announced Monday at Apple's Worldwide Developers Conference, apply to the company's Identifier for Advertisers (IDFA), which assigns a unique number to a user's mobile device. Advertisers have access to the feature and use it in areas including ad targeting, building lookalike audiences, attribution and encouraging consumers to download apps.

IDFA is shared with app makers and advertisers by default, but that will change once iOS 14 rolls out this fall. Then, users must give explicit permission through a popup for app publishers to track them across different apps and websites, or to share that information with third parties.

Facebook was the king of the IDFA (and the Google Advertising ID equivalent on Android): it was the linchpin around which its app install business in particular was built. The company could understand when a user spent a certain amount in a game, for example, look for users that were similar, and then

display an app install ad for that game, and measure how effective it was. In fact, over the last few years, Facebook has simply asked advertisers to specify what return on ad spend they are hoping to achieve, and Facebook does all of the work of figuring out how many ads to display to which users — the entire process is automated.

This part of the business is going to change a lot. Apple was quite clever in their approach: instead of killing the IDFA, which could be construed as anti-competitive, particularly given Apple's expanding app install ad business (which is [expanding beyond App Store search ads](#)), Apple is simply asking users if they would like to be tracked, and letting them render the IDFA useless. Notably, Facebook has declined to even show app install advertisements to the 30% of U.S. iPhone users that turned off their IDFA of their own accord — and now it is opt-in, instead of opt-out.

Still, I wouldn't count Facebook out: to the extent the company is hurt, it seems likely that the universe of 3rd-party ad tech companies that lack Facebook's direct connection with users, both in terms of data collection and ad display, will be in far worse shape, and it is not as if the digital ecosystem — and its associated advertising — is going to disappear. Indeed, [much like GDPR](#), the safe bet is the company with the wherewithal to make lemons out of lemonade. Notably, Apple's alternative for app install ad campaigns, [SKAdNetwork](#), is so limited that there is likely to be tremendous value in whatever company can create the exact sort of automated campaign creation that Facebook is already offering.

It's also worth noting that Apple's crackdown on web cookies has also helped Facebook, as [I explained last month](#); by making it more difficult for 3rd-party payment providers to offer a seamless experience, Apple is opening the door to Facebook taking over payments for direct-to-consumer companies:

Facebook Shops is a perfect example: it is going to succeed because it is good for Shopify's merchants, but the reason it is good for Shopify's merchants is because Facebook and Apple effectively teamed up to make it impossible for Shopify to fix the payment problem on their own.

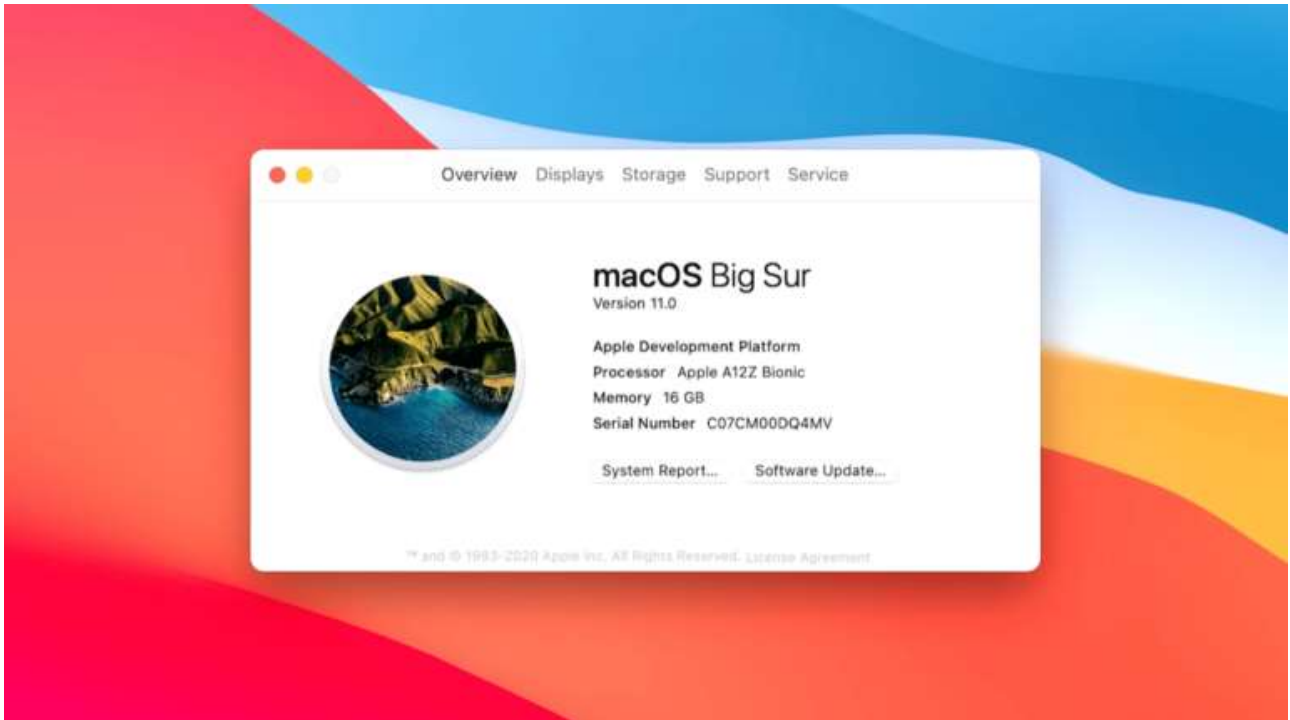
This is what makes the Apple-Facebook dynamic so fascinating: Facebook's biggest opportunities come from filling in the holes in Apple's platform proposition, even as Apple seems opposed to Facebook at every turn.

## Shared Risk

What is particularly notable is how conflict between the two companies threatens their greatest assets.

Start with Apple: while its battle against cookies and effective obliteration of the IDFA are from one perspective deeply rooted in its focus on users, it is impossible to ignore the company's focus on Services revenue. Making the web less useful makes apps more useful, from which Apple can take its share; similarly, it is notable that Apple is expanding its own app install product even as it is knee-capping the industry's. The question is if these attempts to maximize services revenue are in service of the user experience, or Apple's bottom line; the company should take care to remember that the latter follows the former.

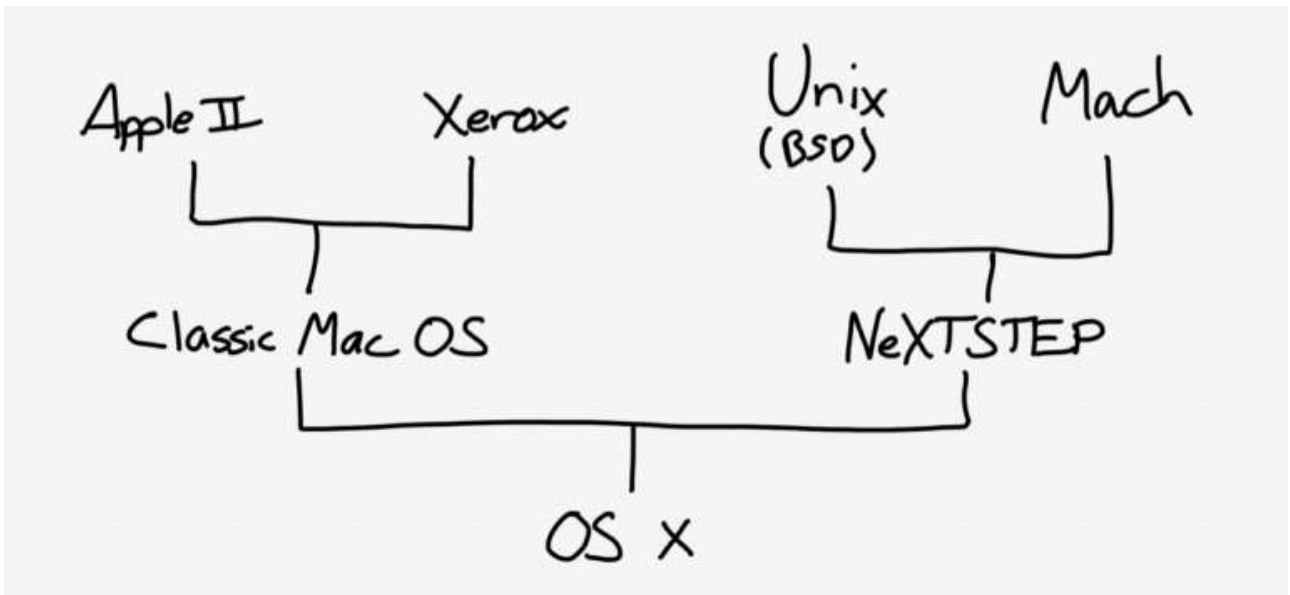
Yesterday, 18 years later, OS X finally reached its own end of the road: the next version of macOS is not 10.16, but 11.0.



There was no funeral.

## The OS X Family Tree

OS X has one of the most fascinating family trees in technology; to understand its significance requires understanding each of its forebearers.



**Unix:** Unix does refer to a specific operating system that originated in AT&T's Bell Labs (the copyrights of which are owned by Novell), but thanks to a settlement with the U.S. government (that was **widely criticized for going easy on the telecoms giant**), Unix was widely-licensed to universities in particular. One of the most popular variants that resulted was the Berkeley Software Distribution (BSD), developed at the University of California, Berkeley.

Facebook, meanwhile, already sees promise in taking business from its best partners, as seen in the announcement of Facebook Shops; there may be a similar temptation when it comes to IDFA, given that the IDFV — **Identifier for Vendor**, which allows a vendor to get the device ID from its own apps — is still available. Might Facebook consider shifting its business model from being an advertising platform for other apps into a WeChat-like publishing model for the most popular games and services?<sup>2</sup> The company should also take care: its service of the long tail has not only made it more of a Bill-Gates platform than Apple, but is also the foundation of the company's strength in crisis.

Regardless, what seems clearer than ever is that it is these two companies, Apple and Facebook, that are driving the industry. That their approaches are so different is in fact why they are the pairing that matters most.

- 
1. I previously quoted this piece in **Beachheads and Obstacles**. [↩]
  2. I do expect a lot of consolidation in the mobile gaming industry, particularly amongst hyper-casual game publishers, for exactly this reason [↩]
- 

## The End of OS X

Tuesday, June 23, 2020

On May 6, 2002, Steve Jobs opened WWDC with a funeral for Classic Mac OS:

Loading...

What all of the variations of Unix had in common was the Unix Philosophy; the **Bell System Technical Journal** explained in 1978:

A number of maxims have gained currency among the builders and users of the Unix system to explain and promote its characteristic style:

1. Make each program do one thing well. To do a new job, build afresh rather than complicate old programs by adding new “features”.
2. Expect the output of every program to become the input to another, as yet unknown, program. Don't clutter output with extraneous information. Avoid stringently columnar or binary input formats. Don't insist on interactive input.
3. Design and build software, even operating systems, to be tried early, ideally within weeks. Don't hesitate to throw away the clumsy parts and rebuild them.
4. Use tools in preference to unskilled help to lighten a programming task, even if you have to detour to build the tools and expect to throw some of them out after you've finished using them.

[...]

The Unix operating system, the C programming language, and the many tools and techniques developed in this environment are finding extensive use within the Bell System and at universities, government laboratories, and other commercial installations. The style of computing encouraged by this environment is influencing a new generation of programmers and system designers. This, perhaps, is the most exciting part of the Unix story, for the increased productivity fostered by a friendly environment and quality tools is essential to meet every-increasing demands for software.

Today you can still run nearly any Unix program on macOS, but particularly with some of the security changes made in Catalina, you are liable to run into permissions issues, particularly when it comes to seamlessly linking programs together.

**Mach:** Mach was a microkernel developed at Carnegie Mellon University; the concept of a microkernel is to run the smallest amount of software necessary for the core functionality of an operating system in the most privileged mode, and put all other functionality into less privileged modes. OS X doesn't have a true microkernel — the BSD subsystem runs in the same privileged mode, for performance reasons — but the modular structure of a microkernel-type design makes it easier to port to different processor architectures, or remove operating system functionality that is not needed for different types of devices (there is, of course, lots of other work that goes into a porting a modern operating system; this is a dramatic simplification).

More generally, the spirit of a microkernel — a small centralized piece of software passing messages between different components — is how modern computers, particularly mobile devices, are architected: multiple specialized chips doing discrete tasks under the direction of an operating system organizing it all.

**Xerox:** The story of Steve Jobs' visiting Xerox is as mistaken as it is well-known; the Xerox Alto and its groundbreaking mouse-driven graphical user interface was well-known around Silicon Valley, thanks to the thousands of demos the Palo Alto Research Center (PARC) did and the papers it had published.

PARC's problem is that Xerox cared more about making money from copy machines than in figuring out how to bring the Alto to market.

That doesn't change just how much of an inspiration the Alto was to Jobs in particular: after the visit he pushed the Lisa computer to have a graphical user interface, and it was why he took over the Macintosh project, determined to make an inexpensive computer that was far easier to use than anything that had come before it.

**Apple:** The Macintosh was not the first Apple computer: that was the Apple I, and then the iconic Apple II. What made the Apple II unique was its explicit focus on consumers, not businesses; interestingly, what made the Apple II successful was VisiCalc, the first spreadsheet application, which is to say that the Apple II sold primarily to businesses. Still, the truth is that Apple has been a consumer company from the very beginning.

This is why the Mac is best thought of as the child of Apple and Xerox: Apple understood consumers and wanted to sell products to them, and Xerox provided the inspiration for what those products should look like.

It was NeXTSTEP, meanwhile, that was the child of Unix and Mach: an extremely modular design, from its own architecture to its focus on object-oriented programming and its inclusion of different "kits" that were easy to fit together to create new programs.

And so we arrive at OS X, the child of the classic Macintosh OS and NeXTSTEP. The best way to think about OS X is that it took the consumer focus and interface paradigms of the Macintosh and layered them on top of NeXTSTEP's technology. In other words, the Unix side of the family was the defining feature of OS X.

## Return of the Mac

In 2005 Paul Graham wrote an essay entitled **Return of the Mac** explaining why it was that developers were returning to Apple for the first time since the 1980s:

All the best hackers I know are gradually switching to Macs. My friend Robert said his whole research group at MIT recently bought themselves Powerbooks. These guys are not the graphic designers and grandmas who were buying Macs at Apple's low point in the mid 1990s. They're about as hardcore OS hackers as you can get.

The reason, of course, is OS X. Powerbooks are beautifully designed and run FreeBSD. What more do you need to know?

Graham argued that hackers were a leading indicator, which is why he advised his dad to buy Apple stock:

If you want to know what ordinary people will be doing with computers in ten years, just walk around the CS department at a good university. Whatever they're doing, you'll be doing.

In the matter of “platforms” this tendency is even more pronounced, because novel software originates with great hackers, and they tend to write it first for whatever computer they personally use. And software sells hardware. Many if not most of the initial sales of the Apple II came from people who bought one to run VisiCalc. And why did Bricklin and Frankston write VisiCalc for the Apple II? Because they personally liked it. They could have chosen any machine to make into a star.

If you want to attract hackers to write software that will sell your hardware, you have to make it something that they themselves use. It’s not enough to make it “open.” It has to be open and good. And open and good is what Macs are again, finally.

What is interesting is that Graham’s stock call could not have been more prescient: Apple’s stock closed at \$5.15 on March 31, 2005, and \$358.87 yesterday;<sup>1</sup> the primary driver of that increase, though, was not the Mac, but rather the iPhone.

## The iOS Sibling

If one were to add iOS to the family tree I illustrated above, most would put it under Mac OS X; I think, though, iOS is best understood as another child of Classic Mac and NeXT, but this time the resemblance is to the Apple side of the family. Or to put it another way, while the Mac was the perfect machine for “hackers”, to use Graham’s term, the iPhone was one of the purest expressions of Apple’s focus on consumers.

The iPhone, as Steve Jobs declared at its unveiling in 2007, runs OS X, but it was certainly not *Mac OS X*: it ran the same XNU kernel, and most of the same subsystem (with some new additions to support things like cellular capability), but it had a completely new interface. That interface, notably, did not include a terminal; you could not run arbitrary Unix programs.<sup>2</sup> That new interface, though, was far more accessible to regular users.

What is more notable is that the iPhone gave up parts of the Unix Philosophy as well: applications all ran in individual sandboxes, which meant that they could not access the data of other applications or of the operating system. This was great for security, and is the primary reason why iOS doesn’t suffer from malware and apps that drag the entire system into a morass, but one certainly couldn’t “expect the output of every program to become the input to another”; until sharing extensions were added in iOS 8 programs couldn’t share data with each other at all, and even now it is tightly regulated.

At the same time, the App Store made principle one — “make each program do one thing well” — accessible to normal consumers. Whatever possible use case you could imagine for a computer that was always with you, well, “There’s an App for That”:

Loading...

Consumers didn't care that these apps couldn't talk to each other: they were simply happy they existed, and that they could download as many as they wanted without worrying about bad things happening to their phone — or to them. While sandboxing protected the operating system, the fact that every app was reviewed by Apple weeded out apps that didn't work, or worse, tried to scam end users.

This ended up being good for developers, at least from a business point-of-view: sure, the degree to which the iPhone was locked down grated on many, but Apple's approach created millions of new customers that never existed for the Mac; the fact it was closed and good was a benefit for everyone.

## macOS 11.0

What is striking about macOS 11.0 is the degree to which it feels more like a son of iOS than the sibling that Mac OS X was:

- macOS 11.0 runs on ARM, just like iOS; in fact the Developer Transition Kit that Apple is making available to developers has the same A12Z chip as the iPad Pro.
- macOS 11.0 has a user interface overhaul that not only appears to be heavily inspired by iOS, but also seems geared for touch.
- macOS 11.0 attempts to acquire developers not primarily by being open and good, but by being easy and good enough.

The seeds for this last point were planted last year with Catalyst, which made it easier to port iPad apps to the Mac; with macOS 11.0, at least the version which will run on ARM, Apple isn't even requiring a recompile: iOS apps will simply run on macOS 11.0, and they will be in the Mac App Store by default (developers can opt-out).

In this way Apple is using their most powerful point of leverage — all of those iPhone consumers, which compel developers to build apps for the iPhone, Apple's rules notwithstanding — to address what the company perceives as a weakness: the paucity of apps in the Mac App Store.

Is the lack of Mac App Store apps really a weakness, though? When I consider the apps that I use regularly on the Mac, a huge number of them are not available in the Mac App Store, not because the developers are protesting Apple's 30% cut of sales, but simply because **they would not work** given the limitations Apple puts on apps in the Mac App Store.

The primary limitation, notably, is the same sandboxing technology that made iOS so trustworthy; that trustworthiness has always come with a cost, which is the ability to build tools that do things that “lighten a task”, to use the words from the Unix Philosophy, even if the means to do so opens the door to more nefarious ends.

Fortunately macOS 11.0 preserves its NeXTSTEP heritage: non-Mac App Store apps are still allowed, for better (new use cases constrained only by imagination and permissions dialogs) and worse (access to other apps and your files). What is notable is that this was even a concern: Apple's recent moves on iOS, particularly around **requiring in-app purchase for SaaS apps**, feel like a drift towards Xerox, a company that was so obsessed with making money it ignored that it was giving demos of the future to its competitors; one wondered if the obsession would filter down to the Mac.

For now the answer is no, and that is a reason for optimism: an open platform on top of the tremendous hardware innovation being driven by the iPhone sounds amazing. Moreover, one can argue (hope?) it is a more reliable driver of future growth than squeezing every last penny out of the greenfield created by the iPhone. At a minimum, leaving open the possibility of entirely new things leaves far more future optionality than drawing the strings every more tightly as on iOS. OS X's legacy lives, for now.

*I wrote a follow-up to this article in **this Daily Update**.*

---

1. Yes, this incorporates Apple's 7:1 stock split [↩]
  2. Unless you jailbroke your phone [↩]
- 

## Apple, ARM, and Intel

Tuesday, June 16, 2020

Mark Gurman **at Bloomberg** is reporting that Apple will *finally* announce that the Mac is transitioning to ARM chips at next week's Worldwide Developer Conference (WWDC):

Apple Inc. is preparing to announce a shift to its own main processors in Mac computers, replacing chips from Intel Corp., as early as this month at its annual developer conference, according to people familiar with the plans. The company is holding WWDC the week of June 22. Unveiling the initiative, codenamed Kalamata, at the event would give outside developers time to adjust before new Macs roll out in 2021, the people said. Since the hardware transition is still months away, the timing of the announcement could change, they added, while asking not to be identified discussing private plans. The new processors will be based on the same technology used in Apple-designed iPhone and iPad chips. However, future Macs will still run the macOS operating system rather than the iOS software on mobile devices from the company.

I use the word “finally” a bit cheekily: while it feels like this transition has been rumored forever, until a couple of years ago I felt pretty confident it was *not* going to happen. Oh sure, the logic of Apple using its remarkable iPhone chips in Macs was obvious, even back in 2017 or so:

- Apple’s A-series chips had been competitive on single-core performance with Intel’s laptop chips for several years.
- Intel, **by integrating design and manufacturing**, earned very large profit margins on its chips; Apple could leverage TSMC for manufacturing and keep that margin for itself and its customers.
- Apple could, as they did with iOS, deeply integrate the operating system and the design of the chip itself to both maximize efficiency and performance and also bring new features and capabilities to market.

The problem, as I saw it, was why bother? Sure, the A-series was catching up on single-thread, but Intel was still far ahead on multi-core performance, and that was before you got to desktop machines where pure performance didn’t need to be tempered by battery life concerns. More importantly, the cost of switching was significant; I wrote in early 2018:

- First, Apple sold 260 million iOS devices over the last 12 months; that is a lot of devices over which to spread the fixed costs of a custom processor. During the same time period, meanwhile, the company only sold 19 million Macs; that’s a much smaller base over which to spread such an investment.
- Second, iOS was built on the ARM ISA from the beginning; once Apple began designing its own chips (instead of buying them off the shelf) there was absolutely nothing that changed from a developer perspective. That is not the case on the Mac: many applications would be fine with little more than a recompile, but high-performance applications written at lower levels of abstraction could need considerably more work (this is the challenge with emulation as well: the programs that are the most likely to need the most extensive rewrites are those that are least tolerant of the sort of performance slowdowns inherent in emulation).
- Third, the PC market is in the midst of its long decline. Is it really worth all of the effort and upheaval to move to a new architecture for a product that is fading in importance? Intel may be expensive and may be slow, but it is surely good enough for a product that represents the past, not the future.

However, the takeaway from the **Daily Update** where I wrote that was that I was changing my mind: ARM Macs felt inevitable, because of changes at both Apple and Intel.

## Apple and Intel

A year before that Daily Update, Apple held **a rather remarkable event** for five writers where the company seemed to admit it had neglected the Mac; from **TechCrunch**:

Does Apple care about the Mac anymore?

That question is basically the reason that we’re here in this room. Though Apple says that it was doing its best to address the needs of pro users, it obviously felt that the way the pro community was reacting to

its moves (or delays) was trending toward what it feels is a misconception about the future of the Mac.

“The Mac has an important, long future at Apple, that Apple cares deeply about the Mac, we have every intention to keep going and investing in the Mac,” says Schiller in his most focused pitch about whether Apple cares about the Mac any more, especially in the face of the success of the iPhone and iPad.

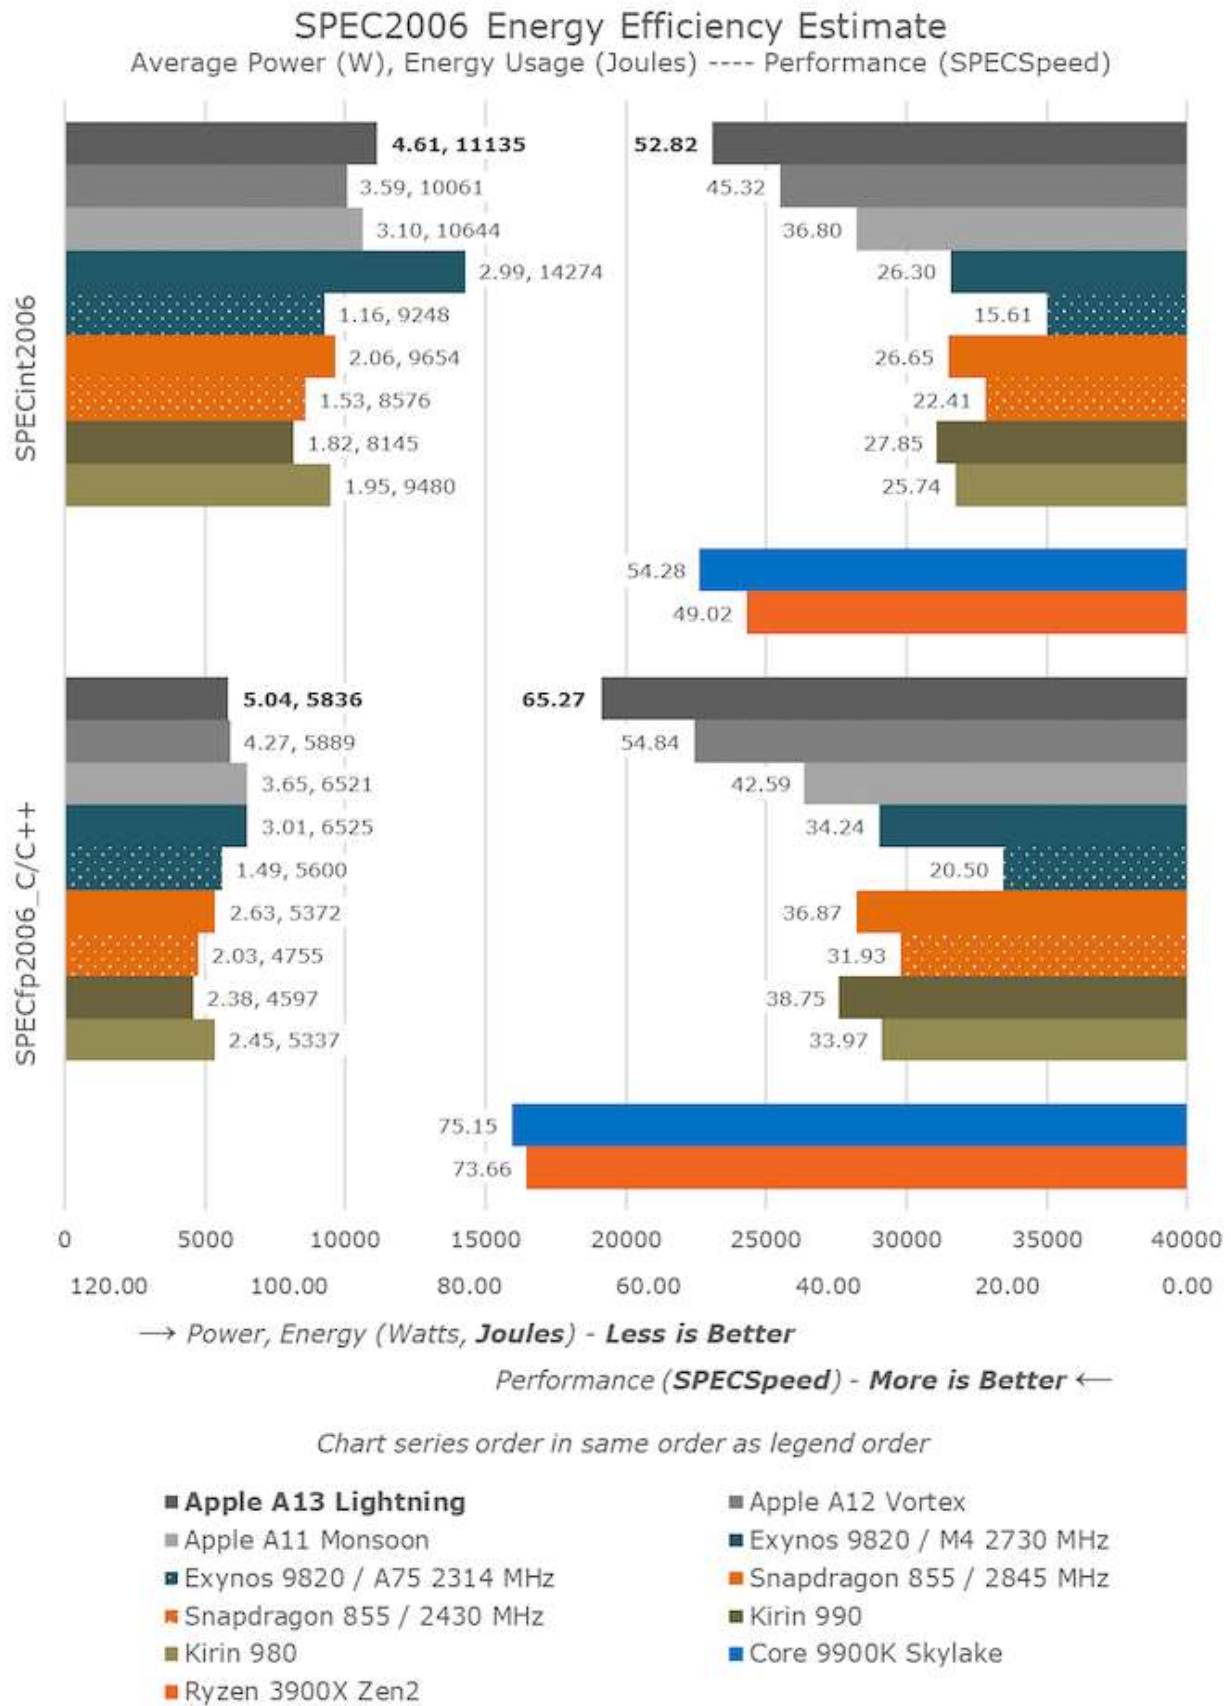
“And if we’ve had a pause in upgrades and updates on that, we’re sorry for that — what happened with the Mac Pro, and we’re going to come out with something great to replace it. And that’s our intention,” he says, in as clear a mea culpa as I can ever remember from Apple.

Yes, Schiller was talking about the Mac Pro, which is what the event was nominally about, but that wasn’t the only Mac long in the teeth, and the ones that had been updated, particularly the laptops, were years into the butterfly keyboard catastrophe; meanwhile there was a steady-stream of new iPhones and iPads with new industrial designs and those incredible chips.

Those seemingly neglected Macs, meanwhile, were stuck with Intel, and Apple saw the Intel roadmap that has only recently become apparent to the world: it has been a map to nowhere. In 2015 Intel started shipping 14nm processors in volume from fabs in Oregon, Arizona, and Ireland; chip makers usually build fabs once per node size, seeking to amortize the tremendous expense over the entire generation, before building new fabs for new nodes. Three years later, though, **Intel had to build more 14nm capacity** after **hiring Samsung** to help it build chips; the problem is that its 10nm chips were delayed by years (the company just started shipping 10nm parts in volume this year).

Meanwhile, TSMC was racing ahead, with 7nm chips in 2017<sup>1</sup>, and 5nm chip production **starting this year**; this, combined with Apple’s chip design expertise, meant that as of last fall iPhone chips were comparable in speed to the top-of-the-line iMac chips. From **Anandtech**:

We’ve now included the latest high-end desktop CPUs as well to give context as to where the mobile is at in terms of absolute performance.



Overall, in terms of performance, the A13 and the Lightning cores are extremely fast. In the mobile space, there's really no competition as the A13 posts almost double the performance of the next best non-Apple SoC. The difference is a little bit less in the floating-point suite, but again we're not expecting any proper competition for at least another 2-3 years, and Apple isn't standing still either.

Last year I've noted that the A12 was margins off the best desktop CPU cores. This year, the A13 has essentially matched best that AMD and Intel have to offer – in SPECint2006 at least. In SPECfp2006 the

A13 is still roughly 15% behind.

The **Intel Core i9-9900K Processor** in those charts **launched at price of \$999** before settling in at a street price of around \$520; it remains the top-of-the-line option for the iMac for an upgrade price of \$500 above the Intel Core i5-8600K, a chip that launched at \$420 and today costs \$220. The A13, meanwhile, probably costs between \$50~\$60.<sup>2</sup>

This is what made next week's reported announcement feel inevitable: Apple's willingness to invest in the Mac seems to have truly turned around in 2017 — not only has the promised Mac Pro launched, but so has an entirely new MacBook line with a redesigned keyboard — even as the cost of sticking with Intel has become not simply about money but also performance.

## The Implications of ARM

The most obvious implication of Apple's shift — again, assuming the reporting is accurate — is that ARM Macs will have superior performance to Intel Macs on both a per-watt basis and a per-dollar basis. That means that the next version of the MacBook Air, for example, could be cheaper even as it has better battery life and far better performance (the i3-1000NG4 Intel processor that is the cheapest option for the MacBook Air is not yet for public sale; it probably costs around \$150, with far worse performance than the A13).

What remains to be seen is just how quickly Apple will push ARM into its higher-end computers. Again, the A13 is already competitive with some of Intel's best desktop chips, and the A13 is tuned for mobile; what sort of performance gains can Apple uncover by building for more generous thermal envelopes? It is not out of the question that Apple, within a year or two, has by far the best performing laptops and desktop computers on the market, just as they do in mobile.

This is where Apple's tight control of its entire stack can really shine: first, because Apple has always been less concerned with backwards compatibility than Microsoft, it has been able to shepherd its developers into a world where this sort of transition should be easier than it would be on, say, Windows; notably the company has over the last decade deprecated its Carbon API and ended 32-bit support with the current version of macOS. Even the developers that have the furthest to go are well down the road.

Second, because Apple makes its own devices, it can more quickly leverage its ability to design custom chips for macOS. Again, I'm not completely certain the economics justify this — perhaps Apple sticks with one chip family for both iOS and the Mac — but if it is going through the hassle of this change, why not go all the way (notably, one thing Apple does not need to give up is Windows support: Windows has run on ARM for the last decade, and I expect Boot Camp to continue, and for virtualization offerings to be available as well; whether this will be as useful as Intel-based virtualization remains to be seen).

What is the most interesting, and perhaps the most profound, is the potential impact on the server market, which is Intel's bread-and-butter. Linus Torvalds, the creator and maintainer of Linux, **explained why he was skeptical** about ARM on the server in 2019:

Some people think that “the cloud” means that the instruction set doesn’t matter. Develop at home, deploy in the cloud. That’s bullshit. If you develop on x86, then you’re going to want to deploy on x86, because you’ll be able to run what you test “at home” (and by “at home” I don’t mean literally in your home, but in your work environment). Which means that you’ll happily pay a bit more for x86 cloud hosting, simply because it matches what you can test on your own local setup, and the errors you get will translate better...

Without a development platform, ARM in the server space is never going to make it. Trying to sell a 64-bit “hyperscaling” model is idiotic, when you don’t have customers and you don’t have workloads because you never sold the small cheap box that got the whole market started in the first place...

The only way that changes is if you end up saying “look, you can deploy more cheaply on an ARM box, and here’s the development box you can do your work on”. Actual hardware for developers is hugely important. I seriously claim that this is why the PC took over, and why everything else died...It’s why x86 won. Do you really think the world has changed radically?

ARM on Mac, particularly for developers, could be a radical change indeed that ends up transforming the server space. On the other hand, the shift to ARM could backfire on Apple: Windows, particularly given [the ability to run a full-on Linux environment without virtualization](#), combined with Microsoft’s developer-first approach, is an extremely attractive alternative that many developers just don’t know about — but they may be very interested in learning more if that is the price of running x86 like their servers do.

## Intel’s Failure

What is notable about this unknown — will developer preferences for macOS lead to servers switching to ARM (which remember, is cheaper and likely more power efficient in servers as well), or will the existing x86 installation base drive developers to Windows/Linux — is that the outcome is out of Intel’s control.

What started Intel’s fall from king of the industry to observer of its fate was its momentous 2005 decision to not build chips for the iPhone; then-CEO Paul Otellini told [Alexis Madrigal at The Atlantic](#) what happened:[3](#)

“We ended up not winning it or passing on it, depending on how you want to view it. And the world would have been a lot different if we’d done it,” Otellini told me in a two-hour conversation during his last month at Intel. “The thing you have to remember is that this was before the iPhone was introduced and no one knew what the iPhone would do...At the end of the day, there was a chip that they were interested in that they wanted to pay a certain price for and not a nickel more and that price was below our forecasted cost. I couldn’t see it. It wasn’t one of these things you can make up on volume. And in hindsight, the forecasted cost was wrong and the volume was 100x what anyone thought.”

What is so disappointing about this excuse is that it runs directly counter to what made Intel great; in 1965, Bob Noyce, then at Fairchild Semiconductor[4](#), shocked the semiconductor world by announcing that Fairchild would price its integrated circuit products at \$1, despite the fact it cost Fairchild far more

than that to produce them. What Noyce understood is that the integrated circuit market was destined to explode, and that by setting a low price Fairchild would not only accelerate that growth, but also drive down its costs far more quickly than it might have otherwise (chips, remember, are effectively zero marginal cost items; the primary costs are the capital costs of setting up manufacturing lines).

That is the exact logic that Otellini “couldn’t see”, so blinded he was by the seemingly dominant PC paradigm and Intel’s enviable profit margins.<sup>5</sup> Worse, those volumes went to manufacturers like TSMC instead, providing the capital for research and development and capital investment that has propelled TSMC into the fabrication lead.

*CORRECTION: A source suggested that this sentence was wrong:*

*What started Intel’s fall from king of the industry to observer of its fate was its momentous 2005 decision to not build chips for the iPhone.*

*XScale, Intel’s ARM chips, were engineered to be fast, not power-efficient, and Intel wasn’t interested in changing their approach; this is particularly striking given that Intel had just recovered from having made the same mistake with the Pentium 4 generation of its x86 chips. Moreover, the source added, Intel wasn’t interested in doing any sort of customization for Apple: their attitude was take-it-or-leave-it for, again, a chip that wasn’t even optimized correctly. A better sentence would have read:*

*Intel’s fall from king of the industry to observer of its fate was already in motion by 2005: despite the fact Intel had an ARM license for its XScale business, the company refused to focus on power efficiency and preferred to dictate designs to customers like Apple, contemplating their new iPhone, instead of trying to accommodate them (like TSMC).*

*What is notable is that doesn’t change the sentiment: the root cause was Intel’s insistence on **integrating design and manufacturing**, certain that their then-lead in the latter would leave customers no choice but to accept the former, and pay through the nose to boot. It was a view of the world that was, as I wrote, “blinded...by the seemingly dominant PC paradigm and Intel’s enviable profit margins.”*

*My apologies for the error, but also deep appreciation for the correction.*

That is why, last month, it was TSMC that was the target of a federal government-led effort to build a new foundry in the U.S.; I explained in **Chips and Geopolitics**:

Taiwan, you will note, is just off the coast of China. South Korea, home to Samsung, which also makes the highest end chips, although mostly for its own use, is just as close. The United States, meanwhile, is on the other side of the Pacific Ocean. There are advanced foundries in Oregon, New Mexico, and Arizona, but they are operated by Intel, and Intel makes chips for its own integrated use cases only.

The reason this matters is because chips matter for many use cases outside of PCs and servers — Intel's focus — which is to say that TSMC matters. Nearly every piece of equipment these days, military or otherwise, has a processor inside. Some of these don't require particularly high performance, and can be manufactured by fabs built years ago all over the U.S. and across the world; others, though, require the most advanced processes, which means they must be manufactured in Taiwan by TSMC.

This is a big problem if you are a U.S. military planner. Your job is not to figure out if there will ever be a war between the U.S. and China, but to plan for an eventuality you hope never occurs. And in that planning the fact that TSMC's foundries — and Samsung's — are within easy reach of Chinese missiles is a major issue.

I think the focus on TSMC was correct, and I am encouraged by TSMC's decision to build a foundry in Arizona, even if they are moving as slowly as they can on a relatively small design; at the same time, what a damning indictment of Intel. The company has not simply lost its manufacturing lead, and is not simply a helpless observer of a potentially devastating shift in developer mindshare from x86 to ARM, but also when its own country needed to subsidize the building of a foundry for national security reasons Intel wasn't even a realistic option, and a company from a territory claimed by China was.

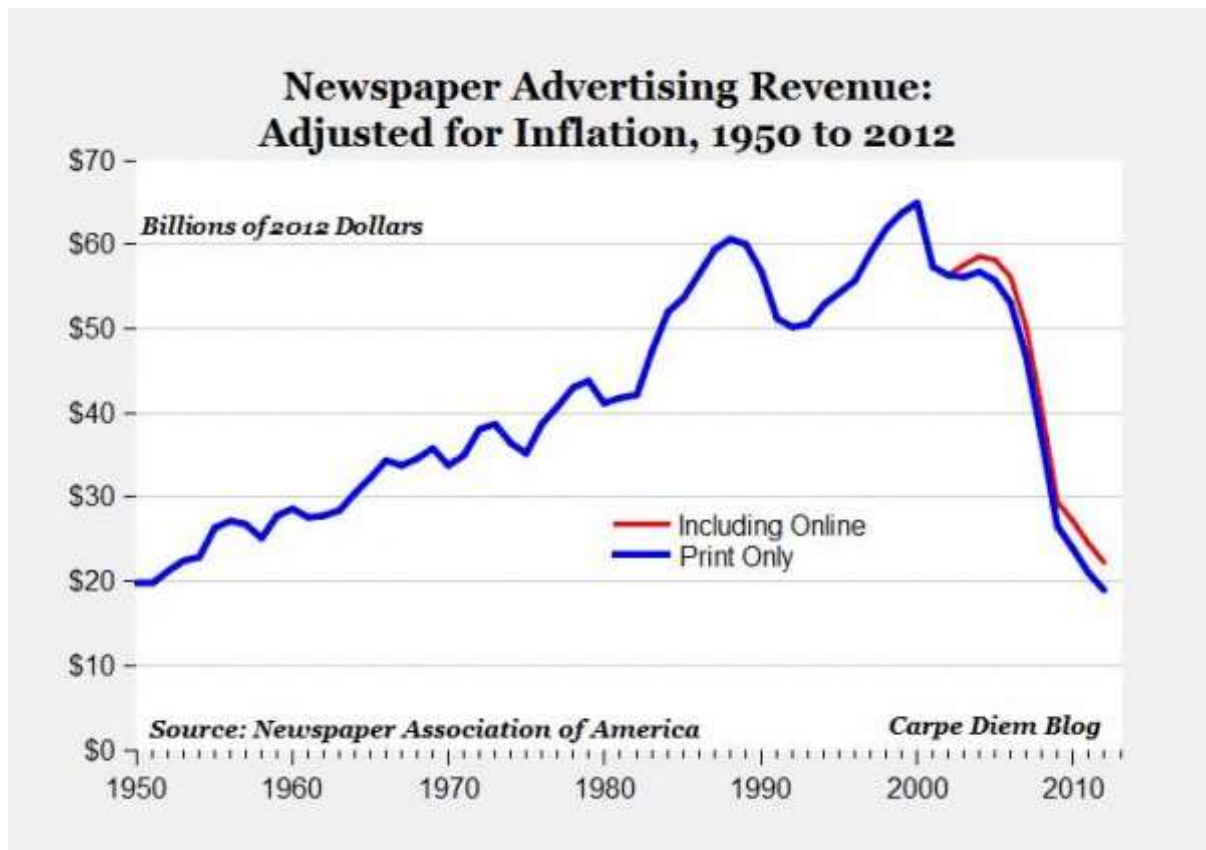
To that end, while I am encouraged by and fully support [this bill by Congress](#) to appropriate \$22.8 billion in aid to semiconductor manufacturers (the amount should be higher), I wonder if it isn't time for someone to start the next great U.S. chip manufacturing company. No, it doesn't really make economic sense, but this is an industry where aggressive federal industrial policy can and should make a difference, and it's hard to accept the idea of taxpayer billions going to a once-great company that has long-since [forgotten](#) what [made it great](#). Intel has prioritized profit margins and perceived lower risk for decades, and it is only now that the real risks of caring about [finances more than fabrication](#) are becoming apparent, for both Intel and the United States.

- 
1. Node sizes are not an exact measure; most industry experts consider TSMC's 7nm node size to be comparable to Intel's 10nm size [[↔](#)]
  2. This number is extremely hard to source; but to the degree I am off it is by the tens of dollars, not hundreds [[↔](#)]
  3. I first used this quote in [Andy Grove and the iPhone SE](#) [[↔](#)]
  4. Noyce and Gordon Moore would form Intel with a large number of Fairchild employees three years later [[↔](#)]
  5. Incredibly, Otellini then doubled-down: Intel actually sold the ARM division that Jobs had wanted access to a year later. [[↔](#)]
- 

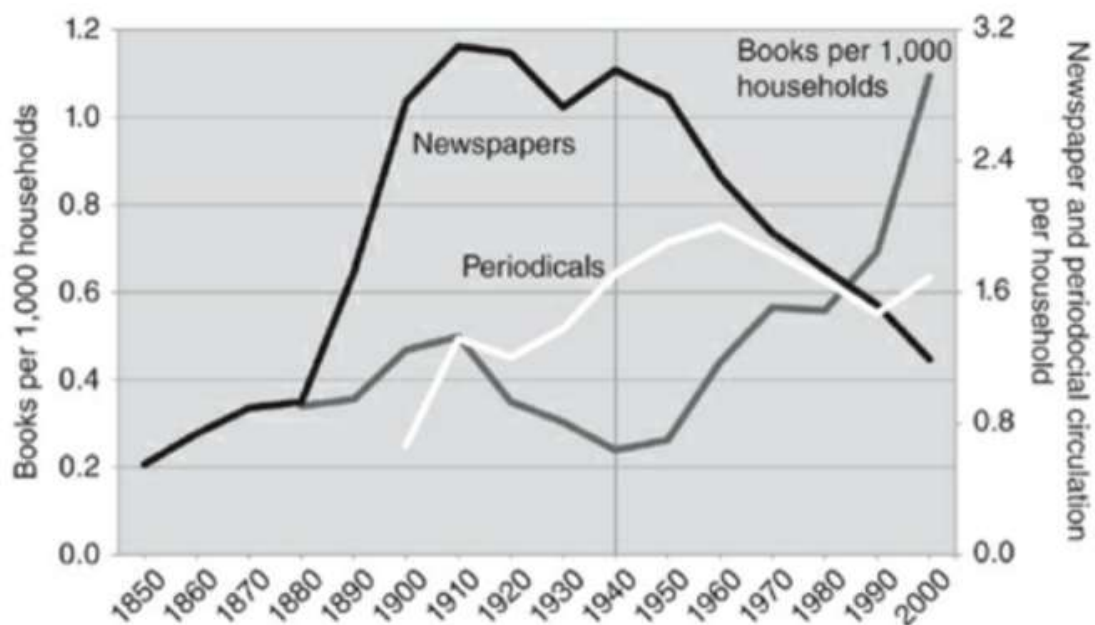
## Never-ending Niches

Tuesday, June 9, 2020

You have almost certainly seen this chart about newspaper advertising revenue since World War II:



The obvious takeaway is that the Internet killed what had been a profitable and growing business; what is interesting, though, is that circulation numbers tell a somewhat different story:



**Figure 6-1.** Book Publications, Newspaper Circulation, and Periodical Circulation Per Household, by Decade, 1850–2009

Sources: HSUS series Dg225, Dg254, Dg256, Dg266, and Ae2; SAUS tables 1129, 1134, 1135. Each year refers to the average over the following decade (e.g. 2000 refers to the average over 2000–2009).

## Time and Reach

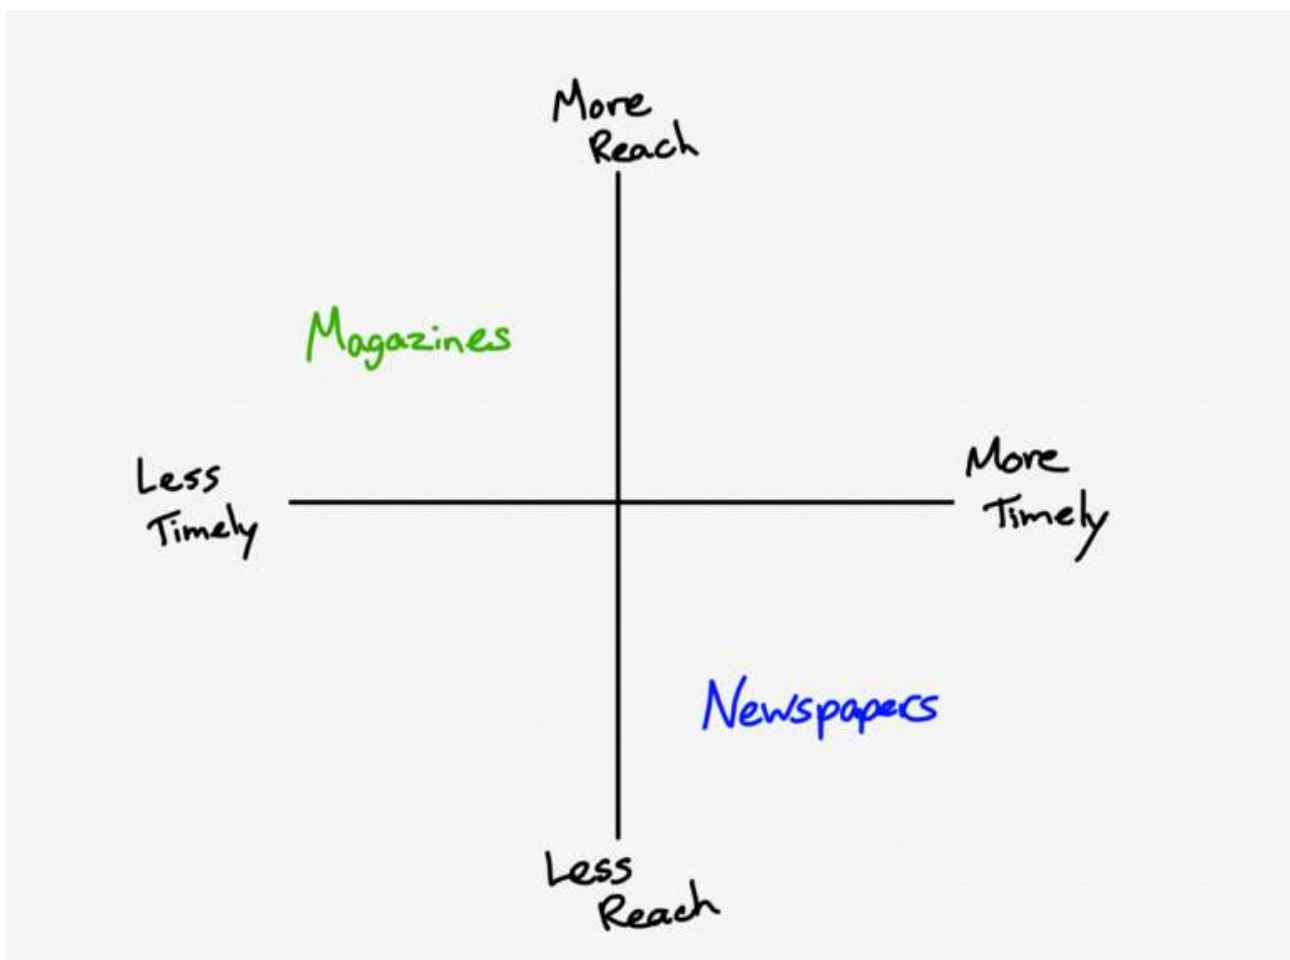
That image is from Robert Gordon's book *The Rise and Fall of American Growth*; one of the more startling facts surrounding this graph is that between 1910 and 1930 the average household purchased 3.1 different newspapers a day. As Gordon notes:

The fastest growth occurred in 1870–1900, by which time newspapers had become firmly established as the main source of information and entertainment for a growing population. Color presses were introduced in the 1890s and were first used to produce color comics and supplements. By the early twentieth century, newspapers had extended their content far beyond the news itself and added “gossip columns, travel and leisure advice, color comics, and sporting results.”

It turned out, though, that the daily deadline inherent to newspapers presented a market opportunity for periodicals. Gordon writes:

The mass-circulation national magazine was a creation of the 1880s and 1890s. Unlike newspapers, for which the circulation area was limited by the need to provide time-sensitive news to a particular metropolitan area, the features contained in magazines could reach readers at a more leisurely pace. Hence magazines were national almost from the beginning in the mid-nineteenth century, and among those with the highest circulations late in the century were *McClure's*, *Collier's*, the *Saturday Evening Post*, and the *Ladies' Home Journal*.

In other words, the market opportunity was defined by the intersection of time and reach; newspapers needed to be timely, but that limited reach, whereas periodicals, by virtue of being weekly or monthly, could also have much greater reach:



These limitations of time and space affected the nature of content as well. While newspapers were focused on “time-sensitive news”, magazines, in a format pioneered by Henry Luce’s *Time* magazine focused more on contextualizing and analyzing the news that you probably already knew about because of your daily newspaper.

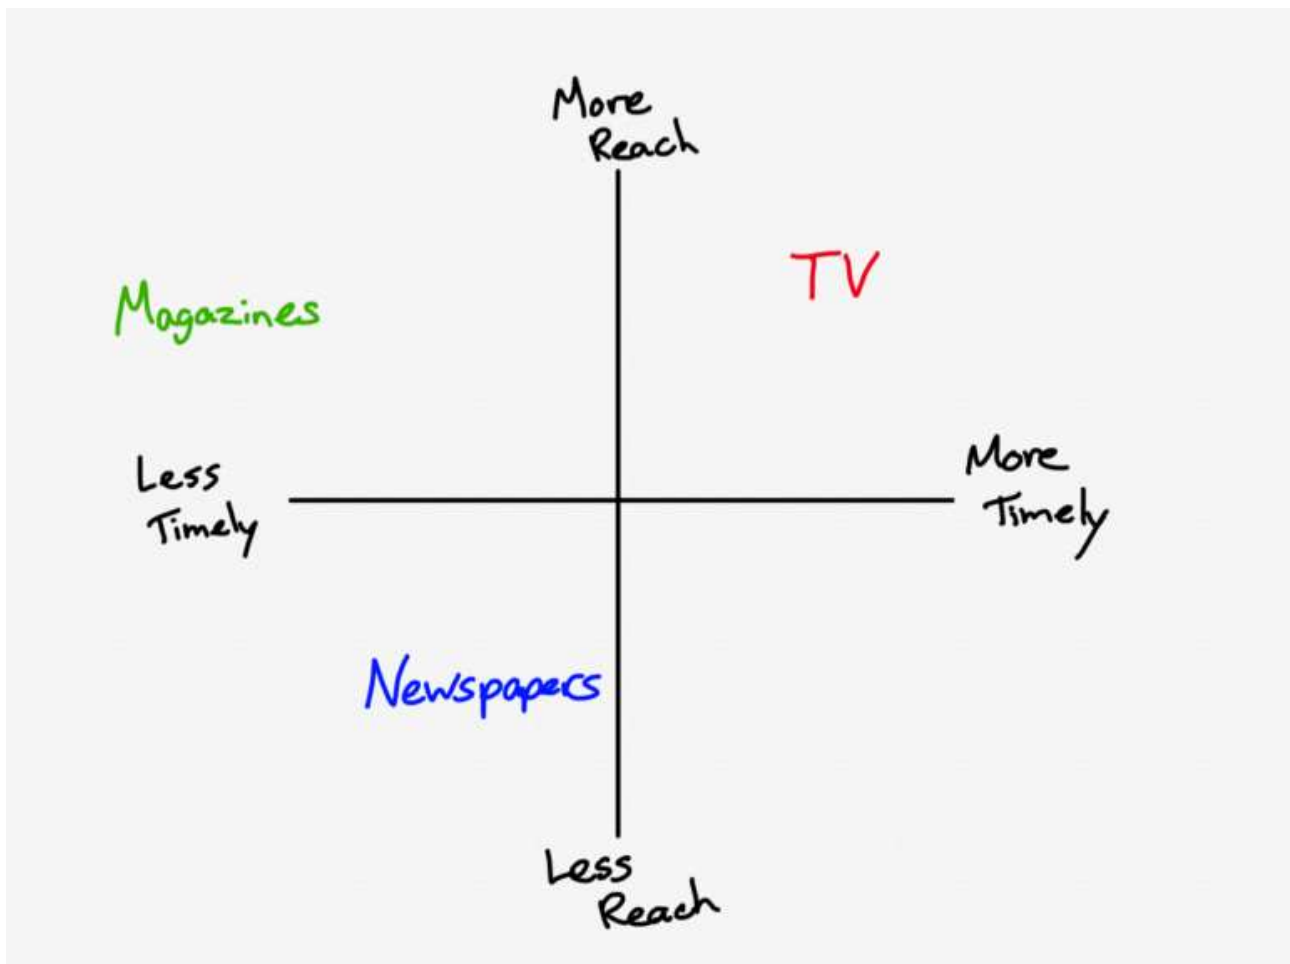
What drove the post-war decline in newspaper circulation was the television. Gordon writes:

Television news, on the other hand, increasingly gained a quality and credibility that the newsreel lacked. The information came first from a familiar anchor, with footage used for the most part as a substantive complement rather than purely as an eye-catching distraction. Indeed, the familiarity of the network anchors helped television news surpass not only the newsreel, but also the newspaper, as the main source of news...

The effect of TV news went beyond the ability to present a familiar, trusted face to deliver the top stories. Like radio, television news was immediate, but it also wielded the power of the image to evoke strong feelings in the viewer. In addition to coverage of the day’s top stories, television news also developed more in-depth journalistic programs...and cable television delivered CNN as the first twenty-four-hour news station in 1980...

Unsurprisingly, after reaching all-time highs in the postwar years, newspaper circulation per household soon began a gradual but continuous decline, dropping from 1.4 per household in 1949 to 0.8 in 1980 and to less than 0.4 in 2010.

This meant that the graph I created above now looked like this:



A daily cadence in a world of CNN was simply the modern version of a weekly cadence for magazines in a world of daily newspapers; what was particularly challenging for newspapers, though, is that TV had further reach as well.

That led the most forward-thinking U.S. newspapers to not just shift towards analysis like magazines had once done, but to also push for more reach of their own, none moreso than the *New York Times*; while the company first formulated its national strategy in 1998, the best articulation of its plan came in its [2003 annual report](#):

As we have mentioned in our previous annual letters, our long-term strategy is to operate the leading news and advertising media in each of the markets in which we compete – both nationally and locally. The centerpiece of this strategy is extending the reach of *The New York Times*'s high-quality journalism into homes and businesses in every city, town, village and hamlet of this country.

Then came the Internet, which accomplished exactly that.

## Evolution Versus Revolution

In [yesterday's Daily Update](#) I described how Jeffrey Katzenberg, the founder of Quibi, mistakenly assumed that mobile was simply the next step in the evolution from motion pictures to television, each of which created new possibilities for, in [Katzenberg's words](#), "the creativity of storytellers [to use] these tools in ways that their inventors had never imagined to amaze audiences."

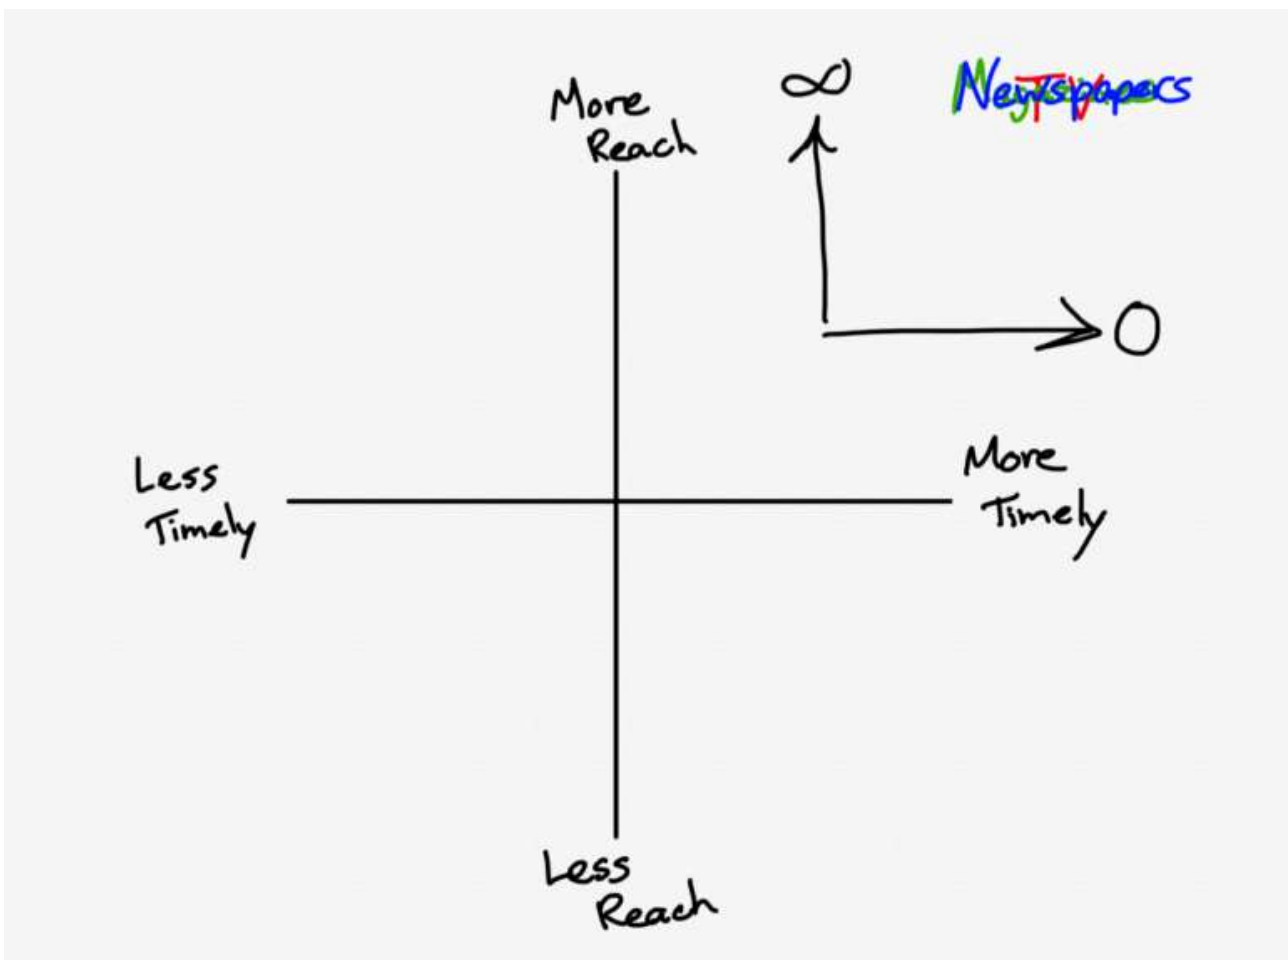
The problem is that mobile was completely different from movies and TV, not because of its form factor, but because of the Internet:

The single most important fact about both movies and television is that they were defined by scarcity: there were only so many movies that would ever be made to fill only so many theater slots, and in the case of TV, there were only 24 hours in a day. That meant that there was significant value in being someone who could figure out what was going to be a hit before it was ever created, and then investing to make it so. That sort of selection and production is what Katzenberg and the rest of Hollywood have been doing for decades, and it's understandable that Katzenberg thought he could apply the same formula to mobile.

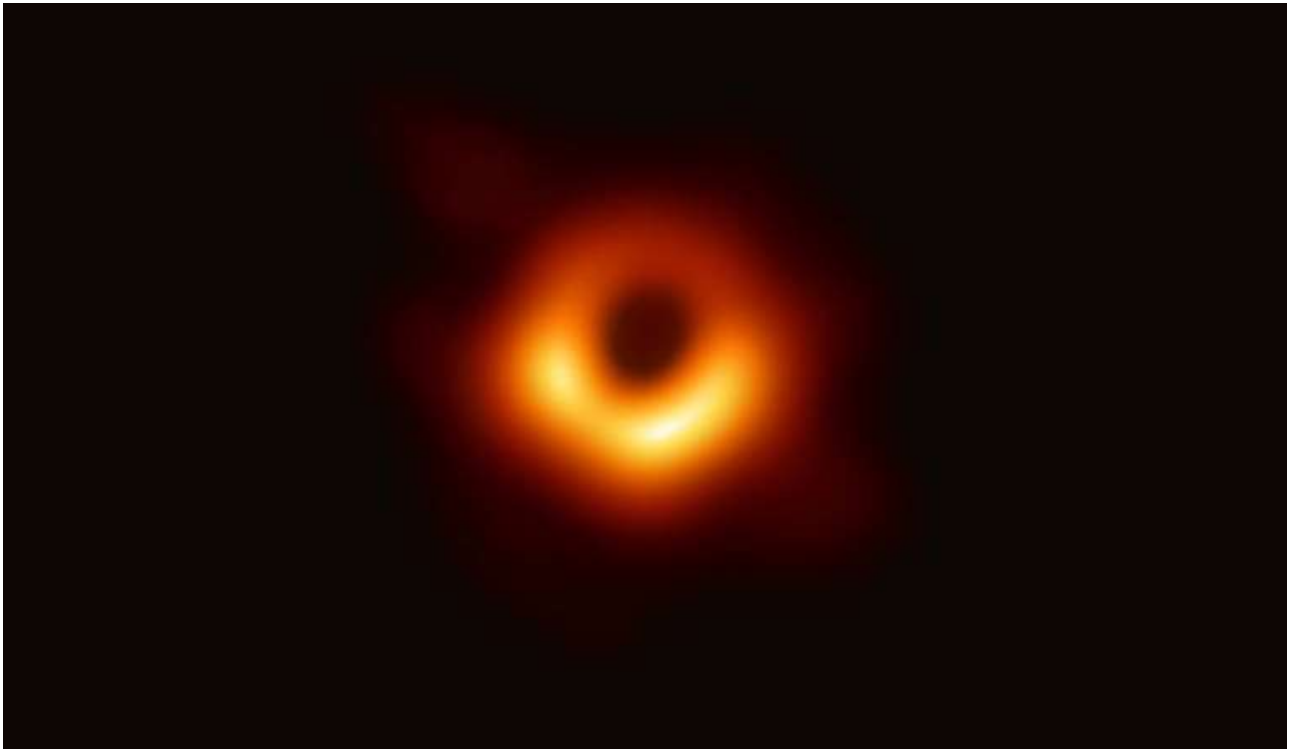
Mobile, though, is defined by the Internet, which is to say it is defined by abundance...The goal is not to pick out the hits, but rather to attract as much content as possible, and then algorithmically boost whatever turns out to be good.

This point cannot be emphasized enough: the Internet is the single most disruptive<sup>1</sup> force of our lifetimes because it does not evolve existing ways of doing things, but completely smashes the assumptions underlying them — assumptions we often didn't even realize existed.

So it was with the Internet and the trade-off between reach and time: suddenly every single media entity on earth, no matter how large or small, and no matter its medium of choice, could reach anyone instantly. To put it another way, reach went to infinity, and time went to zero:



This is, of course, an impossible graph, because zero and infinity cannot be illustrated with axis in two-dimensional space; this is probably a better representation of how time and reach collapsed in on themselves:

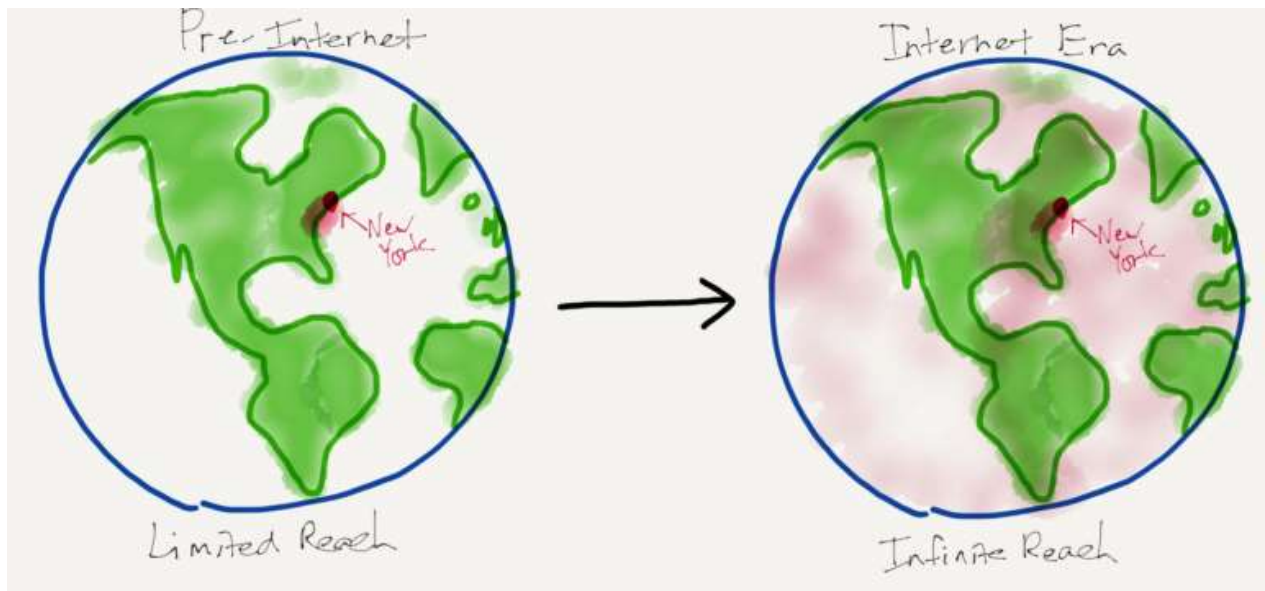


That is the first image of a black hole, and it is certainly an apt metaphor for the Internet: its effect on media assumptions is incalculable and inescapable.

## Competing in an Aggregator World

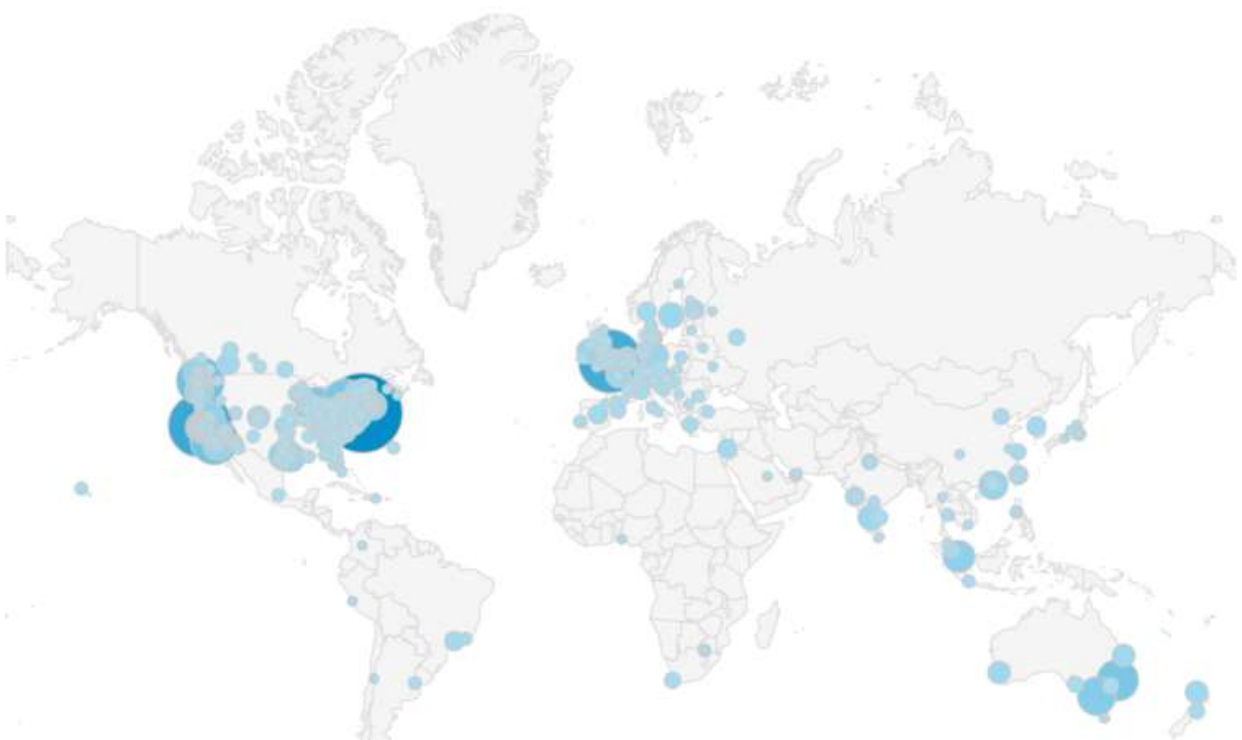
One of the first Stratechery articles about this shift from scarcity to abundance and its impact on media was from 2014 in an article entitled [Economic Power in the Age of Abundance](#).

One of the great paradoxes for newspapers today is that their financial prospects are inversely correlated to their addressable market. Even as advertising revenues have fallen off a cliff – adjusted for inflation, ad revenues are at the same level as the 1950s – newspapers are able to reach audiences not just in their hometowns but literally all over the world.



*Before the Internet, a newspaper like the New York Times was limited in reach; now it can reach anyone on the planet*

The problem for publishers, though, is that the free distribution provided by the Internet is not an exclusive. It's available to every other newspaper as well. Moreover, it's also available to publishers of any type, even bloggers like myself.



*The city-by-city view of Stratechery's readers over the last 30 days.*

To be clear, this is absolutely a boon, particularly for readers, but also for any writer looking to have a broad impact. For your typical newspaper, though, the competitive environment is diametrically opposed to what they are used to: instead of there being a scarce amount of published material, there is an overwhelming abundance. More importantly, this shift in the competitive environment has fundamentally changed just who has economic power.<sup>2</sup>

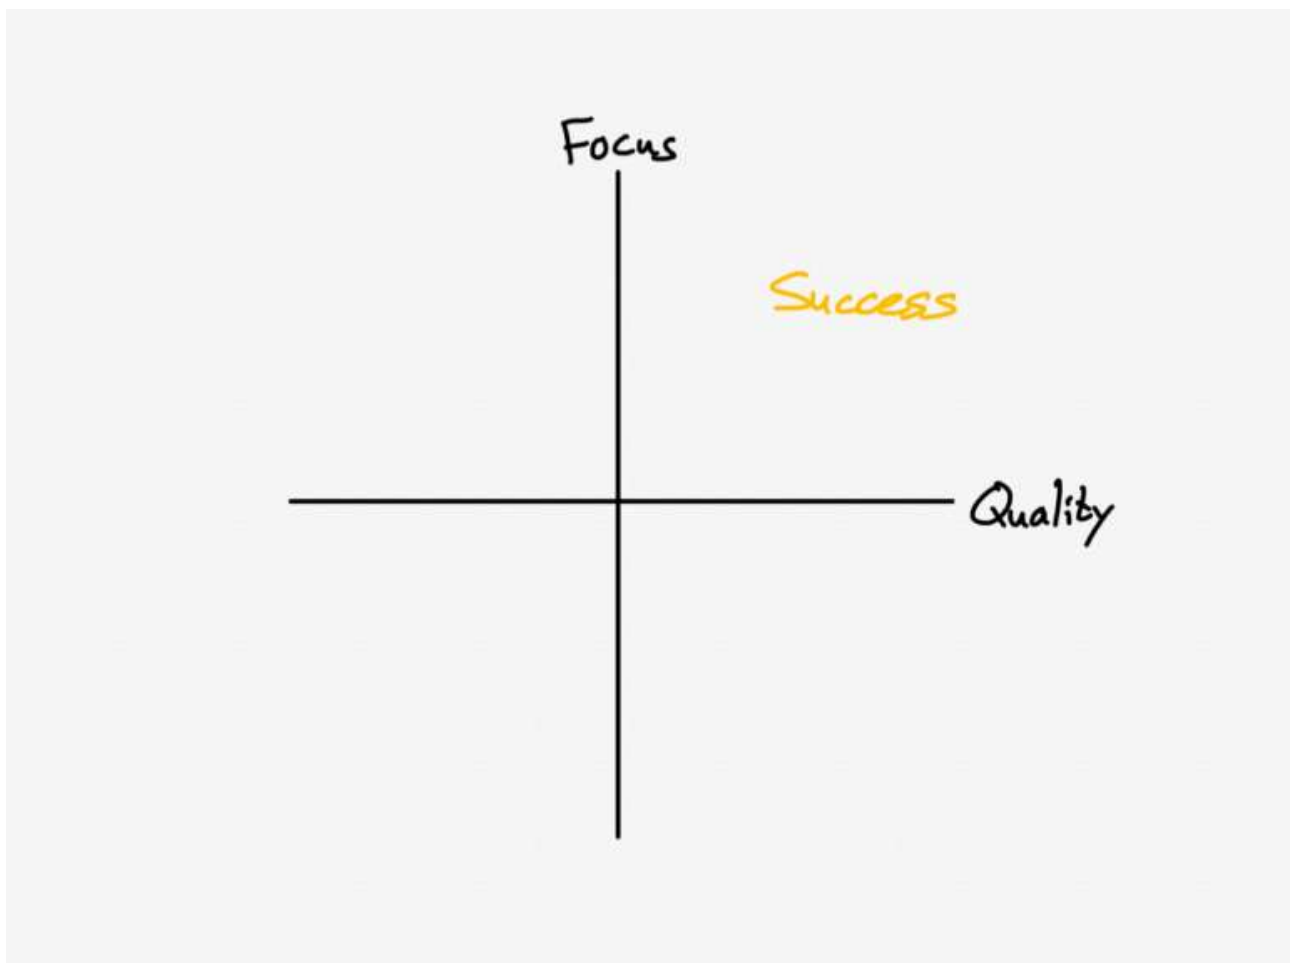
What followed was probably my first clear articulation of **Aggregation Theory**, albeit without the name. The point about effectively infinite competition, though, is a critical one. Neither reach nor timeliness were differentiators, but rather commodities; the companies that dominated on the Internet were those — Google and Facebook in particular — that made sense of the abundance that resulted.

That meant there were three strategies available to media companies looking to survive on the Internet. First, cater to Google. This meant a heavy emphasis on both speed and SEO, and an investment in anticipating and creating content to answer consumer questions. Or you could cater to Facebook, which meant a heavy emphasis on click-bait and human interest stories that had the potential of going viral. Both approaches, though, favored media entities with the best cost structures, not the best content, a particularly difficult road to travel given the massive amounts of content on the Internet created for free.

That left a single alternative: going around Google and Facebook and directly to users.

## Niches and the New York Times

That raises the question as to what are the vectors on which “destination sites” — those that attract users directly, independent of the Aggregators — compete? The obvious two candidates are focus and quality:



What is important to note, though, is that while quality is relatively binary, the number of ways to be focused — that is, the number of niches in the world — are effectively infinite; success, in other words, is about delivering superior quality in your niche — the former is defined by the latter.



This obviously isn't a new concept to Stratechery readers — this is the entire strategic rationale of this site. Again, though, the fact that this is a one-person blog doesn't mean that my competitive situation is any different than that of the *New York Times* or any other media entity on the Internet. In other words, to the extent that the *New York Times* has been successful online — and the company has been **very successful indeed!** — it follows that the company is well-placed in terms of both focus and quality, and in that order.

In this view, the fact that **deeply reported articles about Chinese misinformation on Twitter** are held as being low quality by the Chinese government is immaterial; what matters is that the *New York Times'* audience, which is mostly in the United States, finds it of high quality (I certainly do).

That's an easy example, but there are ones that hit closer to home; for example, I thought this 2018 story that claimed that **Facebook Gave Data Access to Chinese Firm Flagged by U.S. Intelligence** was, **as I wrote at the time**, "deeply flawed at best, and willfully mendacious at worst." It turns out, though, that I am not particularly interested in the "Everything tech does is bad" niche;<sup>3</sup> that story was very high quality for much of the *New York Times'* audience.

I don't bring up this example to complain — quite the contrary! As I wrote in a piece called **In Defense of the New York Times**, after the company wrote an exposé about Amazon's working conditions:

The fact of the matter is that the *New York Times* almost certainly got various details of the Amazon story wrong. The mistake most critics made, though, was in assuming that any publication ever got everything completely correct. Baquet's insistence that good journalism starts a debate may seem like a

cop-out, but it's actually a far healthier approach than the old assumption that any one publication or writer or editor was ever in a position to know "All the News That's Fit to Print."

I'd go further: I think we as a society are in a far stronger place when it comes to knowing the truth than we have ever been previously, and that is thanks to the Internet...the *New York Times* doesn't have the truth, but then again, neither do I, and neither does Amazon. Amazon, though, along with the other platforms that, as described by Aggregation Theory, are increasingly coming to dominate the consumer experience, are increasingly powerful, even more powerful than governments. It is a great relief that the same Internet that makes said companies so powerful is architected such that challenges to that power can never be fully repressed, and I for one hope that the *New York Times* realizes its goal of actually making sustainable revenue in the process of doing said challenging.

Indeed they have, and I see the ongoing criticism of tech as a feature, not a bug.

## Connections and Transformations

There is, in this long-winded explanation, a connection to recent events.

First, that hopeful note about the Internet bringing us closer to the truth by virtue of increasing the amount of information is quite obviously correct. **The light revealing the dust in the air** in terms of the African American experience is coming not from traditional publications, but from the fact that everyone is now a publisher; it turns out the black hole analogy applies not only to analog business models, but also to how the media covered what is clearly not a new problem.

Second, given the nichification of everything, whether by subject matter or sensibility, I am not surprised that the *New York Times* is **finding it difficult** to sustain an opinion section purporting to represent all sides of an issue. This isn't the pre-Internet era, when only a few publications had the reach to plausibly claim they had a duty to show both sides, and more importantly, when that reach defined their competitive advantage. Today all opinions from all people are available everywhere, and the *New York Times*'s ultimate responsibility is to its audience and its reporters.

Third, this discussion explains why **Facebook's calculation should be different**: Facebook (and Google) are not participating in the competition of ideas/attention/monetization, they are defining the terms of that competition. Instead of insisting either company leverage their power explicitly — particularly in terms of politicians that have electoral accountability — more attention should be paid to the fact that that power is completely unaccountable in the first place, and applied in so many ways we cannot see.

What is more worrying is the question of what, if anything, will be connective tissue for society going forward. Infinite niches on as neutral a set of platforms as we can manage makes sense, but by what means do we ensure that people do not disappear into those niches, even if only to decide on how we wish the underlying platforms to be regulated?

Perhaps over time it is geography that will follow business model, instead of the other way around; the shift towards work-from-home is a fascinating development in this regard. What does seem certain is that the past is less of a guide to the future than a reminder that the transformative impact of the Internet is only starting to be felt.

1. In the small ‘d’ sense, not necessarily — although often! — the Christensen sense [↔]
  2. The biggest change in the **city view today**, beyond the relative numbers for the circles, is a far heavier representation in India and a fairly lighter representation in China [↔]
  3. I’ve been pretty consistently in the **tech is an amoral force** camp, criticizing both those that see only evil and **those that see only good** [↔]
- 

## Exponent Podcast: Speech and Systems

Friday, June 5, 2020

On Exponent, the weekly podcast I host with **James Allworth**, we discuss **Dust in the Light**.

Listen to it **here**.

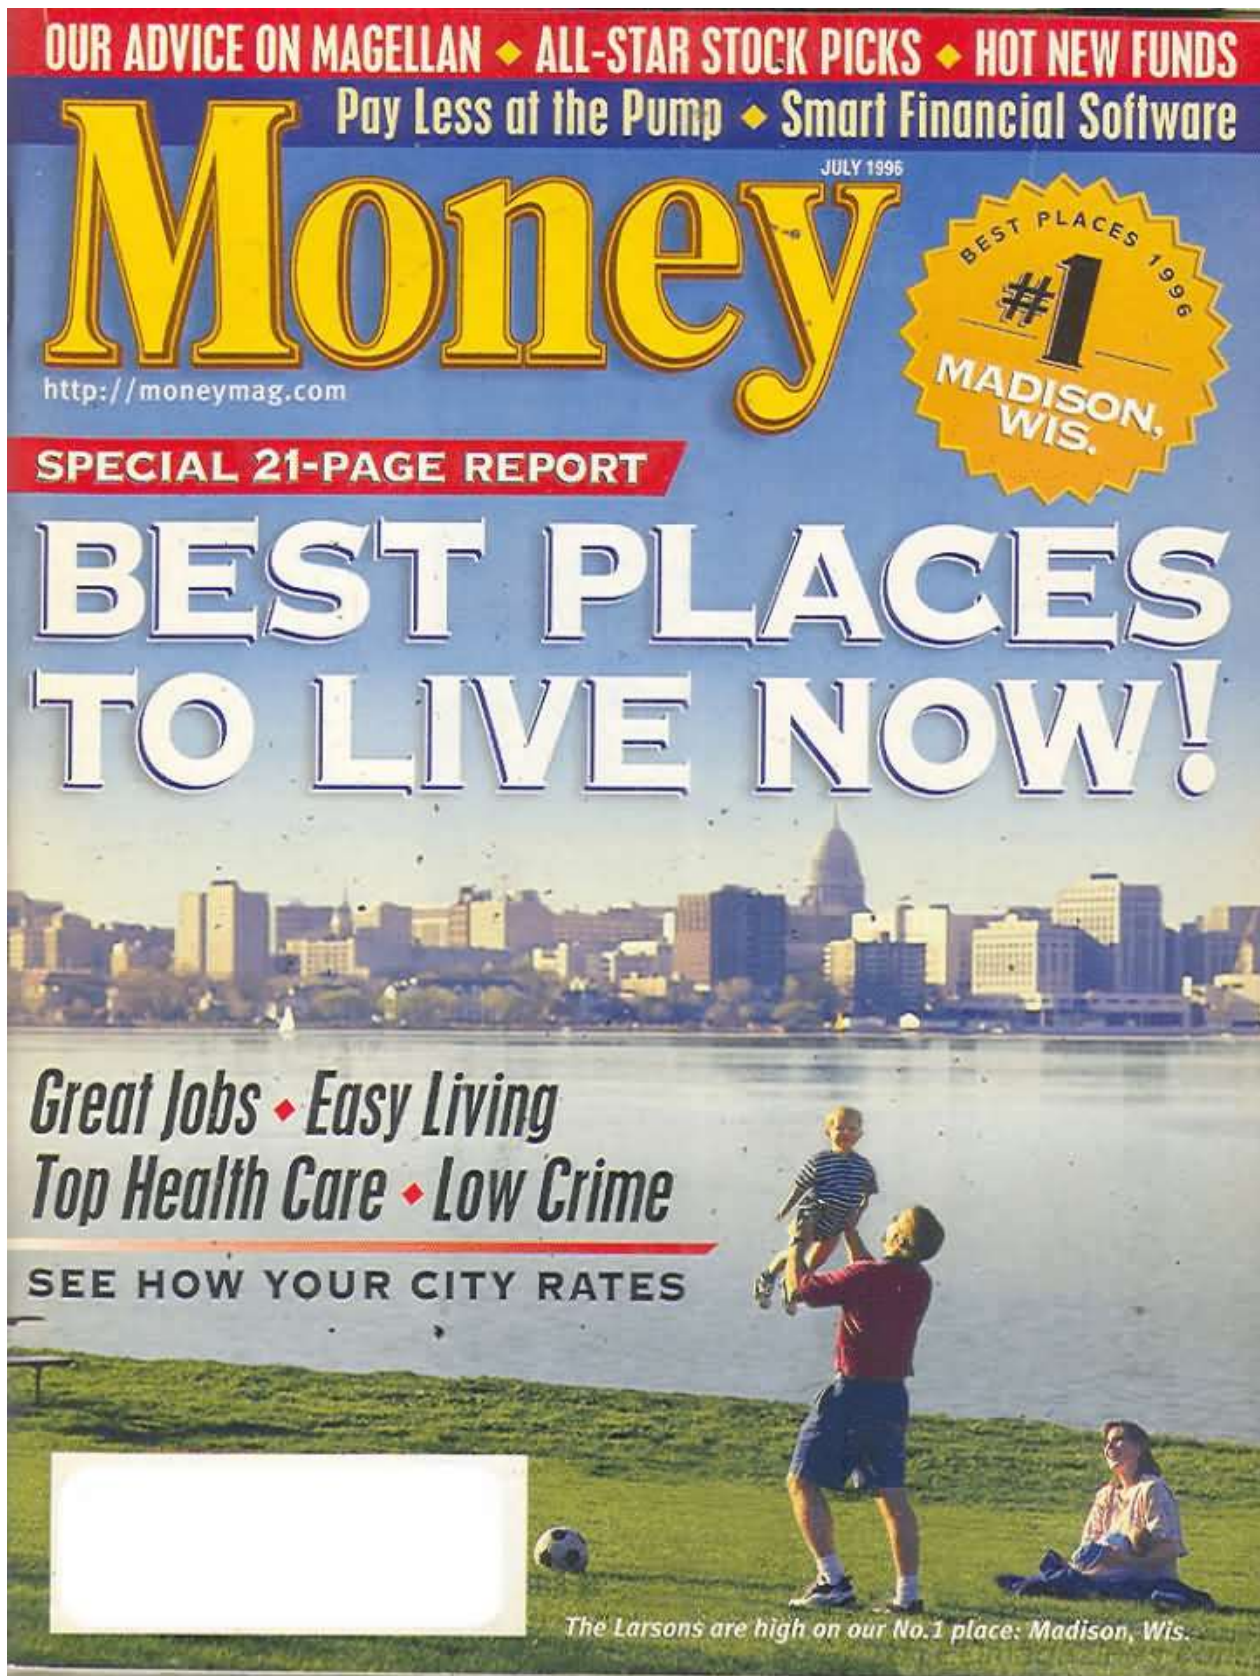
---

## Dust in the Light

Tuesday, June 2, 2020

Everyone in Madison knew to avoid Badger Road.

It was 1996, and the city was celebrating being christened the **best place to live in America by Money Magazine**:



This year, Madison (and the rest of Dane County) earns the No. 1 position among the 300 biggest U.S. metropolitan areas in our 1996 Best Places to Live in America ranking. It snagged the top spot because apparently someone forgot to tell the folks in Madison that life is supposed to be full of trade-offs. The 390,300 residents of Dane County, 80 miles west of Milwaukee in south-central Wisconsin, have a vibrant economy with plentiful jobs, superb health care and a range of cultural activities usually associated with cities twice as big. Yet this mid-size metro area also offers up a low crime rate and palpable friendliness you might assume are available only in, say, Andy Griffith's Mayberry. The news that the great Dane County is top dog this year probably won't surprise the region's residents. More than 90% of Madisonians

rated their quality of life good or very good in a recent survey. Since the cosmopolitan Madison area — the city accounts for about half the county's population — is surrounded by Wisconsin's everpresent dairy farms, it seems only right to toast 1996's No. 1 big cheese with a wedge of aged Wisconsin cheddar.

Still, despite the excellent quality of life, most everyone had, at one time or another, been made aware that the neighborhoods around South Park Street were “dangerous”; that wasn't such a big deal, though, because no one you knew ever went there.

## Madison's Crescent

In 2016, a blogger named Lew Blank observed that the racial distribution of Madison neighborhoods formed a crescent:

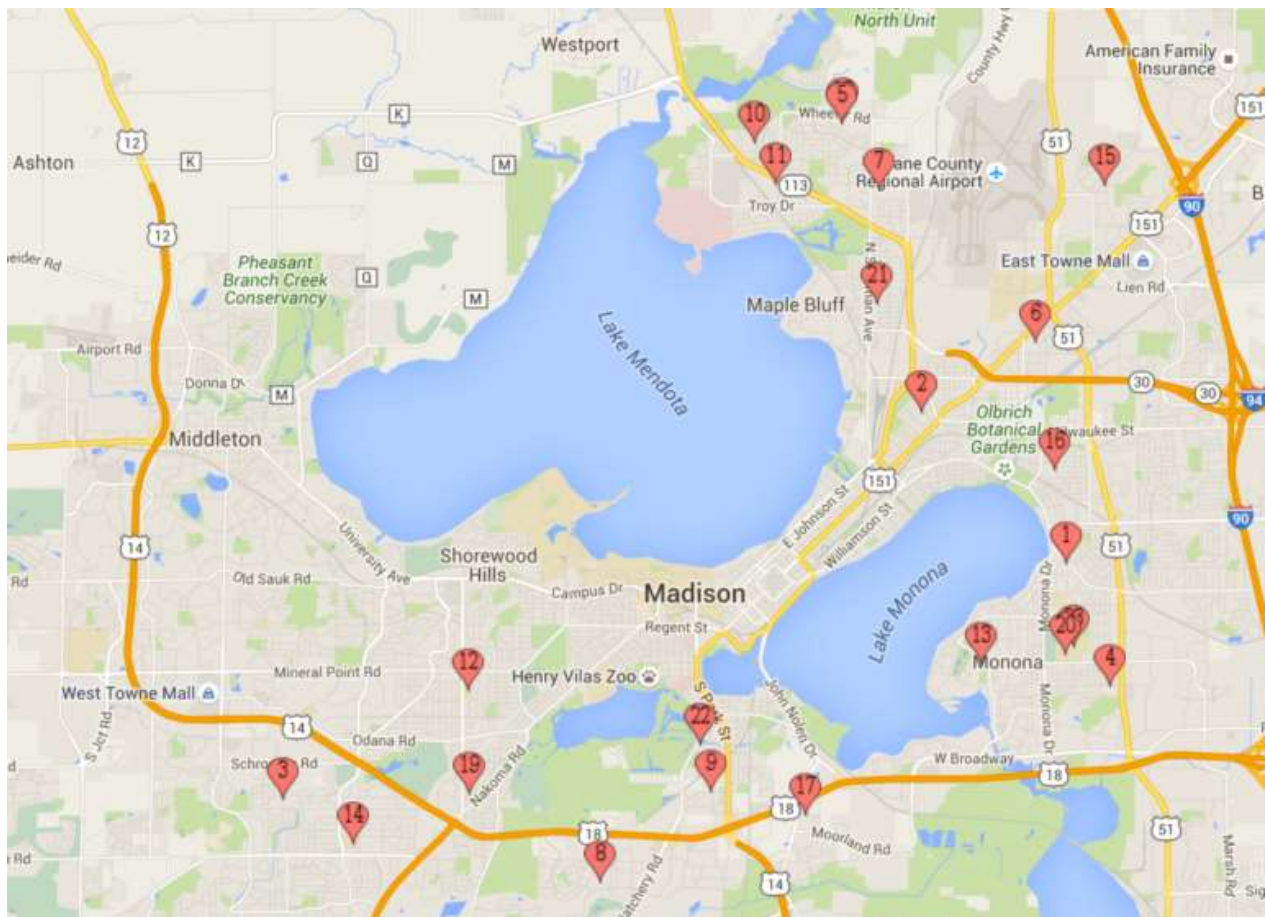
When you look at Madison's Racial Dot Map, you notice a pattern. The bottom and right sides of the map hold the majority of the black and hispanic population. It forms a curve almost — starting in the South Side, crossing along the east side of Lake Monona, and ending at the Northeast Side. I dub this curve-like chain of black and hispanic neighborhoods “The Crescent”.

## THE CRESCENT OF BLACK/HISPANIC NEIGHBORHOODS IN MADISON WI



This crescent was also seen in poverty indicators:

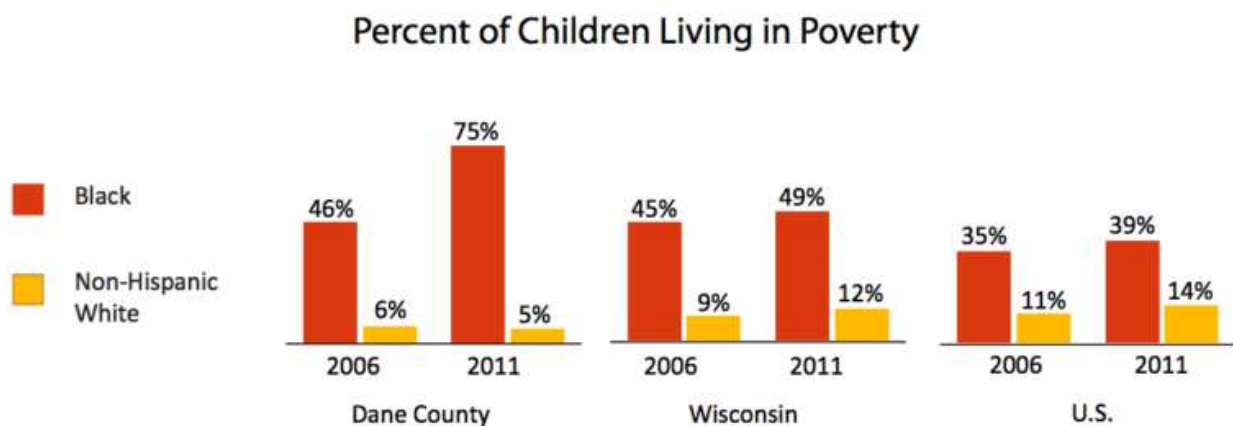
Shown below is a map of every single school in Madison with above average usage of free/reduced lunch programs:



That's right. 23 out of 23 schools in Madison that have above average usage of free/reduced lunch programs all fall along the Crescent.

The deal is, the children who need free/reduced lunch are poor, obviously. So does that mean that the poorest neighborhoods of Madison fall along the Crescent? Unfortunately, that's exactly what it means.

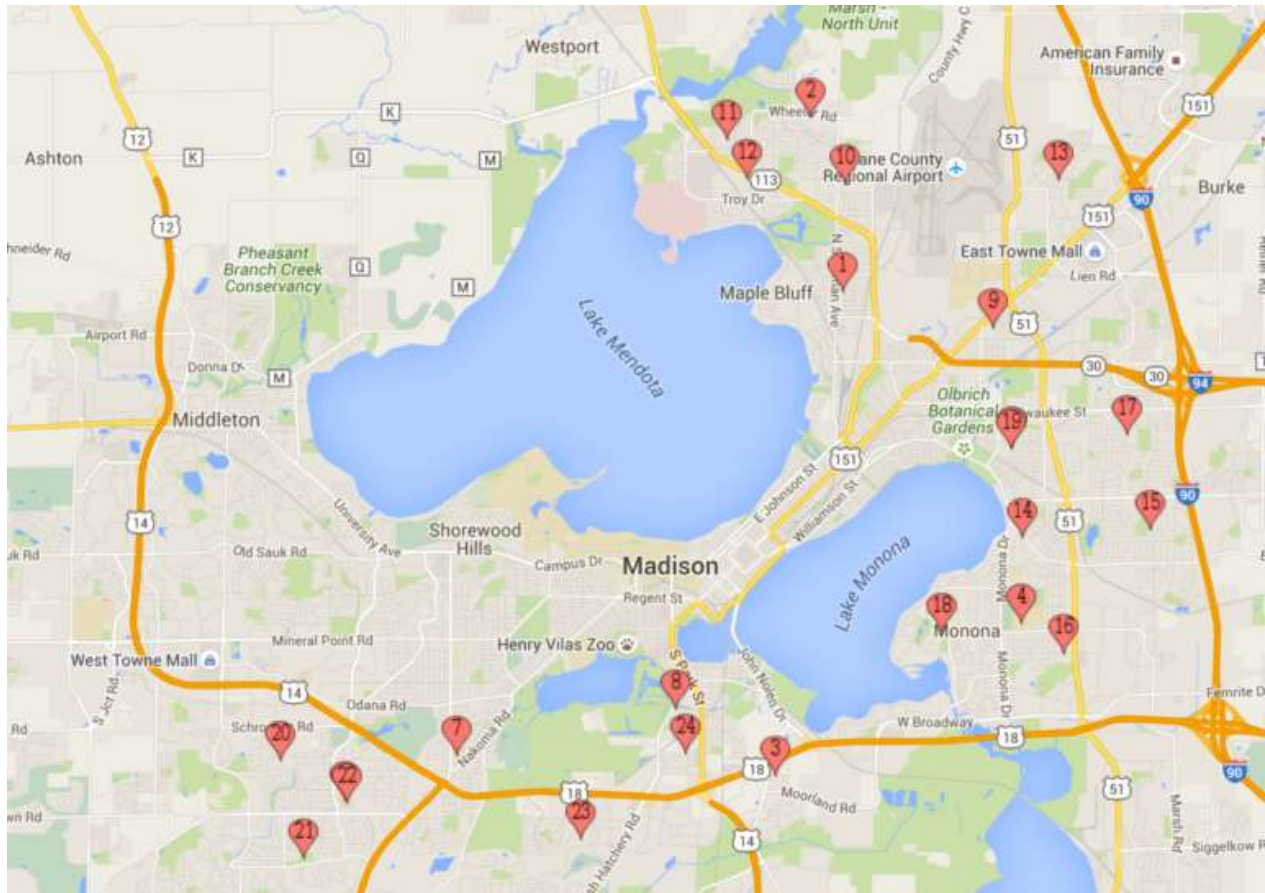
In Madison, a black child is 13 times more likely than a white child to be born into poverty – an insanely high disparity.



So, Madison's black and hispanic neighborhoods (the ones on the Crescent) are its poorest neighborhoods, and Madison's white neighborhoods (the ones not on the Crescent) are its wealthiest neighborhoods.

Unsurprisingly, the crescent was also seen in educational outcomes:

It's clear, then, that schools in the Crescent of poor black/hispanic neighborhoods would be expected to have below-average academic success. And unfortunately, the map below of all Madison schools with below-average reading proficiency rates indicates that this is exactly the case.



Believe it or not, 24 out of 24 schools with below-average reading proficiency rates fall along the Crescent.

And, of course, crime; I personally added the red box to Blank's final map to indicate the Badger Road area I mentioned above:

The map below of the addresses of incarcerated Madisonians shows that incarceration in Madison tends to be clustered around the Crescent.

## Addresses of People in Custody of Dane County Sheriff on a Date in 2015

Shown:

Addresses of ~600 of ~675 people with Dane County addresses.

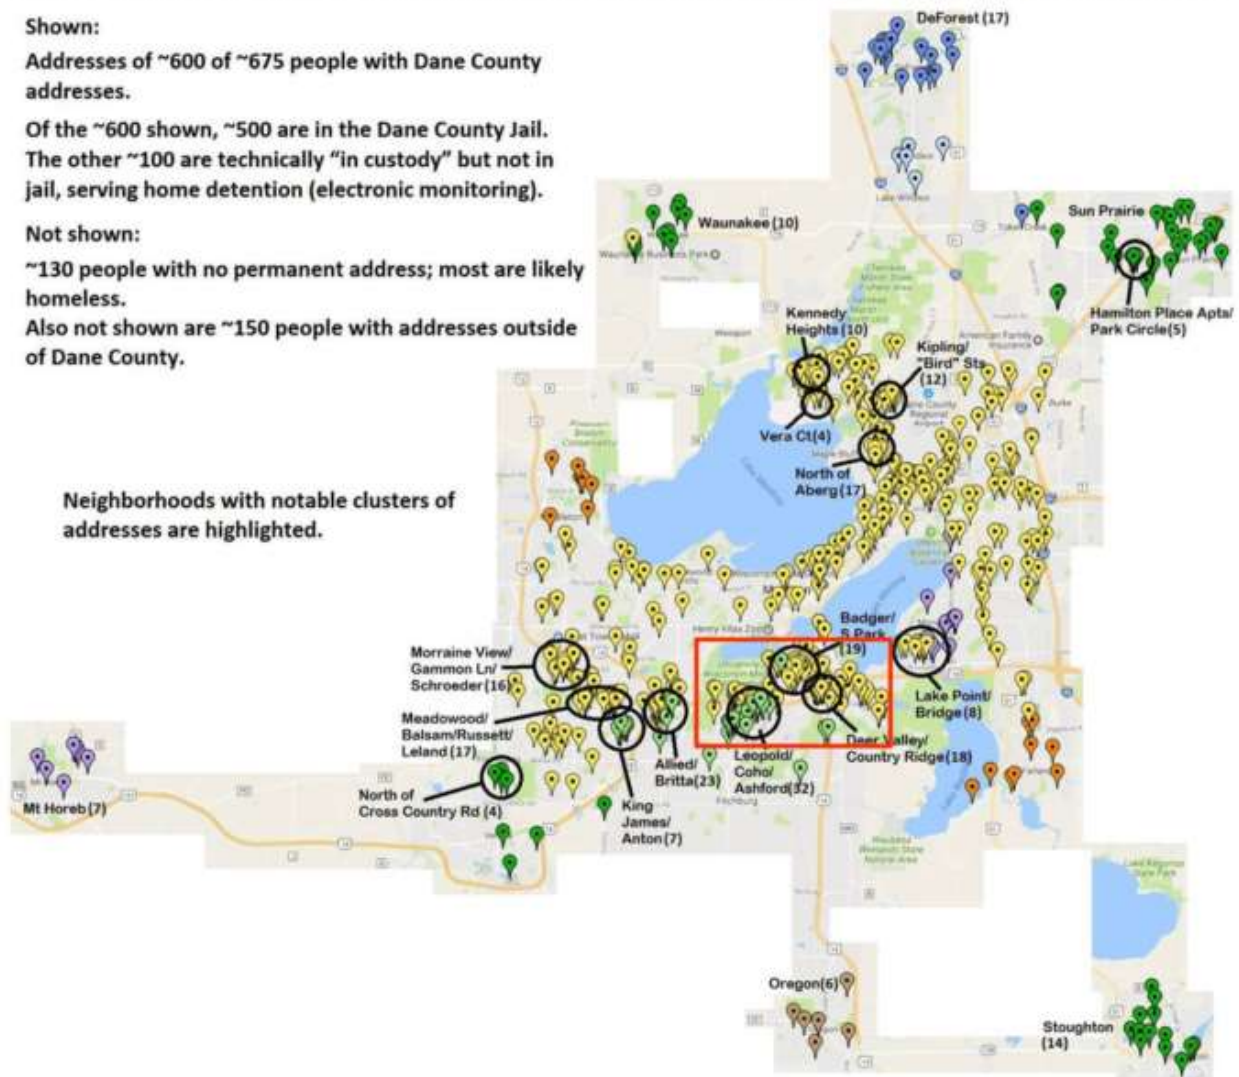
Of the ~600 shown, ~500 are in the Dane County Jail. The other ~100 are technically “in custody” but not in jail, serving home detention (electronic monitoring).

Not shown:

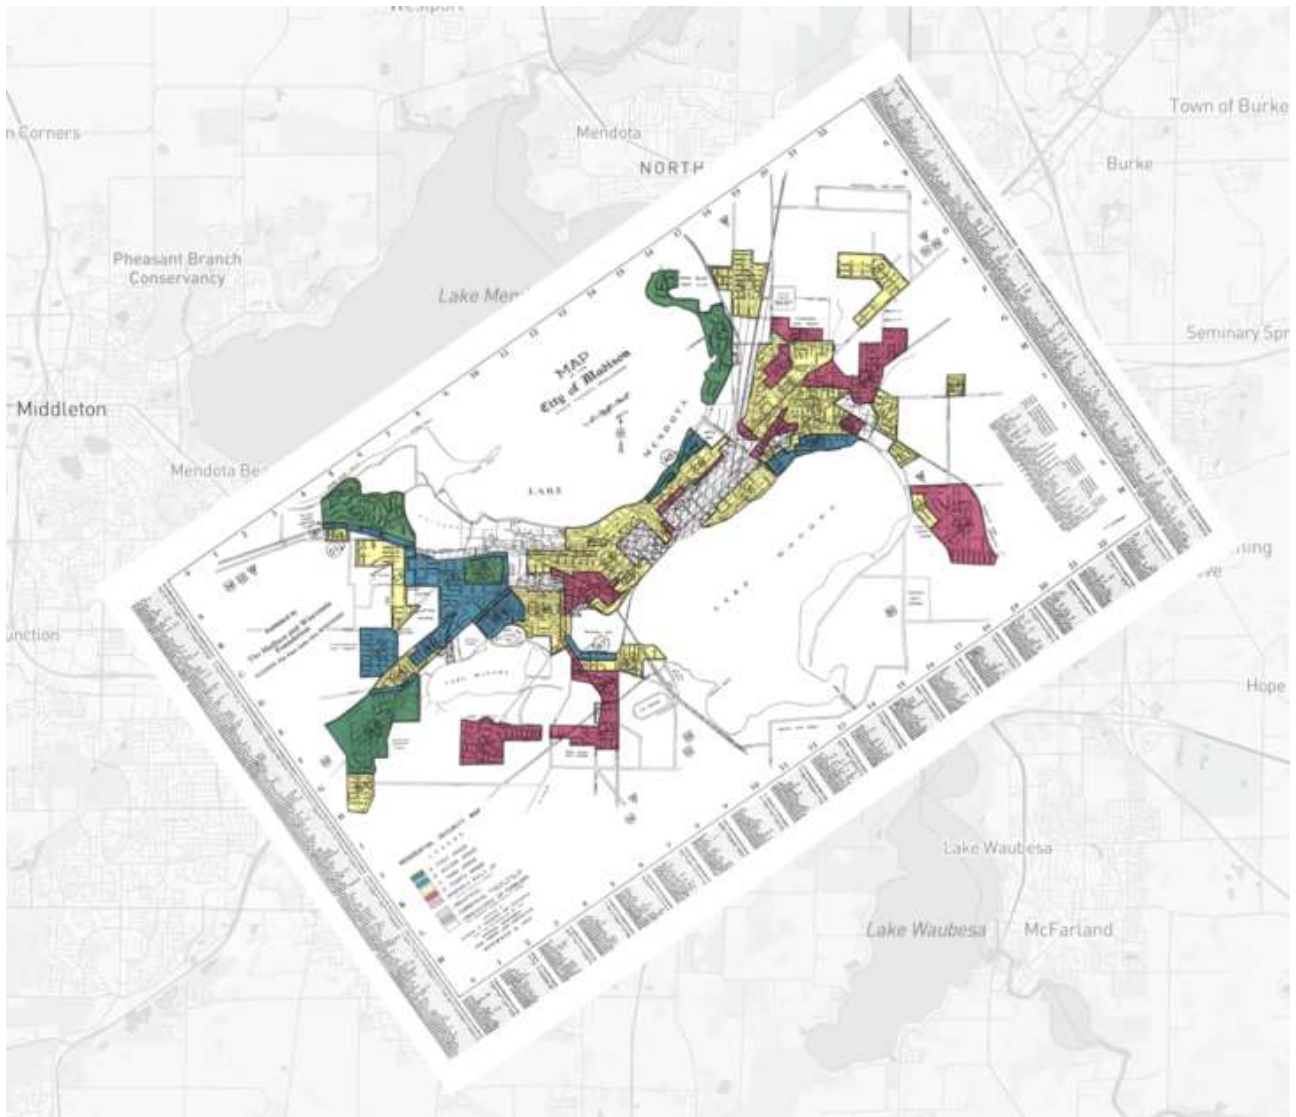
~130 people with no permanent address; most are likely homeless.

Also not shown are ~150 people with addresses outside of Dane County.

Neighborhoods with notable clusters of addresses are highlighted.

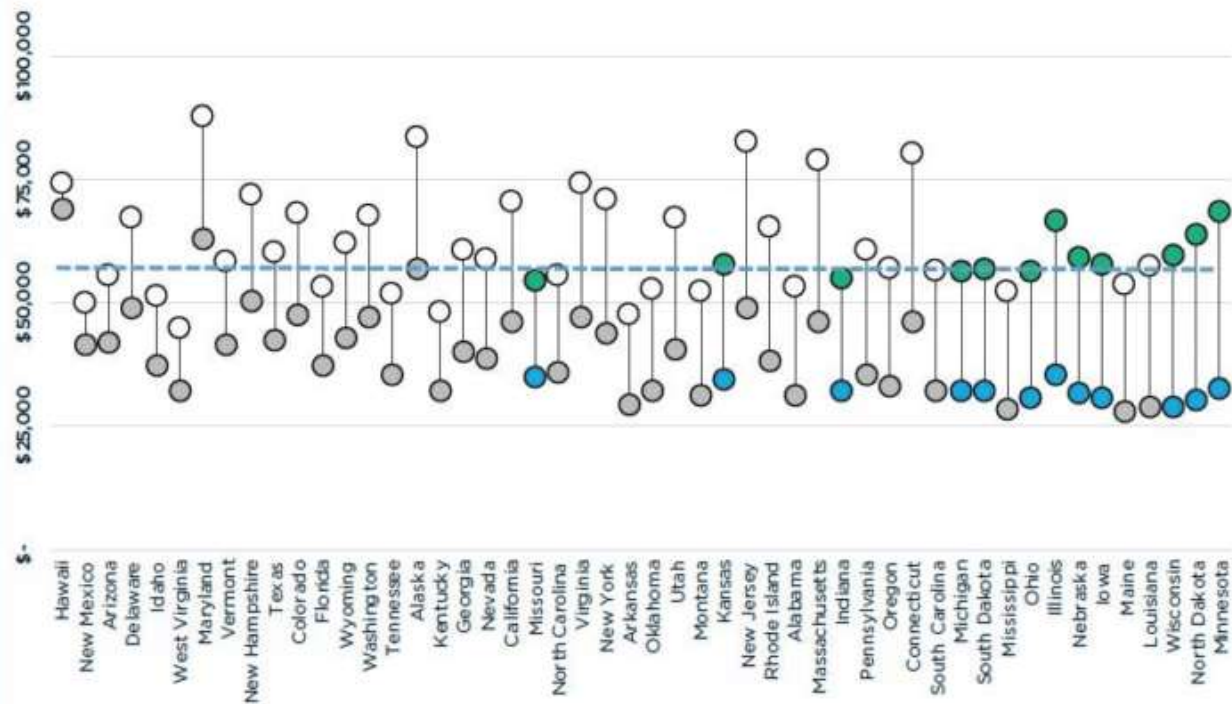


There was one map that Blank was missing though: the notorious Home Owners' Loan Corporation map. The Home Owners' Loan Corporation was a federal agency formed as part of President Franklin D. Roosevelt's New Deal; its purpose was to refinance home mortgages, but as part of the process, the Corporation mapped out U.S. neighborhoods by risk, and by risk, HOLC all-too-frequently meant percentage of non-white people, particularly African Americans. Here was the map of Madison (laid on top of a current map in order to deliver the a north-is-up perspective):



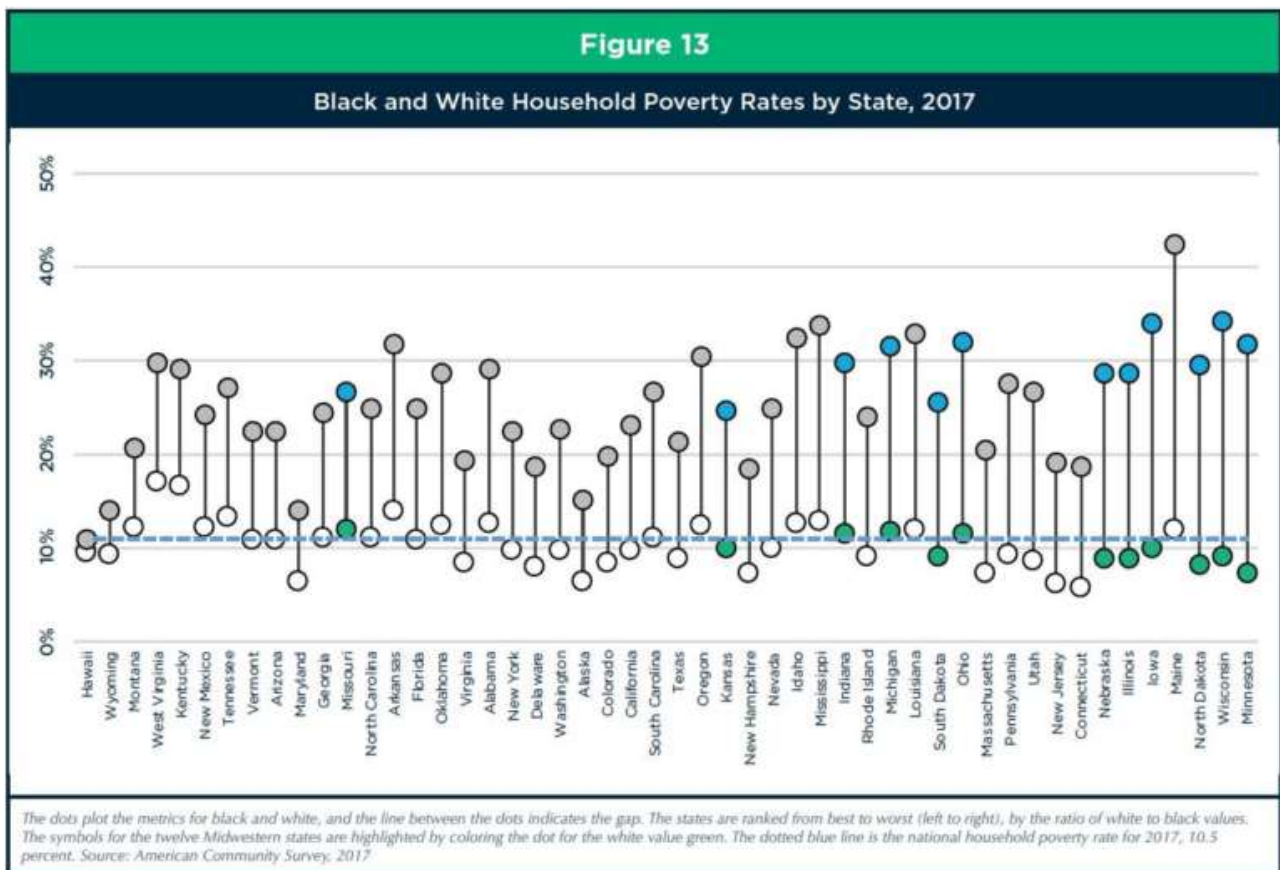
Once you see the crescent, you can't unsee it. And, relatedly, you can't escape the impact on Madison. In 2018 African Americans made up 7% of the population but 43% of arrests and 46% of Dane County Jail inmates; African American students were 18% of the school district, but received 57% of suspensions; 10 percent of African American students received an "advanced" or "proficient" score on the math portion of Wisconsin's standardized testing, while 61% of white students did the same (the proportions were similar for all subjects). Meanwhile the average house price in the Burr Oaks neighborhood (which includes Badger Road) is \$145,300; the Madison average is \$300,967.

The pattern is even worse in Milwaukee, Wisconsin's largest city, and the most segregated city in the country; small wonder that Wisconsin ranks so highly when it comes to the disparity between black and white median household income:

**Figure 12****Black and White Median Household Income by State, 2017**

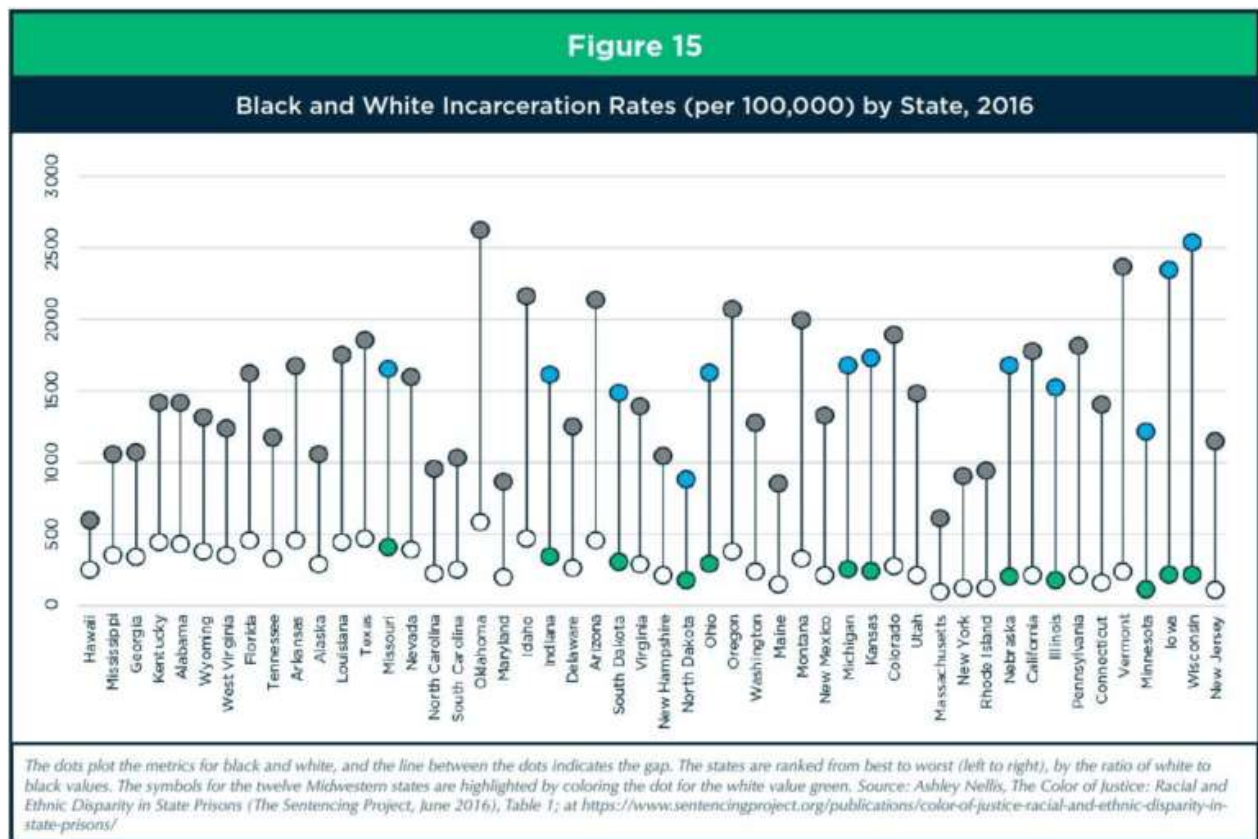
The dots plot the metrics for black and white, and the line between the dots indicates the gap. The states are ranked from best to worst (left to right), by the ratio of white to black values. The symbols for the twelve Midwestern states are highlighted by coloring the dot for the white value green. The dotted blue line is the national median household income for 2017, \$57,652. Source: American Community Survey, 2017

And poverty rates:



And, as the paper from which these charts are drawn puts it:

Racial disparities in rewards and returns (opportunity, compensation, security) are matched by racial disparities in punishment...Five (WI, IA, MN, IL, NE) of the ten worst-performing states, ranked by the ratio between black and white rates, are in the Midwest. Such disparities, glaring in their own right, also have profound impacts on individuals, families and communities. Incarceration short circuits equal citizenship—the right to vote, educational and employment opportunities, access to housing—in deep and lasting ways.



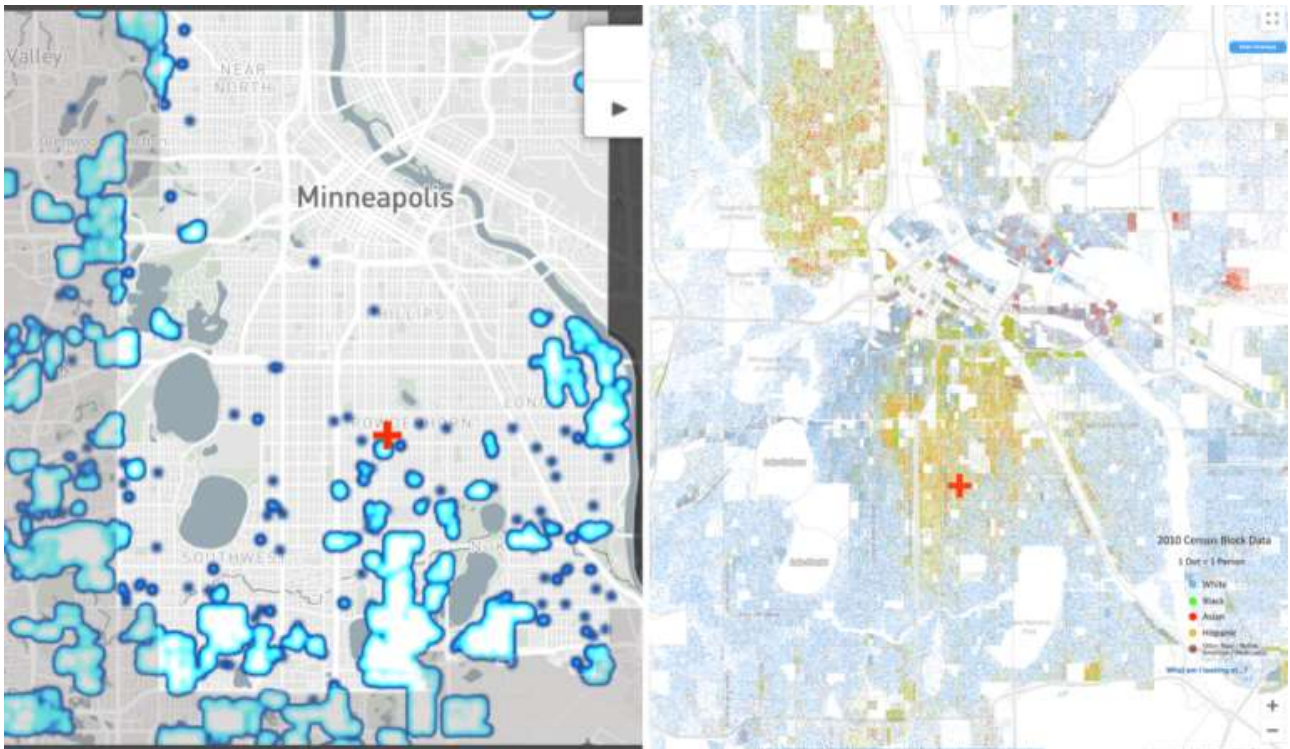
The one state competing with Wisconsin for the highest measurements of disparity is the neighbor to the west: Minnesota.

## Minneapolis and George Floyd

While red-lining helped shape segregation in many cities, Minneapolis was pre-emptive about its discrimination; **beginning in the 1910s** Minneapolis real estate deeds started to include “Covenants” that explicitly excluded African Americans. A **team from the University of Minnesota** has been researching real estate deeds to uncover these covenants, and created this striking time-lapse of their spread:

Loading...

Racial covenants were ruled unconstitutional by the Supreme Court in 1948, but the effect remains; compare the racial covenant map to the racial dot map Blank referenced above — the blue (which is white people) adheres to the blue of racial covenants:



That red cross, meanwhile, is the location of the homicide of George Floyd, in the decidedly non-blue portion of the map. “Homicide” was the word used by the [Hennepin County Medical Examiner](#), which ruled that Floyd’s cause of death was “Cardiopulmonary arrest complicating law enforcement subdual, restraint, and neck compression”; it is up to prosecutors and a jury to decide if that homicide constitutes murder.

The rest of the country did not take so long: nearly all have seen [the video](#) of Minneapolis police officer Derek Chauvin with his knee on Floyd’s neck for 8 minutes and 46 seconds, even as Floyd first complains he cannot breathe, and then, for the final two minutes and 53 seconds, falls silent.

## Dust in the Air

The first version of the Hennepin County Medical Examiner’s autopsy, at least the part quoted in [the criminal complaint against Chauvin](#), read a bit differently:

The Hennepin County Medical Examiner (ME) conducted Mr. Floyd’s autopsy on May 26, 2020. The full report of the ME is pending but the ME has made the following preliminary findings. The autopsy revealed no physical findings that support a diagnosis of traumatic asphyxia or strangulation. Mr. Floyd had underlying health conditions including coronary artery disease and hypertensive heart disease. The combined effects of Mr. Floyd being restrained by the police, his underlying health conditions and any potential intoxicants in his system likely contributed to his death.

The underlying health conditions and intoxicants are still in the final report; what has changed is their relative prominence in explaining Floyd’s death. One suspects that in a different world — say, the world

that was Minneapolis for most of the 20th century — said underlying health conditions and intoxicants would have been held to be the cause of death, not “Other significant conditions.” Perhaps there would be a two paragraph story in the *Star Tribune* on page A17, or more likely Floyd’s death would have disappeared into a police filing cabinet, a non-event as far as most of Minneapolis was concerned. At best there would be a murmur to avoid that sketchy Powderhorn neighborhood, a rarely-visited barely-remembered exception to Minneapolis’ status as one of **the best cities in America**.

Those who knew Floyd or witnessed his death would know better, of course. They would, as Kareem Abdul-Jabbar wrote in the **Los Angeles Times**, shout “Not @#\$\$%! again!” Abdul-Jabbar explains:

African Americans have been living in a burning building for many years, choking on the smoke as the flames burn closer and closer. Racism in America is like dust in the air. It seems invisible — even if you’re choking on it — until you let the sun in. Then you see it’s everywhere. As long as we keep shining that light, we have a chance of cleaning it wherever it lands. But we have to stay vigilant, because it’s always still in the air.

What made the Floyd story different than all of the surely similar examples that went before it is the Internet, specifically the combination of cameras on smartphones and social networks. The former means any incident can be recorded on a whim; the latter means that said recording can be spread worldwide instantly. That is exactly what happened with the Floyd homicide: the initial video was captured on a smartphone and posted on **Facebook**, triggering a level of attention to the Floyd case that in all likelihood changed the nature of the autopsy and led to the pressing of charges against Chauvin — a chance, in Abdul-Jabbar’s words, of cleaning at least one spec of that omnipresent dust.

## Trump’s Tweets

Notably, this is not why Facebook is in the news this week; yesterday hundreds of employees staged a virtual walkout to protest the fact that the company committed, in their mind, a sin of omission: not deleting posts from President Trump. Those posts are copies of Trump tweets, three of which Twitter modified in some way last week. The first two were Trump allegations that voting by mail had a high risk of fraud; **Twitter attached** a “Get the facts” label that led to a page disputing Trump’s claim.

The more serious intervention came early Friday morning, when Twitter obscured a Trump tweet because it, in their determination, “violated the Twitter Rules about glorifying violence.”



Twitter — at least as far as citing its rules is concerned — apparently objected to the phrase “when the looting starts, the shooting starts,” which is associated with a brutal segregationist police chief **from Miami**; Trump claimed to not know the saying’s history, but honestly, arguing about that phrase feels like a distraction from Trump’s all-capitalization use of the descriptor “thugs”, a word with **significant racial undertones**. That certainly seemed to be what the protesting Facebook employees picked up on; from the **New York Times**:

“The hateful rhetoric advocating violence against black demonstrators by the US President does not warrant defense under the guise of freedom of expression,” one Facebook employee wrote in an internal message board, according to a copy of the text viewed by The New York Times. The employee added: “Along with Black employees in the company, and all persons with a moral conscience, I am calling for Mark to immediately take down the President’s post advocating violence, murder and imminent threat against Black people.” The Times agreed to withhold the employee’s name.

What is notable about that New York Times story is that it reproduces the post (and tweet!) that the employees want taken down:



President Trump's post on Facebook.



Mr. Trump's post on Twitter, which the platform modified.

So did the story I linked to above about that Miami police chief, and countless other publications. Indeed, it seems rather obvious that Twitter's action — and those objecting to Facebook's lack of action — ensured that Trump's tweet would be far more widely read than it might have been otherwise.

It is not clear that this is a bad thing. Trump's tweet is abominable, but sadly, of a piece with far too many presidents. [The Associated Press wrote in 2019](#):

Throughout American history, presidents have uttered comments, issued decisions and made public and private moves that critics said were racist, either at the time or in later generations. The presidents did so both before taking office and during their time in the White House...

This extends far beyond the founding fathers, most of whom owned slaves, to the 20th century:

The Virginia-born Woodrow Wilson worked to keep blacks out of Princeton University while serving as that school's president...

Democrat Lyndon Johnson assumed the presidency in 1963 after the assassination of John F. Kennedy and sought to push a civil rights bill amid demonstrations by African Americans...But according to tapes of his private conversations, Johnson routinely used racist epithets to describe African Americans and some blacks he appointed to key positions.

His successor, Republican Richard Nixon, also regularly used racist epithets while in office in private conversations...As with Johnson, many of Nixon's remarks were unknown to the general public until tapes of White House conversations were released decades later. Recently the Nixon Presidential Library released an October 1971 phone conversation between Nixon and then California Gov. Ronald Reagan, another future president...Reagan, in venting his frustration with United Nations delegates who voted against the U.S., dropped some racist language.

The part about secret tapes is notable: much of this racism — this dust in the air — was in darkened rooms, unseen by the public. Trump, if nothing else, has no need for secret tapes: we have his very public Twitter account, and all indications, particularly in terms of pre-COVID polling, suggest that it **massively weakened** his reelection bid.

The president's threats, meanwhile, continue: **yesterday** Trump demanded governors around the country crack down on the looting that has in several cities followed peaceful protests, saying he would call in the military otherwise. That certainly seems to be, in broad strokes, in line with the tweet Twitter hid — does it matter that Trump stated his position on a conference call and in the Rose Garden instead of a tweet?

In fact, that is what is so striking about the demands that Facebook act on this particular post (beyond the extremely problematic prospect of an unaccountable figure like Zuckerberg unilaterally deciding what is and is not acceptable political speech): the preponderance of evidence suggests that these demands have nothing to do with misinformation, but rather reality. The United States really does have a president named Donald Trump who uses extremely problematic terms — in all caps! — for African Americans and quotes segregationist police chiefs, and social media, for better or worse, is ultimately a reflection of humanity. Facebook deleting Trump's post won't change that fact, but it will, at least for a moment, turn out the lights, hiding the dust.

## A Gargantuan Force

It is hard to be optimistic about anything at this moment in time. My regular refrain from the beginning of the coronavirus crisis is that the most likely outcome will be the acceleration of trends that were

already happening. That is particularly scary given what I wrote back when Stratechery started in 2013, in a post called **Friction**:

Count me with those who believe the Internet is on par with the industrial revolution, the full impact of which stretched over centuries. And it wasn't all good. Like today, the industrial revolution included a period of time that saw many lose their jobs and a massive surge in inequality. It also lifted millions of others out of sustenance farming. Then again, it also propagated slavery, particularly in North America. The industrial revolution led to new monetary systems, and it created robber barons. Modern democracies sprouted from the industrial revolution, and so did fascism and communism. The quality of life of millions and millions was unimaginably improved, and millions and millions died in two unimaginably terrible wars.

Another comparison point is the printing press, which I wrote about **last year** in the context of Facebook:

Just as important, though, particularly in terms of the impact on society, is the drastic reduction in fixed costs. Not only can existing publishers reach anyone, anyone can become a publisher. Moreover, they don't even need a publication: social media gives everyone the means to broadcast to the entire world. Read again Zuckerberg's description of the Fifth Estate:

*People having the power to express themselves at scale is a new kind of force in the world — a Fifth Estate alongside the other power structures of society. People no longer have to rely on traditional gatekeepers in politics or media to make their voices heard, and that has important consequences.*

It is difficult to overstate how much of an understatement that is. I just recounted how the printing press effectively overthrew the First Estate, leading to the establishment of nation-states and the creation and empowerment of a new nobility. The implication of overthrowing the Second Estate, via the empowerment of commoners, is almost too radical to imagine.

And yet, look again at this past week: a century of institutionalized racism in Minneapolis was not necessarily overthrown, but certainly overwhelmed in the case of George Floyd, because of a post on Facebook. Both peaceful protests and wanton destruction and looting were likely organized on social media. Video of both were circulated around the world via ubiquitous smartphone cameras on said social networks. The Internet is an amoral force — it can effect both positive and negative outcomes — but what cannot be underestimated is how gargantuan a force it is.

To that end, while there is much to fear, there is room for hope as well. I am grateful that I can no longer unsee Madison's crescent, thanks to a blog post. I am angered by the video of Floyd's death, and appalled at the dust in the air that yes, I was privileged enough to avoid without a second thought. And no matter what upheaval lies ahead, I am certain that the light that illuminates that dust so brightly can never be put away. There are no more gatekeepers, oftentimes for worse, but also for better.

# Platforms in an Aggregator World

Tuesday, May 26, 2020

In the month since I wrote [The Anti-Amazon Alliance](#), there has been two significant announcements from two of the principals in that alliance:

- Shopify announced the Shop app
- Facebook announced Facebook Shops

Many have argued that these announcements are related to each other: Shopify needs to build a customer-facing application in order to take-on Facebook, and the announcement of Facebook Shops is why. My takeaway is the opposite: the inevitability of Facebook Shops — thanks in part to a surprising culprit — is precisely why Shopify should not spend much time on end-user acquisition.

## The Anti-Amazon Alliance Redux

In *The Anti-Amazon Alliance* I noted that brick-and-mortar retail served two functions: discovery and distribution. On the Internet, though, those functions have been split between two different value chains: Facebook is the top-of-funnel for discovery, while Amazon dominates search-driven distribution; the genesis of that article was driven by the news that Google was re-focusing Google Shopping away from pay-to-play to being a search engine that listed products from anyone, thus joining the *Anti-Amazon Alliance*.

One of the most important companies in making both of the non-Amazon value-chains work is Shopify, which is a platform, not an Aggregator. Google and Facebook may collect the customers, but it is Shopify (and WooCommerce, its open-source competitor) that provides the infrastructure for merchants to actually sell things online; Shopify is also working to help merchants get the things they sell into customers hands with the Shopify Fulfillment Network, which I wrote about in [Shopify and the Power of Platforms](#). At last week's [Reunite Conference](#) the company announced that it now owned-and-operated seven warehouses across the U.S., has built an R&D center to improve warehouse operations, and has incorporated [6 River Systems](#) in the form of “Chuck” robots to help with order fulfillment.

One particularly clever component of the Shopify Fulfillment Network is that the same templates a merchant uses to customize their website also translate seamlessly to custom packaging; when you receive a package from a Shopify merchant — the goal is two days, anywhere in the world — it appears as if it came from that merchant, not from Shopify. That's part and parcel of being a platform: you win when those on your platform win, even if customers never even know you exist.

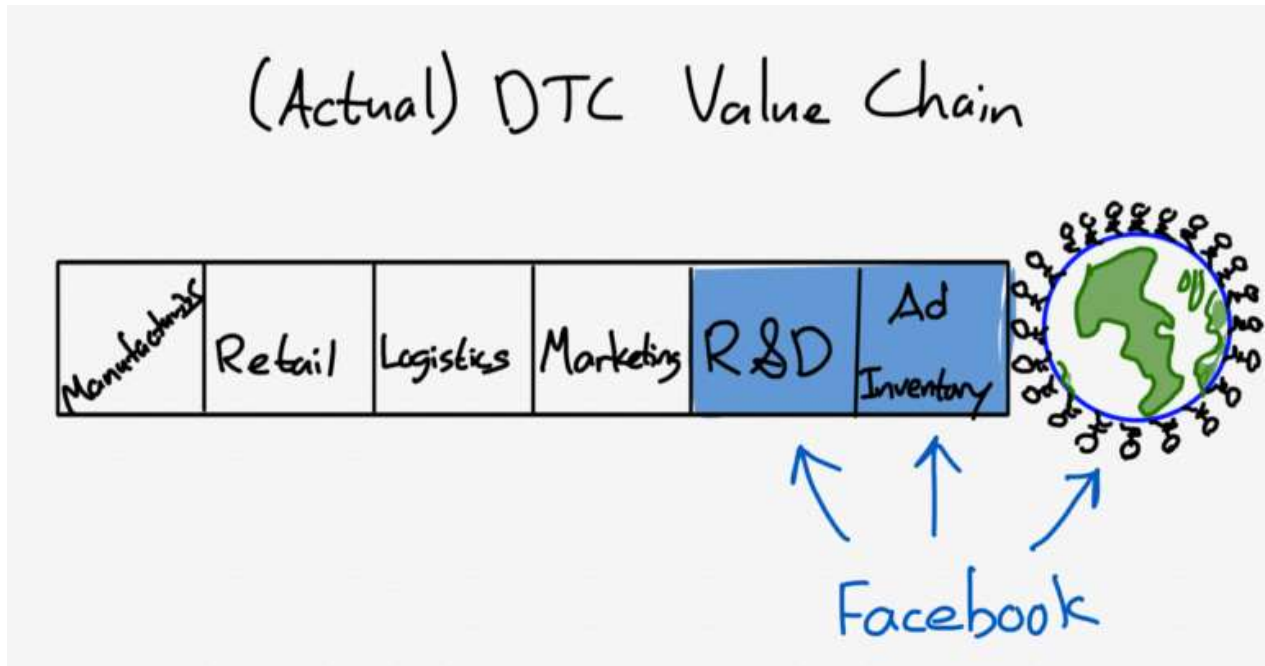
This is also why [I didn't agree](#) with Shopify's decision to rebrand their *Arrive* tracking app as the *Shop* app: it demotes merchants relative to Shopify itself (and the attempt to maintain some sort of brand prominence confuses the user experience).

## Facebook Shops

Those that disagreed with my opinion on the Shop App generally cited Facebook and the fact that the social network is capturing an increasing amount of value from the direct-to-consumer space. I wrote

about why this is happening earlier this year in [Email Addresses and Razor Blades](#):

The problem is that in the process of depending on Google and Facebook for marketing, the DTC companies gave up their planned integration in the value chain, and the associated profits, to Facebook and Google:



The actual integrated players — Google and Facebook — integrate customers and research and development to dominate marketing; DTC may have online retail operations, but that is a modularized — and thus commoditized — part of the value chain (and meanwhile, Amazon was in the process of integrating retail and logistics).

However, suggesting that Shopify take responsibility for directly acquiring customers (as opposed to building tools to help merchants do so) is not only out of step with Shopify's position in the value chain, but there is also no particularly good reason to believe that Shopify will be any better at this than their merchants. What Facebook is capable of in terms of customer acquisition is not trivial — that is precisely why they are able to charge a premium for it.

What is more concerning for Shopify — at least at first glance — is the possibility of Facebook backwards integrating, a la Facebook Shops. From the [Facebook Newsroom](#):

Facebook Shops make it easy for businesses to set up a single online store for customers to access on both Facebook and Instagram. Creating a Facebook Shop is free and simple. Businesses can choose the products they want to feature from their catalog and then customize the look and feel of their shop with a cover image and accent colors that showcase their brand. This means any seller, no matter their size or budget, can bring their business online and connect with customers wherever and whenever it's convenient for them.

People can find Facebook Shops on a business' Facebook Page or Instagram profile, or discover them through stories or ads. From there, you can browse the full collection, save products you're interested in

and place an order — either on the business’ website or without leaving the app if the business has enabled checkout in the US.

And just like when you’re in a physical store and need to ask someone for help, in Facebook Shops you’ll be able to message a business through WhatsApp, Messenger or Instagram Direct to ask questions, get support, track deliveries and more. And in the future, you’ll be able to view a business’ shop and make purchases right within a chat in WhatsApp, Messenger or Instagram Direct.

Bad news for Shopify, right? Except note the last paragraph:

We’re also working more closely with partners like Shopify, BigCommerce, WooCommerce, ChannelAdvisor, CedCommerce, Cafe24, Tienda Nube and Feedonomics to give small businesses the support they need. These organizations offer powerful tools to help entrepreneurs start and run their businesses and move online. Now they’ll help small businesses build and grow their Facebook Shops and use our other commerce tools.

While Shopify did [write a blog post about Facebook Shops](#), the partnership merited a mere 18 seconds in the Reunited Keynote; this doesn’t feel like something Shopify is super thrilled about, and for understandable reasons.<sup>1</sup>

## Shopify’s Business Model

Shopify’s original business model is labeled “Subscription Solutions”; merchants pay a subscription fee to use the Shopify platform — the price [ranges from \\$29 to \\$299 per month](#) — and can use the payment provider of their choice. When Shopify IPO’d *Subscription Solutions* was 60% of their \$67 million in quarterly revenue.

Over the last five years, though, “Merchant Solutions” — which are a percentage of transactions, usually from using Shopify Payments — has been the primary growth driver. Last quarter it was *Merchant Solutions* that was 60% of Shopify’s \$470 million in quarterly revenue.

This shift in Shopify’s business model is almost certainly why the company appears to be uncomfortable with this evolution of their “partnership” with Facebook. Sure, Facebook Shop integration will be a feature of Shopify’s *Subscription Solutions*, but Shopify will be locked out of sales made via Facebook Checkout, which means no *Merchant Solutions* revenue, and by extension, no participating in the upside of its merchants’ growth.

What makes this a particularly bitter pill to swallow for Shopify is that Facebook’s move is really good for those merchants (and a far better solution than the overly complicated Instagram Shopping beta that didn’t integrate with outside services). At this point I’m going to assume that those of you who have bought a product from an Instagram ad outnumber those of you who haven’t,<sup>2</sup> and the truth is that it is a pretty janky experience tediously filling in all of your shipping and billing details. Sometimes it is easier to just abandon your purchase for that item you had no idea existed 30 seconds ago.

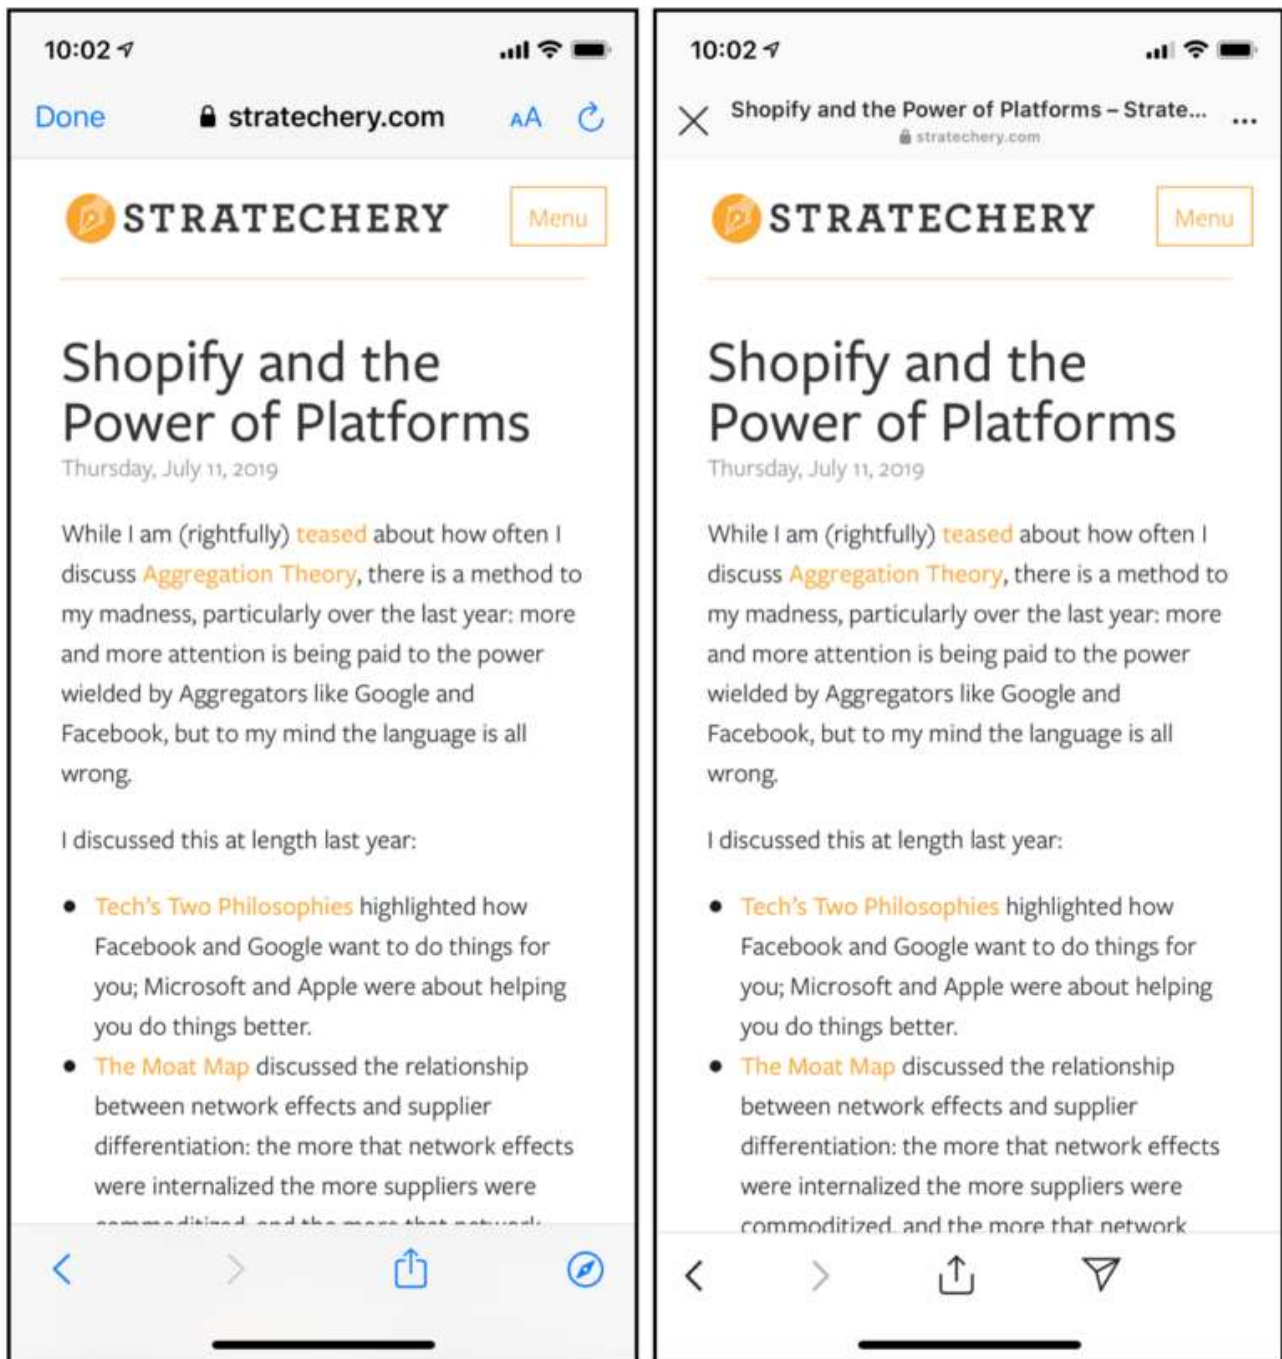
Facebook Checkout, on the other hand, would mean a purchase that is completed in seconds, not minutes. After all, if anyone knows who you are, it is Facebook!3 That is going to mean a lot more conversions on Facebook and Instagram ads in particular — and Shopify isn't going to see an incremental cent.

## Firefighters and Arsonists

It is very important to note that Facebook is setting itself up as the solution for a problem it played a major role in creating. To examine iOS specifically, Apple provides a user-friendly solution to the payment problem: Apple Pay. Apple Pay, though, **is limited** to Safari and

**SFSafariViewController** objects; the latter is the most straightforward way for developers to add a browser to their application — it looks like Safari, and has a button in the lower-right corner to load the page in Safari.

Compare and contrast an **SFSafariViewController** webview (in this case, from Twitter's app) to what you get in Instagram:



What Instagram<sup>4</sup> has done is basically build their own browser within Instagram. It still uses iOS's WebKit engine — that is an iOS requirement, which means that Chrome uses WebKit as well — but it is not otherwise controlled by Apple. And that, Apple has decreed, means no Apple Pay, which means that the purchase experience is worse in Instagram and Facebook than it could have been had either Facebook used `SFSafariViewController` or if Apple relaxed its policies around Apple Pay.

Notably, both Facebook and Apple are motivated by opposite sides of the same coin: Facebook uses its own browser because that allows it to capture far more data than they could from `SFSafariViewController`; Apple, meanwhile, puts a sandbox around `SFSafariViewController` for the purpose of keeping user data away from 3rd party developers, and incentivizes usage by both making `SFSafariViewController` far easier to use and also including benefits like Apple Pay integration.

The conclusion I am sure many of you have already arrived at is to be grateful for Apple's policies and irritated at Facebook's. After all, Apple is just looking out for the user, and Facebook is looking to exploit them, right?

In fact, it's a bit more complicated than that, and cookies are a good example. Remember that a whole bunch of those Instagram advertisements are from Shopify merchants, all of whose websites are hosted on the same infrastructure. Imagine if Shopify could set a cookie when a user made a purchase on one website, and then when the user visited another Shopify website they were already logged in with all of their payment information ready to go. Notably, Shopify has built this infrastructure — it's called **Shop Pay** — but you have to log in on every distinct Shopify website.

The reason for this rigamarole is that for the last three years Apple has been leading the charge towards the elimination of 3rd-party cookies of the sort I just described. There are good reasons for this, specifically the elimination of 3rd-party trackers primarily loaded by advertisements. There are, though, real casualties along the way, including Shopify and its merchants.

I wrote about these tradeoffs last year in **Privacy Fundamentalism**:

Technology can be used for both good things and bad things, but in the haste to highlight the bad, it is easy to be oblivious to the good. Manjoo, for example, works for the New York Times, which makes most of its revenue from subscriptions; given that, I'm going to assume they do not object to my including 3rd-party resources on Stratechery that support my own subscription business?

This applies to every part of my stack: because information is so easily spread across the Internet via infrastructure maintained by countless companies for their own positive economic outcome, I can write this Article from my home and you can read it in yours. That this isn't even surprising is a testament to the degree to which we take the Internet for granted: any site in the world is accessible by anyone from anywhere, because the Internet makes moving data free and easy.

This, unfortunately, better described the world of 2017 than the world of 2020. Yes, websites are still freely accessible, and thank goodness for that, but it is ever more difficult for those websites to be supported by critical infrastructure like Shopify (or, in my case, Stripe and WordPress). Privacy is a good thing, but so is entrepreneurship and competition; maximizing one without any consideration of the others leads to unintended outcomes.

Facebook Shops is a perfect example: it is going to succeed because it is good for Shopify's merchants, but the reason it is good for Shopify's merchants is because Facebook and Apple effectively teamed up to make it impossible for Shopify to fix the payment problem on their own.

This makes me sad: the part of the Internet that fills me with the most optimism are platforms like Shopify and Stripe and Substack that make it possible for individual entrepreneurs to try and build their own businesses with world-class tools at their disposal; making it hard to utilize those tools primarily benefits established companies and apps.<sup>5</sup>

## Shopify's Platform Prospects

This is the harsh reality for Shopify and anyone else competing in a value chain with an Aggregator: you are not going to beat Facebook at their own game, particularly given the technical limitations increasingly facing infrastructure providers. At the end of the day overcoming Facebook's skill at acquiring customers must be the responsibility of the merchants themselves: what works — and what will always work, as long as the web exists — is creating something so compelling that people will go to you directly, and yes, fill out a payment form.

This, though, is why I remain so excited about the Shopify Fulfillment Network. I'm not completely sold on Shopify's approach — I thought it might make more sense to try and create a common interface for merchants to interact with independent 3PL providers, as opposed to building out Shopify own logistics service (but I could very well be wrong about this) — but I absolutely endorse the company making massive investments in this space.

First, this is a service that no merchant can build on its own; it is a perfect example of how a platform can create something for an ecosystem that would not exist otherwise. Moreover, this is specifically what is needed to fulfill the promise I laid out last year, of Shopify as an Amazon competitor, not because users choose to go to Shopify, but because they have no reason to know that Shopify exists.

Second, this is a service that Facebook Shops is going to make more valuable, not less: all of those merchants increasing sales thanks to Facebook Checkout still need to ship their things, and there is zero chance that Facebook ever integrates into the real world.

Third, this is a service no one else is going to build. Yes, it is very difficult to build fulfillment centers and develop robots and employ lots of workers, but within that difficulty is an escape from competition, i.e. a far more reliable way to make sustainable profits in the long run.

I am also bullish about Shopify **moving into banking and financial services**; yes, there is certainly more competition here with the likes of Stripe — another Shopify partner<sup>6</sup>, but this is another area where taking on the sorts of risks that an Aggregator never would (and never should) makes sense for a platform provider.

## A Reason to Build

In **my response** to Marc Andreessen's **It's Time to Build** essay, I suggested three different ways the tech industry could make more of a difference in the world of atoms, not just bits; the second was as follows:

Second, invest in real-world companies that differentiate investment in hardware with software. This hardware could be machines for factories, or factories themselves; it could be new types of transportation, or defense systems. The possibilities, at least once you let go of the requirement for 90% gross margins, are endless.

That reminded me of these two “warnings” in Shopify's **most recent earnings report**:

- Our expectation is that the gross margin percentage of merchant solutions will decline in the short term as we develop Shopify Fulfillment Network and 6 River Systems Inc. (“6RS”).
- Our expectation is that the continued growth of merchant solutions may cause a decline in our overall gross margin percentage.

In short, building fulfillment centers and developing robots and employing lots of workers isn’t great for margins. It is, though, at least in the ultra-competitive markets that Shopify operates in, good for building a moat. This is certainly something that Amazon discovered long ago; I wrote in 2018 in [Amazon Go and the Future](#):<sup>7</sup>

This willingness to spend is what truly differentiates Amazon, and the payoffs are tremendous. I mentioned telecom companies in passing above: their economic power flows directly from massive amounts of capital spending; said power is limited by a lack of differentiation. Amazon, though, having started with a software-based horizontal model and network-based differentiation, has not only started to build out its vertical stack but has spent massive amounts of money to do so. That spending is painful in the short-term — which is why most software companies avoid it — but it provides a massive moat.

In this view, the upside of building in the real world is rooted in just how difficult it is; given that we have reached [The End of the Beginning](#), it may be the highest upside available.

- 
1. This paragraph originally said that Shopify didn’t write a blog post; I apologize for the error [↩]
  2. Please don’t tell me if you haven’t 🤖 [↩]
  3. In all seriousness, the company was honest that [they will use Facebook Shop behavior for ad targeting](#), which lends credence to their promise to keep purchase methods private. [↩]
  4. And Facebook, but not WhatsApp, which simply loads links in Safari [↩]
  5. This, by the way, suggests that Substack was smart to keep everyone on the same domain; I generally believe that you should always have your own domain, but the truth is that Apple’s overzealous cookie policy has increasingly made that a liability [↩]
  6. Interestingly, the wording in Shopify’s annual report about their relationship has shifted in the last year, although not dramatically [↩]
  7. That article, by the way, [wrongly predicted](#) that Amazon would not license Amazon Go, and also [mischaracterized Marx’s view of automation](#). The latter hurts more. [↩]
- 

## Chips and Geopolitics

Tuesday, May 19, 2020

The debate around who belongs on the Mount Rushmore of tech would be a long one; what is certain is that Morris Chang should be on the list. He certainly leads the way in terms of impact relative to name recognition.

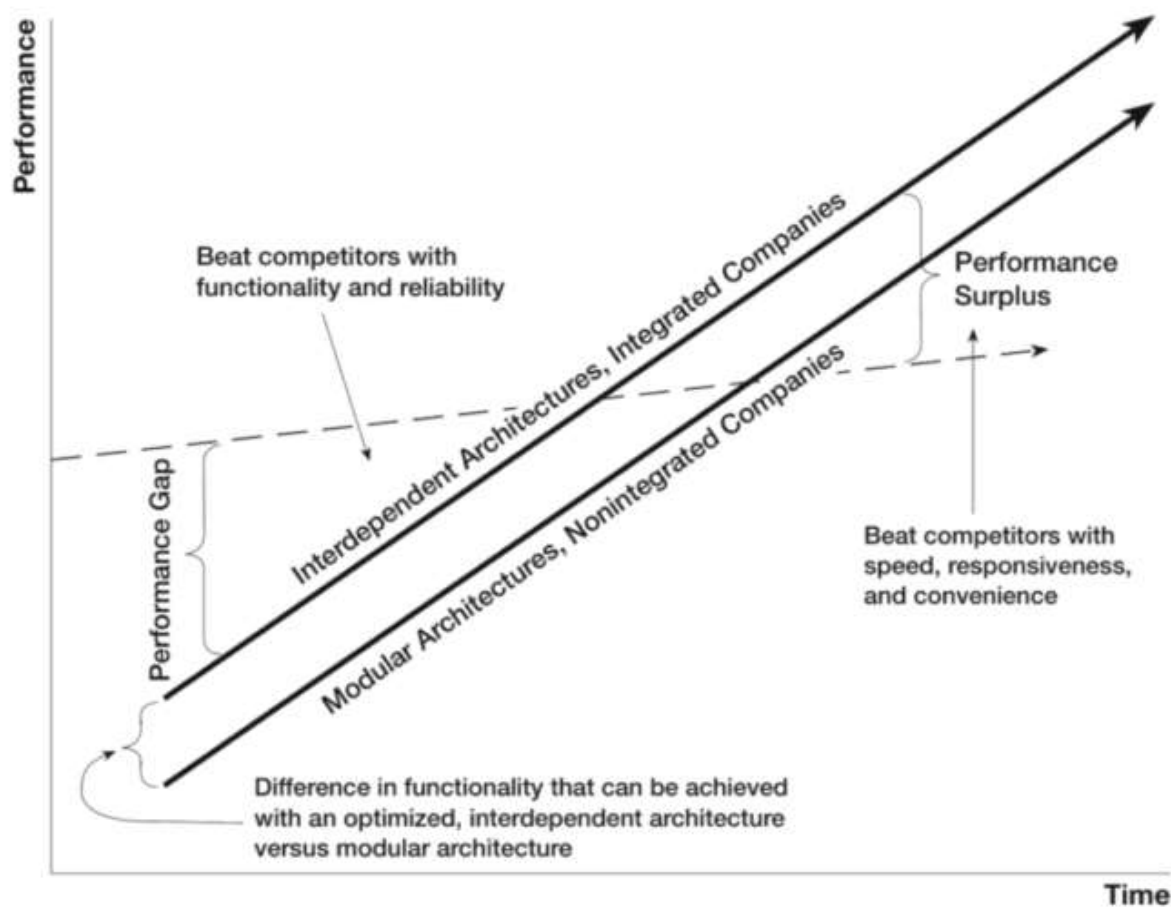
## Integration and Modularization

Clayton Christensen, in 2003's *The Innovator's Solution*, explained how the natural course of industries was from interdependent architectures to modular ones:

Customers will not buy your product unless it solves an important problem for them. But what constitutes a “solution” differs across the two circumstances in Figure 5-1: whether products are not good enough or are more than good enough. The advantage, we have found, goes to integration when products are not good enough, and to outsourcing — or specialization and dis-integration when products are more than good enough.

**FIGURE 5 - 1**

### Product Architectures and Integration



The left side of Figure 5-1 indicates that when there is a performance gap — when product functionality and reliability are not yet good enough to address the needs of customers in a given tier of the market — companies must compete by making the best possible products. In the race to do this, firms that build their products around proprietary, interdependent architectures enjoy an important competitive advantage against competitors whose product architectures are modular, because the standardization inherent in modularity takes too many degrees of design freedom away from engineers, and they cannot not optimize performance.

This makes intuitive sense: optimizing everything results in better performance, at the cost of long-term reliability and flexibility. Sure, long-term reliability and flexibility are nice-to-have, but they are lesser priorities. Once that top priority is met, though, these secondary priorities come to the forefront.

Overshooting does not mean that customers will no longer pay for improvements. It just means that the type of improvement for which they will pay a premium price will change. Once their requirements for functionality and reliability have been met, customers begin to redefine what is not good enough. What becomes not good enough is that customers can't get exactly what they want exactly when they need it, as conveniently as possible. Customers become willing to pay premium prices for improved performance along this new trajectory of innovation in speed, convenience, and customization. When this happens, we say that the basis of competition in a tier of the market has changed.

This is a big problem for firms that are dominant in a market undergoing this transition; after all, the reason said firms are dominant is because they are the highest performing, which is to say that they are highly integrated, and to unwind said integration is usually untenable for both business model and more deep-rooted cultural reasons. That opens the door for new entrants:

The pressure of competing along this new trajectory of improvement forces a gradual evolution in product architecture, as depicted in Figure 5-1 — away from the interdependent, proprietary architectures that had the advantage in the not-good-enough era toward modular designs in the era of performance surplus. Modular architectures help companies to compete on the dimensions that matter in the lower-right portions of the disruption diagram. Companies can introduce new products faster because they can upgrade individual subsystems without having to redesign everything. Although standard interfaces invariably force compromises in system performance, firms have the slack to trade away some performance with these customers because functionality is more than good enough.

Modularity has a profound impact on industry structure because it enables independent, nonintegrated organizations to sell, buy, and assemble components and subsystems. Whereas in the interdependent world you had to make all of the key elements of the system in order to make any of them, in a modular world you can prosper by outsourcing or by supplying just one element. Ultimately, the specifications for modular interfaces will coalesce as industry standards. When that happens, companies can mix and match components from best-of-breed suppliers in order to respond conveniently to the specific needs of individual customers.

Taiwan Semiconductor Manufacturing Company (TSMC), the company Chang founded in 1987, is arguably the single best example of the process Christensen described.

## Intel and TSMC

Intel invented the microprocessor in 1971, and for decades to come, it was not good enough. The 4-bit Intel 4004 was followed by the 8-bit Intel 8008, and then the Intel 8080. Then, in 1978, came the Intel 8086, a 16-bit processor that was backwards compatible with programs written for the 8080 and 8008. That was followed by the Intel 80286, and in 1985, the 32-bit Intel 80386. It was the 80386 that defined the baseline x86 instruction set that undergirds modern processors in most laptops, desktops, and

servers, but x86 has its roots in the 8008. Intel, by integrating design, manufacture, and software from the 1970s, would go on to define and dominate the processor market for decades.

It would take a very long time for this integrated approach to overshoot the market. Intel's 80386 was succeeded by the 80486, then the Pentium, and every release made computers so much faster that use cases unimaginable only one or two years prior suddenly seemed within reach, if only Intel could continue its rate of improvement. And, to the company's credit — and with a solid push from AMD into a 64-bit variant that retained backwards compatibility to the 80386 — Intel did just that.

Still, Intel made general purpose processors; processors that were created for a specific task would be much faster, at least in theory, but it was hard to get started: Chang, then a long-time executive at Texas Instruments, observed in the 1980s that it cost \$50~\$100 million dollars to start a new chip company, primarily because of the cost of manufacturing. You could contract production from Intel or Texas Instruments or Motorola, but it wasn't reliable — and they were also your competitor!

A few years later, in 1987, Chang was invited home to Taiwan, and asked to put together a business plan for a new government initiative to create a semiconductor industry. Chang **explained in an interview with the Computer History Museum** that he didn't have much to work with:

I paused to try to examine what we have got in Taiwan. And my conclusion was that [we had] very little. We had no strength in research and development, or very little anyway. We had no strength in circuit design, IC product design. We had little strength in sales and marketing, and we had almost no strength in intellectual property. The only possible strength that Taiwan had, and even that was a potential one, not an obvious one, was semiconductor manufacturing, wafer manufacturing. And so what kind of company would you create to fit that strength and avoid all the other weaknesses? The answer was pure-play foundry...

In choosing the pure-play foundry mode, I managed to exploit, perhaps, the only strength that Taiwan had, and managed to avoid a lot of the other weaknesses. Now, however, there was one problem with the pure-play foundry model and it could be a fatal problem which was, "Where's the market?"

What happened is exactly what Christensen would describe several years later: TSMC created the market by "enabl[ing] independent, nonintegrated organizations to sell, buy, and assemble components and subsystems." Specifically, Chang made it possible for chip designers to start their own companies:

When I was at TI and General Instrument, I saw a lot of IC [*Integrated Circuit*] designers wanting to leave and set up their own business, but the only thing, or the biggest thing that stopped them from leaving those companies was that they couldn't raise enough money to form their own company. Because at that time, it was thought that every company needed manufacturing, needed wafer manufacturing, and that was the most capital intensive part of a semiconductor company, of an IC company. And I saw all those people wanting to leave, but being stopped by the lack of ability to raise a lot of money to build a wafer fab. So I thought that maybe TSMC, a pure-play foundry, could remedy that. And as a result of us being able to remedy that then those designers would successfully form their own companies, and they will become our customers, and they will constitute a stable and growing market for us.

It worked. Graphics processors were an early example: Nvidia was started in 1993 with only \$20 million, and never owned its own fab.<sup>1</sup> Qualcomm, after losing millions manufacturing its earliest designs, spun off its chip-making unit in 2001 to concentrate on design, and Apple started building its own chips without a fab a decade later. Today there are thousands of chip designers in all kinds of niches creating specialized chips for everything from appliances to fighter jets, and none of them have their own foundry.

There was one other thing that happened along the way: as I detailed in [2018's Intel and the Danger of Integration](#), TSMC eventually surpassed Intel in not just flexibility but also pure performance:

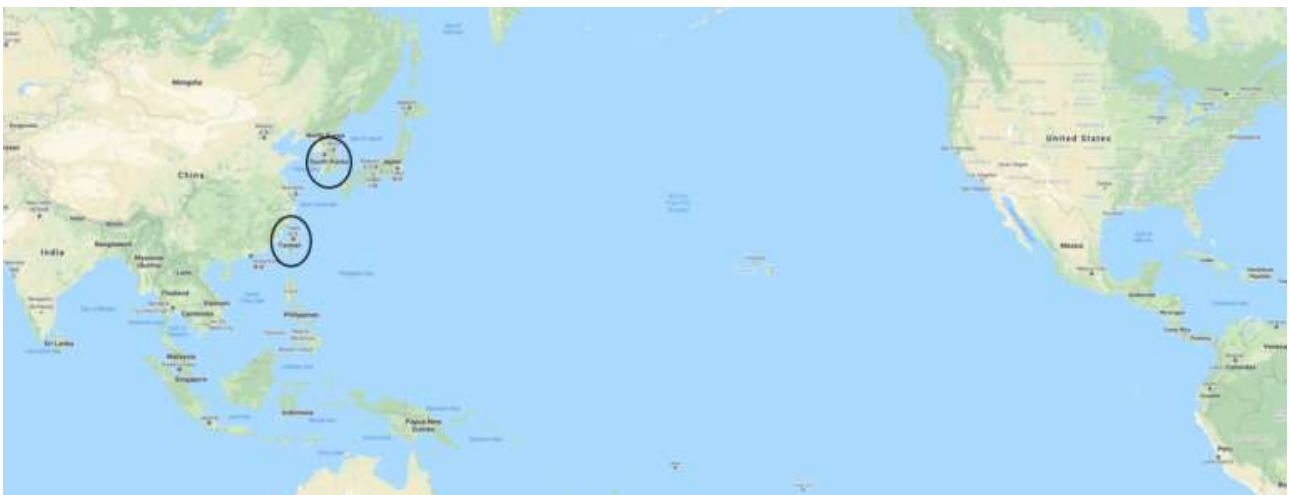
In time, though, TSMC got better, in large part because it had no choice: soon its manufacturing capabilities were only one step behind industry standards, and within a decade had caught-up (although Intel remained ahead of everyone). Meanwhile, the fact that TSMC existed created the conditions for an explosion in “fabless” chip companies that focused on nothing but design...the increased business let TSMC invest even more in its manufacturing capabilities.

In short, TSMC is the best chipmaker in the world, no matter what vector of performance you care about. And with that came an entirely new class of problems, not just for TSMC, but also Taiwan.

## Geopolitical Concerns

The international status of Taiwan is, as they say, complicated. So, for that matter, are U.S.-China relations. These two things can and do overlap to make entirely new, even more complicated complications.

Geography is much more straightforward:



Taiwan, you will note, is just off the coast of China. South Korea, home to Samsung, which also makes the highest end chips, although mostly for its own use, is just as close. The United States, meanwhile, is on the other side of the Pacific Ocean. There are advanced foundries in Oregon, New Mexico, and Arizona, but they are operated by Intel, and Intel makes chips for its own integrated use cases only.

The reason this matters is because chips matter for many use cases outside of PCs and servers — Intel’s focus — which is to say that TSMC matters. Nearly every piece of equipment these days, military or

otherwise, has a processor inside. Some of these don't require particularly high performance, and can be manufactured by fabs built years ago all over the U.S. and across the world; others, though, require the most advanced processes, which means they must be manufactured in Taiwan by TSMC.

This is a big problem if you are a U.S. military planner. Your job is not to figure out if there will ever be a war between the U.S. and China, but to plan for an eventuality you hope never occurs. And in that planning the fact that TSMC's foundries — and Samsung's — are within easy reach of Chinese missiles is a major issue.

China, meanwhile, is investing heavily in catching up, although Semiconductor Manufacturing International Corporation (SMIC), its Shanghai-based champion, only just started manufacturing on a 14nm process, years after TSMC, Samsung, and Intel. In the long run, though, the U.S. faced a scenario where China had its own chip supplier, even as it threatened the U.S.'s chip supply chain.

## TSMC's Announcement

This was the context for last week's announcement that TSMC is building a fab in the United States. From the [Wall Street Journal](#):

Taiwan Semiconductor Manufacturing Co., the world's largest contract manufacturer of silicon chips, said Friday it would spend \$12 billion to build a chip factory in Arizona, as U.S. concerns grow about dependence on Asia for the critical technology. TSMC said the project, disclosed earlier Thursday by The Wall Street Journal, has the support of the federal government and the state of Arizona. It comes as the Trump administration has sought to jump-start development of new chip factories in the U.S. due to rising fears about the U.S.'s heavy reliance on Taiwan, China and South Korea to produce microelectronics and other key technologies.

TSMC made the decision to go ahead with the project at a board meeting on Tuesday in Taiwan, according to people familiar with the matter, adding that both the State and Commerce Departments are involved in the plans. Construction will begin next year with production targeted for 2024, the company said in a statement. TSMC's new plant would make chips branded as having 5-nanometer transistors, the tiniest, fastest and most power-efficient ones manufactured today. TSMC just started rolling out 5-nanometer chips at a factory in Taiwan in recent months. TSMC said the plant would make 20,000 wafers a month, making it a relatively small facility for a company that made more than 12 million wafers last year alone. TSMC's Fab 18 in Taiwan, which currently produces its 5-nanometer chips, was targeted for 100,000 wafers a month when it broke ground in 2018.

First off, while this announcement has superficial similarities to the star-crossed Foxconn factory in Wisconsin, that project reeked of political theater from the start, and, more pertinently, never made much sense for anyone involved. The current outcome — [empty innovation centers and a still-unfinished factory](#) that has already been re-purposed — was frankly the default outcome.

This TSMC project is different for several reasons. First, you don't halfway build a foundry; TSMC is either in for billions, or they're in for nothing. Second, it seems clear that the federal government is contributing significantly to the cost. And third, that is exactly what the federal government should do, because the national security implications are real.

This does raise the question about just how committed TSMC is to this project. As the Wall Street Journal notes, the Arizona fab is quite small, relatively speaking, and while 5-nanometer chips are top-of-the-line today, they won't be in 2024, when the fab opens. Moreover, it is worth noting that TSMC has a fab in Washington that it opened in 1998; it still operates, but TSMC didn't make any additional investments in the U.S. until now.

I think, though, that this is an overly pessimistic reading of this news, at least from a U.S. perspective. First off, of course TSMC is going to start small, and with technology it has already figured out how to build. It is one thing to build a massive "gigafab" next door to the ones you have already built in Taiwan, even as your best employees, who have pushed TSMC to the top over the last thirty years, figure out the next processing node; it is quite another to attempt something similar across the ocean.

What is a much bigger deal, though, is that the Taiwan of 2020 is not last in line when it comes to processor technology, but first, and the government — which retains a significant ownership stake in TSMC — has been committed to keeping TSMC's best technology in Taiwan. That this move is happening at all suggests the sort of momentous choice not simply on TSMC's part but also Taiwan's that is hard to undo: when it comes to the U.S. and China, ambiguously sitting in the middle, selling to both, was no longer an option.

## Lessons for Tech

There are three big lessons for tech specifically and America broadly in this news.

First, while we **learned in 2016** that technology was inseparable from domestic politics, the lesson in 2020 should be that technology is inseparable from geopolitics. It is chips that gave Silicon Valley its name, and everything about this chip decision is about geopolitics, not economics.

Second, at some point every tech company is going to have to make a choice between the U.S. and China. It is tempting to blame the tension between the two countries on Trump, but the truth is that China, particularly under Xi Jinping, has been significantly hardening its rhetoric and actions since before Trump was elected, and has been committed to not just catching but surpassing the U.S. in technology **for years**. There is a **fundamental clash of values** between the West and China, and it is clear that **China is interested in exporting theirs**. At some point everyone will be stuck in the middle, like TSMC, and **Switzerland won't be an option**.

Third, Intel, **much like Compaq**, is an allegory for where the U.S. seems to have lost its way. Locked in an endless pursuit of efficiency and shareholder value, the U.S. gave up its flexibility and resiliency in favor of top-end performance. Intel is one of the most advanced chip makers in the world, but it turns out that capability is far too constrained to its own needs to be of general applicability. Worse, to the extent Intel was willing to become a contract manufacturer, it **wanted the federal government to pay for it**, the better to satisfy shareholders. The government, rightly, in my mind, chose an operator that was actually used to operating in the world as it is, not once was.

At the same time, TSMC's justifiable carefulness in building a U.S. fab gives Intel an opportunity. Back in 2013, in one of the first Stratechery articles, **I urged the company** to embrace manufacturing and give up its integration, margins be damned. Intel specifically, and the U.S. generally, would be in far better

shape had they acted then. As the saying goes, though, the second best time to start is now — and that applies not only to Intel, which should spend the money to get into contract manufacturing on its own, but also to the U.S. The world has changed, and it's time to act accordingly.

---

1. The very first Nvidia chips were manufactured by SGS-Thomson Microelectronics, but have been manufactured by TSMC from the original GeForce on [[↔](#)]
- 

## Exponent Podcast: Open, Free, and Spotify

Friday, May 15, 2020

On Exponent, the weekly podcast I host with [James Allworth](#), we discuss [Dithering and Open Versus Free](#).

Listen to it [here](#).

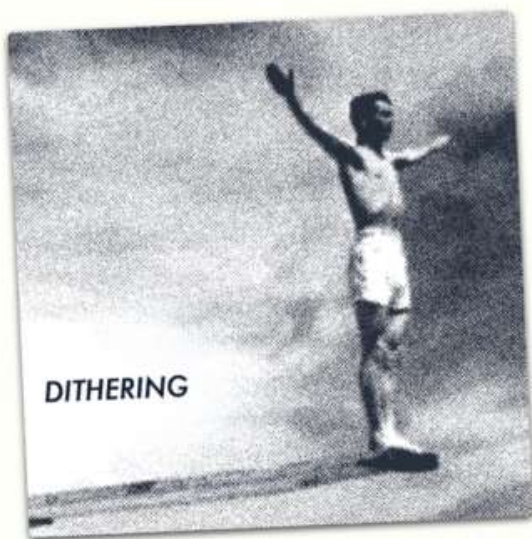
---

## Dithering and Open Versus Free

Tuesday, May 12, 2020

The big news, in case you haven't yet heard: [John Gruber](#) and I have launched a new podcast called [Dithering](#):

# DITHERING



A new podcast from **Ben Thompson** and **John Gruber**. Three episodes per week, 15 minutes per episode. Not a minute less, not a minute more.

Dithering costs \$5/month or \$50/year. If you're a Stratechery subscriber, it costs \$3/month or \$30/year to [add it as a bundle](#).<sup>1</sup>

Dithering covers some of the same topics as Stratechery — the [Dithering web page](#) has a descriptive lists of topics — but in the conversational style that many of you have enjoyed [on my appearances on Gruber's The Talk Show podcast](#), or in the Daily Update Interview I did with Gruber [last Thursday](#).<sup>2</sup> Expect less in-depth analysis than a typical Stratechery post, and more back-and-forth, with the occasional foray into non-tech topics. All in fifteen minutes, exactly. It's perfect for your dishwashing commute!

That time limit is certainly a challenge (that is why we recorded 20 episodes before we launched — the entire back catalog is available to subscribers), but we really wanted to experiment with what a podcast might be. We purposely don't have show notes or much of a web page, and we have created evocative cover art embedded in each episode's MP3, because the canonical version of Dithering is in your podcast player. This is as pure a podcast as can be — and that means open, even if it isn't free.

## Open != Free

I'm used to dealing with the seeming contradiction between *open* and *free*: [back in 2014](#) I started selling an email I called the [Daily Update](#). There was no special app required, and while Daily Updates were [archived](#) on the web, I took care to not shove a paywall in your face; if you wanted more content from me you could pay for more, and I would send you an email over the open SMTP protocol, that landed in the email client you already used.

This combination of *open* and *for-pay* turned out to be extraordinarily powerful: even as *closed* but *free* feeds like Facebook were turning into pay-to-play for publishers, email remained the only feed that everyone checked every day that didn't have a gatekeeper, which made it the best possible means of delivering the value proposition I was charging for — a proposition I most clearly defined in 2017's [The Local News Business Model](#):

It is very important to clearly define what a subscriptions means. First, it's not a donation: it is asking a customer to pay money for a product. What, then, is the product? It is not, in fact, any one article (a point that is missed by the misguided focus on micro-transactions). Rather, a subscriber is paying for the regular delivery of well-defined value.

The importance of this distinction stems directly from the economics involved: the marginal cost of any one Stratechery article is \$0. After all, it is simply text on a screen, a few bits flipped in a costless arrangement. It makes about as much sense to sell those bit-flipping configurations as it does to sell, say, an MP3, costlessly copied.

So you need to sell something different.

In the case of MP3s, what the music industry finally learned — after years of kicking and screaming about how terribly unfair it was that people “stole” their music, which didn't actually make sense because digital goods are [non-rivalrous](#) — is that they should sell convenience. If streaming music is free on a marginal cost basis, why not deliver all of the music to all of the customers for a monthly fee?

This is the same idea behind nearly every large consumer-facing web service: Netflix, YouTube, Facebook, Google, etc. are all predicated on the idea that content is free to deliver, and consumers should have access to as much as possible. Of course how they monetize that convenience differs: Netflix has subscriptions, while Google, YouTube, and Facebook deliver ads (the latter two also leverage the fact that content is free to create). None of them, though, sells discrete digital goods. It just doesn't make sense.

## Aggregators Versus Publishers

This model is pretty good for consumers: they get access to an abundance of content for a set price. It's great for the Aggregators: because they have so many consumers, the suppliers of content are forced to accede to the Aggregator's terms, even as Aggregators are best placed to serve advertisers. That is another way of saying that it is the individual content maker that is getting the short end of the stick:

- On Spotify, individual artists make fractions of a cent per play, and their payout is based on their share of all Spotify plays; if you have a super-fan that listens to nothing but your songs, you still only get a few pennies.
- On Netflix, show creators are getting bigger payments up front, but in return they are giving up residuals and international rights; Netflix owns all of the upside.
- YouTube is actually one of the more creator-friendly Aggregators: what you earn is pretty closely tied to how many views you achieve. That, though, means a hamster wheel lifestyle of constantly churning out content and begging for subscribers, even as it requires ever more views to achieve the same amount of money. And, of course, YouTube could de-monetize you at any time, for any reason.
- Google helps consumers find content, but because (all of the) consumers start with search, so do advertisers; Facebook lets consumers make content, and then favors it over professionally produced links.

It is important to note that, the constant griping of traditional gatekeepers notwithstanding, Aggregators are by definition good for most content creators; after all, *everyone* is now a content creator, whereas previously publishing was reserved for those who had access to physical assets like printing presses, recording studios, or broadcast towers. That means most people are **publishing for the first time** (with effects **both good and bad**).

It also means that traditional publishers face more competition for attention, and, as long as they rely on Aggregators, an inherently unstable source of income: one big song, show, video, or article can make some money, but without an ongoing connection and commitment from the consumer to the content creator, it is increasingly impossible to make a living.

## Subscriptions and Open Protocols

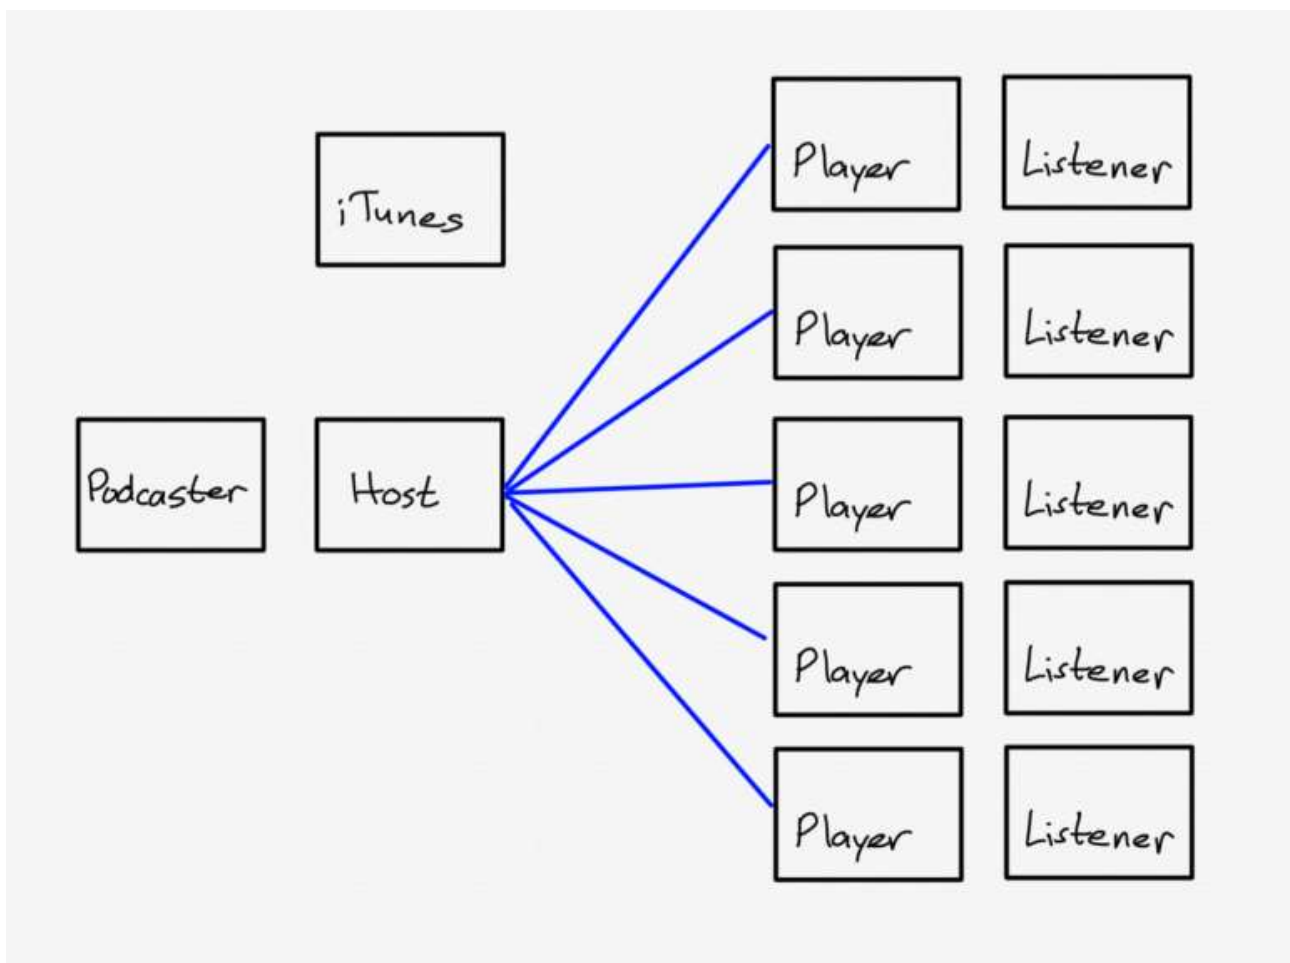
This is why subscriptions — “paying for the regular delivery of well-defined value” — are so important. I defined every part of that phrase:

- *Paying*: A subscription is an ongoing commitment to the production of content, not a one-off payment for one piece of content that catches the eye.

- *Regular Delivery*: A subscriber does not need to depend on the random discovery of content; said content can be delivered to the subscriber directly, whether that be email, a bookmark, or an app.
- *Well-defined Value*: A subscriber needs to know what they are paying for, and it needs to be worth it.

This runs in the opposite direction of a Spotify-type model, even as it takes advantage of the same foundation of zero marginal costs. If an email is an artifact of hard work creating something people are interested in, the open ecosystem of HTTP and SMTP drives the costs of delivering that artifact to zero. There is no massive streaming infrastructure to build, nor endless data centers in the cloud — this can all be rented for not much money at all — which means that the cost structure of an independent creator can be dramatically lower than any traditional publisher, even as **their addressable market is the same size**.

HTTP and SMTP, though, are not the only open protocols available to publishers: RSS is another, and it is the foundation of the podcast ecosystem. Most don't understand that podcasts are not hosted by Apple, but rather that iTunes is a directory of RSS feeds hosted on servers all over the Internet. When you add a podcast to your podcast player, you are simply adding an RSS feed that includes information about the show, and a link for where to download new episodes.



This, if you squint, looks a lot like email: create something that listeners find valuable on an ongoing basis, and deliver it into a feed they already check, i.e. their existing podcast player. That is *Dithering*: while you have to pay to get **a feed customized to you**, that feed can be put in your favorite podcast app, which means *Dithering* fits in with the existing open ecosystem, instead of trying to supplant it.

Well, almost all podcast apps: Spotify is an exception.<sup>3</sup>

## Spotify's Facebook Play

Podcasting, as I wrote last year, looks a lot like the early web; iTunes is the Yahoo directory, and advertising is punching the monkey:

The current state of podcast advertising is a situation not so different from the early web: how many people remember this?



These ads were elaborate affiliate marketing schemes; you really could get a free iPod if you signed up for several credit cards, a Netflix account, subscription video courses, you get the idea. What all of these marketers had in common was an anticipation that new customers would have large lifetime values, justifying large payouts to whatever dodgy companies managed to sign them up.

The parallels to podcasting should be obvious: why is Squarespace on seemingly every podcast? Because customers paying monthly for a website have huge lifetime values. Sure, they may only set up the website once, but they are likely to maintain it for a very long time, particularly if they grabbed a “free” domain along the way. This makes the hassle of coordinating ad reads and sponsorship codes across a plethora of podcasts worth the trouble; it’s the same story with other prominent podcast sponsors like ZipRecruiter or SimpliSafe.

Some are content with this state of affairs, and I understand the sentiment: the early web, annoying banner ads notwithstanding, was in many respects a nicer place as well, and some folks even made money. The problem, though, is it didn’t last: once Google and Facebook figured out that the best way to advertise was to aggregate users and deliver targeted ads, the open web withered; it is only in the last few years that, thanks to email, independent publishing is making a return.

I strongly believe that podcasting is approaching a similar precipice; I wrote a year ago that Spotify wants to be the podcast Aggregator:

What I think Spotify senses, though, is that while podcasts, at least in theory, solve many of their business model problems, Spotify is also uniquely positioned to solve the problems of many podcasters/suppliers. To wit:

- Increasing advertising revenue for the entire industry requires a centralized player that can leverage a large userbase. Spotify is still a distant second to Apple in podcasts, but they are growing fast. Just as

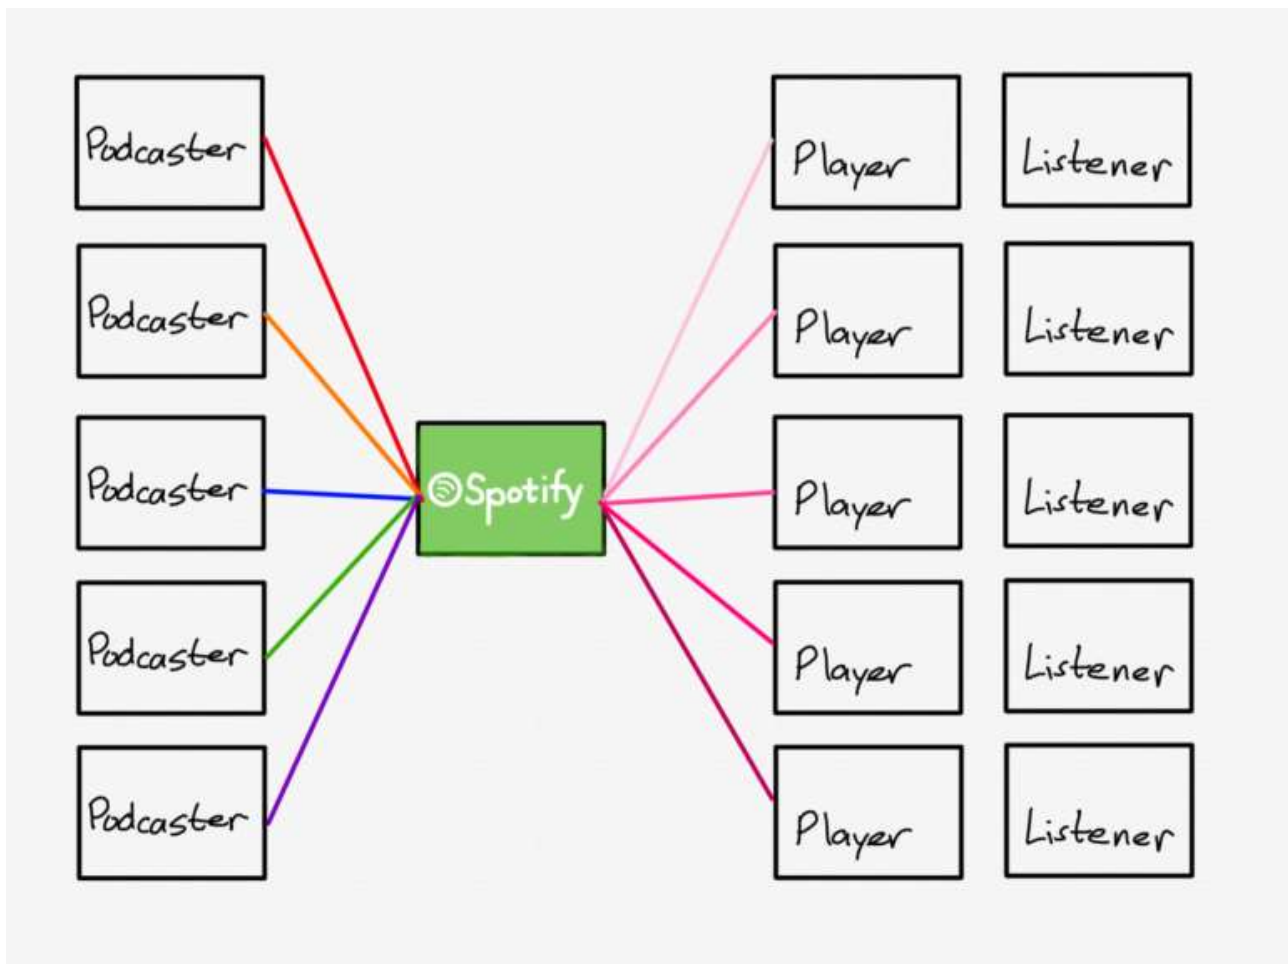
importantly, Spotify already has a strongly growing advertising business — again, larger than the entire podcast market — that it can extend to podcasts.

- The open nature of podcasts means it is very difficult to monetize users directly; Spotify, though, has already built an entire infrastructure around monetizing users directly. Podcasts exclusive to Spotify can likely make meaningful money from Spotify subscribers that still gives Spotify far higher margin than music.

Spotify CEO Daniel Ek made clear that this was the goal after the company acquired The Ringer; from the Q1 2020 [earnings call](#):

When we look at the overall opportunity, it is pretty clear that we haven't added Internet-level monetization yet to audio. So, all the things that you've come to expect in video and display in terms of measurability, in terms of just targeting, a lot of that is lacking in podcasts today. And you've seen it time and time again. As you add those capabilities, you generally can raise CPMs across the board, because advertisers feel more certain about the results that they're getting. And if we do that, that's going to be a tremendous benefit for all the podcasting creators, but it's also going to be a tremendous benefit for Spotify.

This is why Spotify [adjusted its accounting last quarter](#) to recognize the cost of its owned-and-operated podcast production as an expense for its advertising business; the company isn't primarily focused on acquiring paying subscribers, although that is a nice side effect of increased engagement on Spotify. The real goal is to intermediate podcasters and listeners and take over podcast advertising just like Facebook and Google took over web advertising.



That, though, is bad for openness — indeed, Spotify isn't open at all. You can't simply add an RSS feed to Spotify, as you can most other podcast players. Rather, podcasters have to submit their feeds to Spotify and agree to the service's **terms of service**, which can be changed at any time at Spotify's sole discretion. Sure, the terms are relatively benign today; they could include the right to insert advertising tomorrow. Even if that doesn't happen, though, Spotify still is not open: they can take down your content or choose not to play it, just as Facebook could not show your page unless you were willing to pay-to-play.

This is where, as I noted when **I launched the Daily Update podcast**, it is important to distinguish between my role as analyst, podcaster, and publisher:

Analyst Ben says it is a good idea for Spotify to try and be the Facebook of podcasting...Writer/Podcaster Ben certainly sees the allure: having my podcast available to Spotify's 271 million monthly active users would be great; for that matter, having this Daily Update read by everyone I could reach on Facebook would be great as well. I've already put in the work, why not reach everyone? Indeed, were I supported by advertising, that would be the imperative.

Publisher Ben, though, remembers that my business model is predicated on a higher average revenue per user (thanks to subscriptions), not a higher number of users; that means making tradeoffs, and foregoing wide reach is one of them. That, by extension, means not agreeing to Spotify's terms for Exponent, and accepting that leveraging RSS to have per-subscriber feeds makes having the Daily Update Podcast on Spotify literally impossible. More broadly, owning my own destiny as a publisher means avoiding Aggregators and connecting directly with customers.

*Dithering* is another effort driven by Publishers Ben and John; if we are to maintain a thriving podcast ecosystem that is open, we must figure out monetization, and from my perspective, that means subscriptions. The fact that Spotify won't even allow *Dithering* to be played on their app only increases the urgency: if the choice is free and closed versus for-pay and open I will always push for the latter — three times a week, 15 minutes per episode.

Some additional notes on *Dithering* and the service powering it:

- Yes, there are several companies building out paid podcasting services. None of them had all of the features I wanted, including full control over hosting, the ability to create bundles, and customized feeds based on subscription status; to paraphrase [Alan Kay](#), I'm serious about figuring out how paid podcasts might be a sustainable business not only for me but for the entire ecosystem, and that meant building my own software to maximize experimentation.
- My good friends at [Model Rocket](#) did most of the actual work on the podcast service; they are, if I may use the term, rock stars, and you haven't heard the last of our collaboration. [Brad Ellis](#) (who designed the Stratechery logo), created *Dithering*'s look-and-feel in collaboration with Gruber.
- The *Stratechery* + *Dithering* bundle originally launched with both podcasts on the same feed; I think there is tremendous potential in this approach, but a positive user experience would require podcast players to adopt a new standard we proposed. That seems unlikely given that...
- It is frustrating the degree to which many players don't abide by the [current podcast standard](#). Few apps, for example, respect the `<itunes:image>` tag, which shows per-episode artwork in the feed (and Apple's own Podcasts app doesn't even show MP3 artwork).
- To that end, last night we pushed a new update that splits *Stratechery* and *Dithering* into two feeds, even if you subscribe to the bundle. Visit the [podcast management page](#) to add the independent *Dithering* feed.

Finally, *Stratechery* and the Daily Update are not going anywhere. Indeed, I am more inspired than ever — building something is nice in that way. Also, don't miss [Gruber's post about Dithering](#), and note that both [Exponent](#) and [The Talk Show](#) remain as free podcasts. Free is fine! — but we should not forget that open is more important.

- 
1. Unfortunately due to the limitations of our membership software, we can't offer monthly subscriptions to annual subscribers; if you are an annual subscriber and add on *Dithering*, and realize you don't like it, we will refund you your remaining 11 months [[↔](#)]
  2. That interview is free-to-listen-to even if you aren't a Daily Update subscriber; create an account and add a Stratechery feed [here](#). [[↔](#)]
  3. And, to be fair, Google Podcasts and Stitcher; this analysis applies to them as well [[↔](#)]
- 

[← Older Posts](#)

On the business, strategy, and impact of technology.

© Stratechery LLC 2020 | [Terms of Service](#) | [Privacy Policy](#)