2018 ESSAY COLLECTION

andrewchen.co

CONTENTS

What’s next for marketplace startups? Reinventing the $10 trillion service economy, that’s what.

[Dear readers, this essay is on the future of marketplaces. Is there still room for marketplace startups to innovate? We answer, emphatically, yes! Am excited to share a vision on the past and future of the service economy, in a collaboration by my a16z colleague [Li Jin](). From “Unbundling Craiglist” to “Uber for X” – we lay it all out in a single framework. Hope you enjoy our thinking! -A]

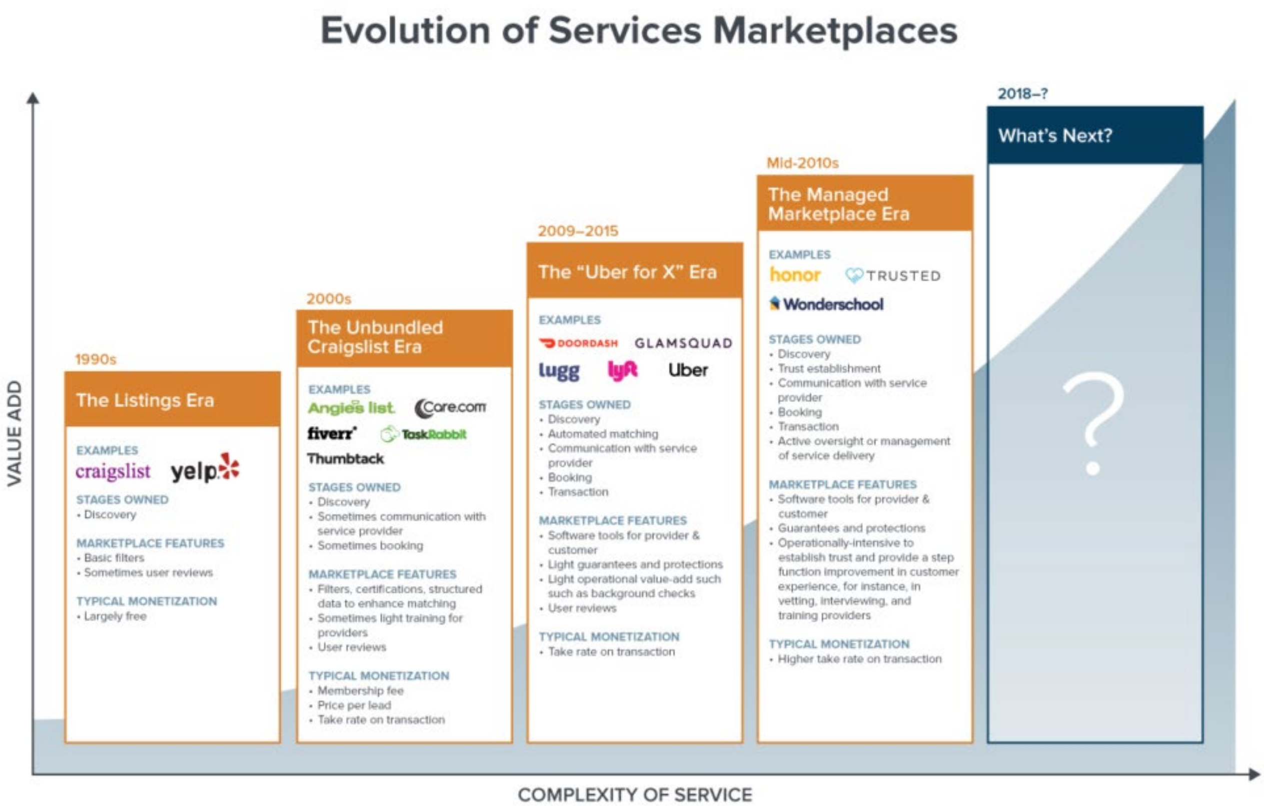

Above: 4 eras of marketplaces focused on the service economy – and what’s next

Goods versus Services – why a breakthrough is coming

Marketplace startups have done incredibly well over the first few decades of the internet, reinventing the way we shop for goods, but have been less successful services. In this essay, we argue that a breakthrough is on its way: While the first phase of the internet has been about creating marketplaces for goods, the next phase will be about reinventing the service economy. Startups will build on the lessons and tactics to crack the toughest service industries – including regulated markets that have withstood digital transformation for decades. In doing this, the lives of 125 million Americans who work in the services-providing industries will join the digital transformation of the economy.

In the past twenty years, we've transformed the way people buy goods online, and in the process created Amazon, eBay, JD.com, Alibaba, and other e-commerce giants, accounting for trillions of dollars in market capitalization. The next era will do the same to the $9.7 trillion US consumer service economy, through discontinuous innovations in AI and automation, new marketplace paradigms, and overcoming regulatory capture.

The service economy lags behind: while services make up 69% of national consumer spending, the Bureau of Economic Analysis estimated that just 7% of services were primarily digital, meaning they utilized internet to conduct transactions.

We propose that a new age of service marketplaces will emerge, driven by unlocking more complex services, including services that are regulated. In this essay, we'll talk about:

- ***Why services are still primarily offline***
- ***The history of service marketplace paradigms***
  - *The Listings Era*
  - *The Unbundled Craigslist Era*
  - *The "Uber for X" Era*
  - *The Managed Marketplace Era*
- ***The future of service marketplaces***
  - *Regulated services*
  - *Five strategies for unlocking supply in regulated markets*
- ***Future opportunities***

Let's start by looking at where the service economy is right now and why it's resisted a full scale transformation by software.

Get essays about marketplaces and more by [subscribing to my newsletter.]

Software eating the service economy, but it's been slow

We've all had the experience of asking friends for recommendations for a great service provider, whether it be a great childcare provider, doctor, or hair stylist. Why is that? Why aren't we discovering and consuming these services in the same digital way we've come to expect for goods?

Despite the rise of services in the overall economy, there are a few reasons why services have lagged behind goods in terms of coming online:

- ***Services are complex and diverse, making it challenging to capture relevant information in an online marketplace***
- ***Success and quality in services is subjective***
- ***Fragmentation – small service providers lack the tools or time to come online***
- ***Real-world interaction is at the heart of services delivery, which makes it hard to disaggregate parts of a purchase that might be done online***

Let's unpack each reason below:

First, on the complexity and diversity of services, services are performed by providers who vary widely, unlike goods which are manufactured to a certain spec. Even the names of services can vary: what one home cleaning service calls a "deep clean" can be different from another provider's definition. This lack of standardization makes it difficult for a service marketplace to capture and organize the relevant information.

Second, services are often complex interactions without a clear yardstick of success or quality. The customer experience of a service is often subjective, making traditional marketplace features like reviews, recommendations, and personalization more difficult to implement. Sometimes just getting the job completed (as in rideshare) is sufficient to earn a 5-star review, whereas other higher-stakes services, like childcare, have complex customer value functions, including safety, friendliness, communicativeness, rapport with child, and other subjective measures of success.

Third, small service providers often lack the tools or time to come online. In many service industries, providers are small business owners with low margins; contrast this with goods manufacturing where there are economies of scale in production, and thus consolidation into large consumer products companies. As a result of industry fragmentation, service providers often don't have time or budget to devote to key business functions, such as responding to customer requests, promoting and marketing themselves, maintaining a website, and other core functions. While major e-commerce platforms have taken on the role of distribution, merchandising, and fulfilling orders for goods, there are few platforms that service providers can plug into to manage their businesses and reach customers.

Fourth, real-world interaction is central to services, which can pull other steps of the services funnel into the offline world as well. Many services are produced and consumed simultaneously in real-world interactions, whereas goods entail independent stages of production, distribution, and consumption. The various stages of the goods value chain can be easily unbundled, with e-commerce marketplaces comprising the discovery, transaction, and fulfillment steps. Conversely, since the production and consumption of services usually occur simultaneously offline, the discovery, distribution, and transaction pieces are also often integrated into the offline experience.

For instance, since getting a haircut entails going to a salon and having interactions with the providers there, the stages of the value chain that precede and follow that interaction (discovery, booking, and payment) also often get incorporated into the in-person experience.

All of these factors make it very hard for services to come online as comprehensively and widely as commerce – but there's hope. We've seen multiple eras of bringing the service economy online, and we're on the verge of a breakthrough!

The 4 eras of Service Marketplaces, and what's next

There have been 4 major generations of service marketplaces, but coverage of services and providers remains spotty, and many don't provide end-to-end, seamless consumer experiences. Let's zoom out and talk through each historical marketplace paradigm, and what we've learned so far.

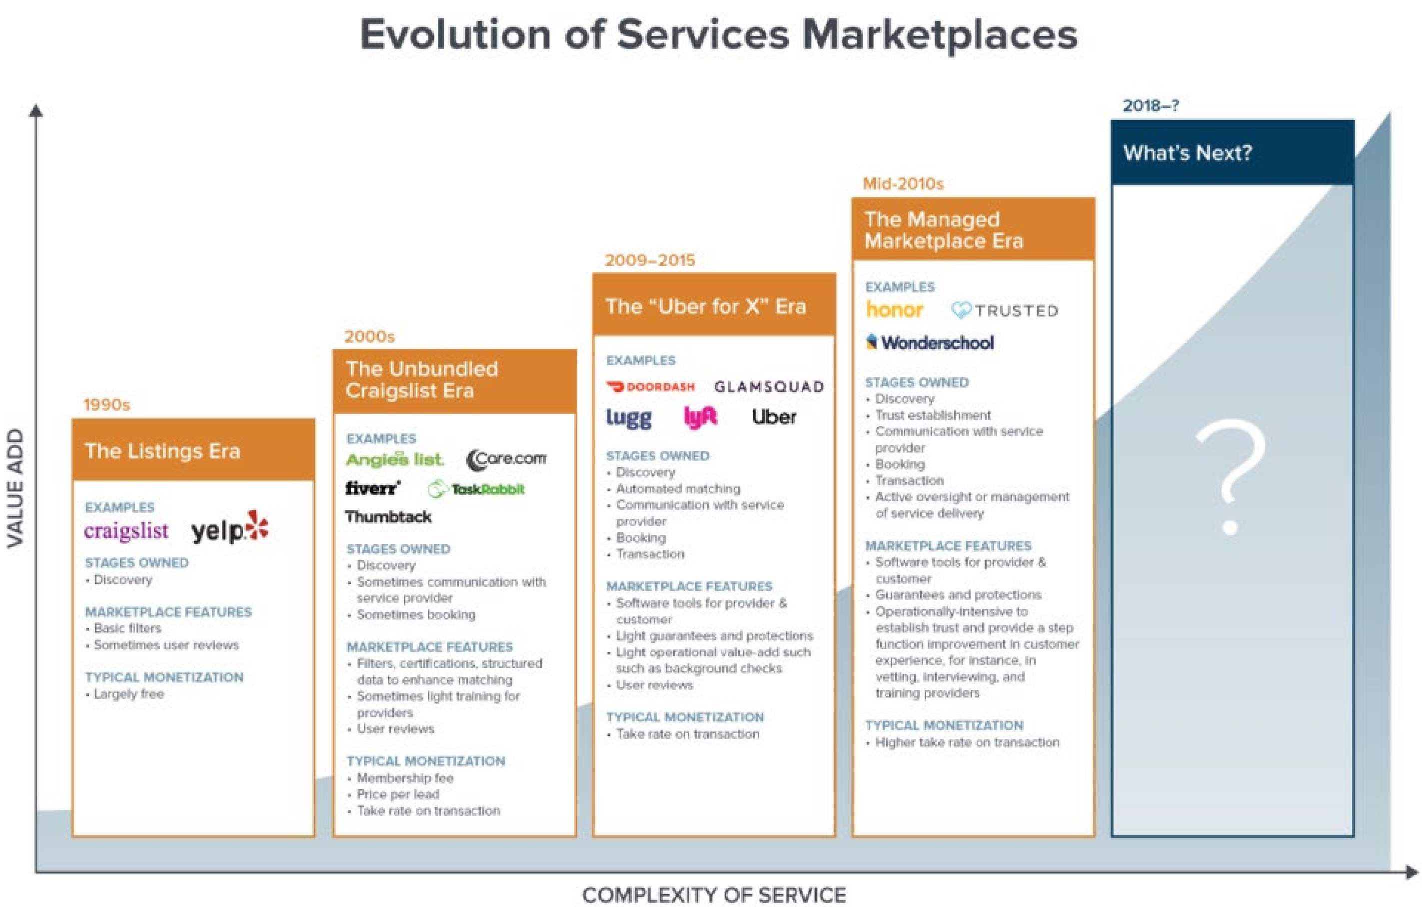

Above, you can see that there have roughly been four major eras of marketplace innovation when it comes to the service economy.

1. The Listings Era (1990s)

The first iteration of bringing services online involved unmanaged horizontal marketplaces, essentially listing platforms that helped demand search for supply and vice versa. These marketplaces were the digital version of the Yellow Pages, enabling visibility into which service providers existed, but placing the onus on the user to assess providers, contact them, arrange times to meet, and transact. The dynamic here is "caveat emptor"–users assume the responsibility of vetting their counterparties and establishing trust, and there's little in the way of platform standards, protections, or guarantees.

Craigslist's Services category is the archetypal unmanaged service marketplace. It includes a jumble of house remodeling, painting, carpet cleaners, wedding photographers, and other services. But limited tech functionality means that it feels disorganized and hard to navigate, and there's no way to transact or contact the provider without moving off the platform.

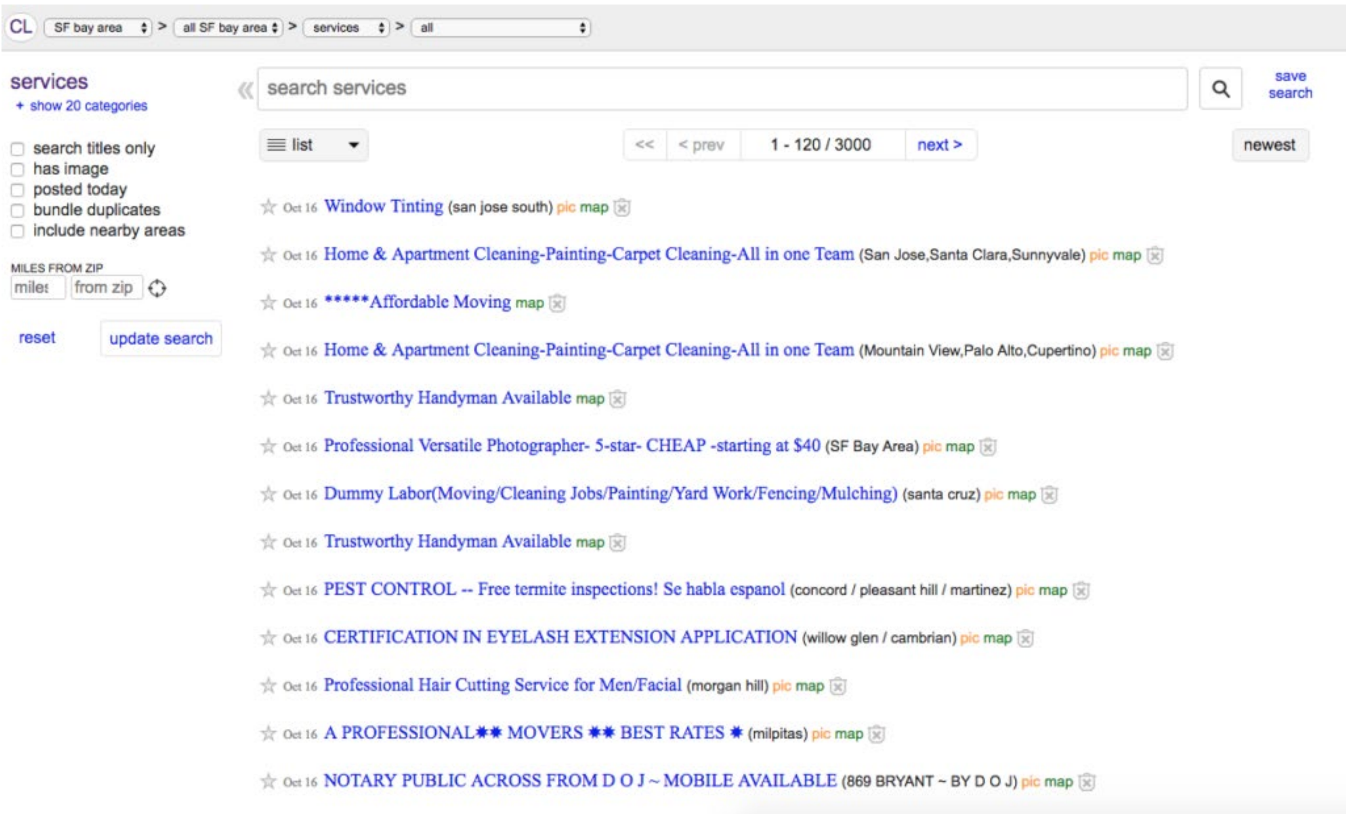

We've all had the experience of a listings-oriented product, like Craigslist. You find something you want, but everything else – trust/reviews/payments/etc – that's all up to you!

2. The Unbundled Craigslist Era (2000s)

Companies iterated on the horizontal marketplace model by focusing on a specific sub-vertical, enabling them to offer features tailored to a specific industry. We've all seen the diagram of various companies picking off Craigslist verticals – it looks something like this:

As a reaction to the "Wild West" nature of Craigslist, to improve the customer experience, each startup would create value-add via software. For instance, Care.com carves off the Childcare section of Craigslist, and provides tech value-add in the form of filters, structured information, and other features to improve the customer experience of finding a local caregiver. It's a huge leap in terms of user experience over Craigslist's Childcare section.

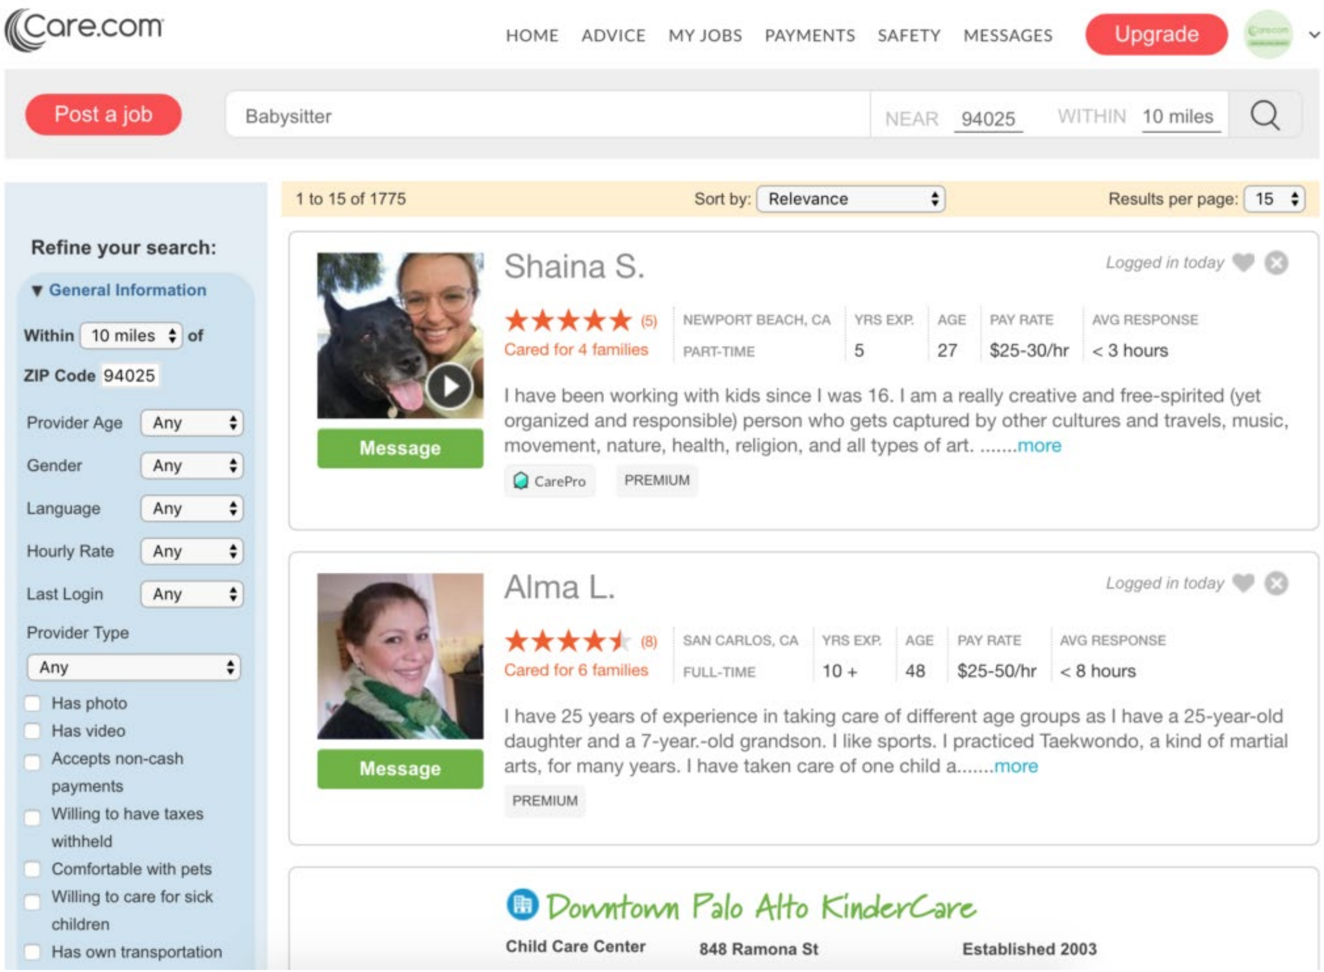

Angie's List, a home services site founded in 2005, carves off Craigslist's household services category. The platform has features including reviews, profiles, certified providers, and an online quote submission process. But the marketplace doesn't encompass the entire end-to-end experience: users turn to Angie's List for discovery, but still need to message or call providers and coordinate offline.

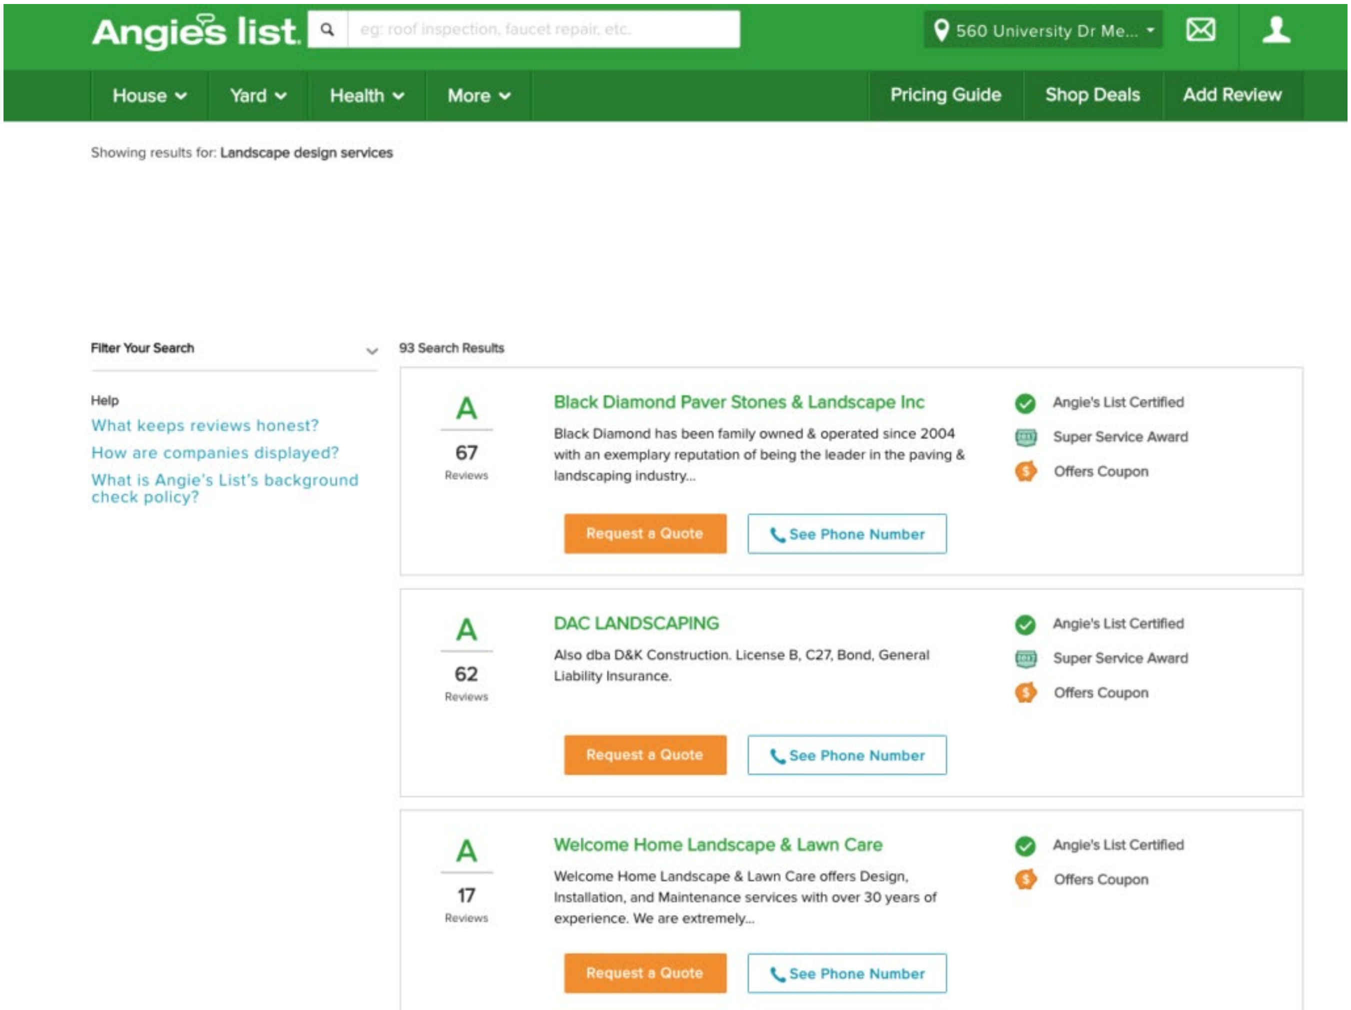

Unmanaged vertical marketplaces like Angie's List go a step beyond Craigslist and take on some value-add services like certifying providers when they meet certain standards, but customers still need to select and contact the service provider, place their trust in the provider rather than the platform, and transact offline.

Like previous listing sites, these platforms in this era try to use the 'wisdom of the crowds' to promote trust. These platforms have a network effect in that more reviews means more users and more reviews. But user reviews have their limitations, as every user has a unique value function that they're judging a service against. Without standardized moderation or curation, and without machine learning to automate this process, customers have the onus of sifting through countless reviews and selecting among thousands of providers.

3. The “Uber for X” Era (2009-)

In the early 2010s, a wave of on-demand marketplaces for simple services arose, including transportation, food delivery, and valet parking. These marketplaces were enabled by widespread mobile adoption, making it possible to book a service or accept a job with the tap of a button.

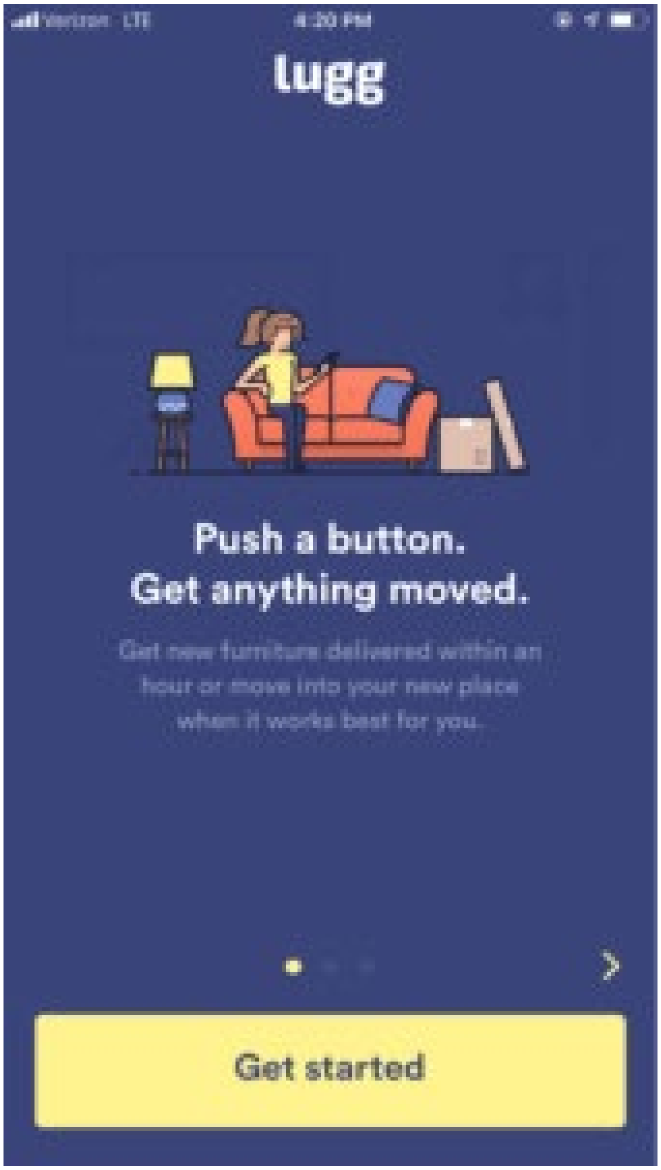

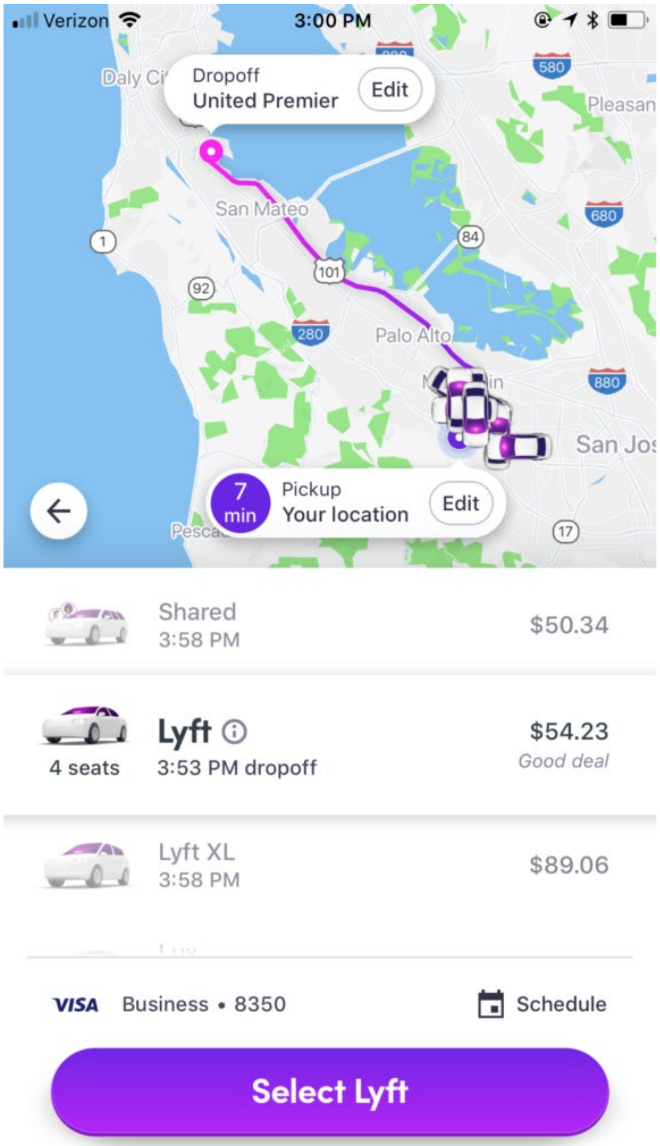

Companies like Handy, Lugg, Lyft, Rinse, Uber and many others made it efficient to connect to service providers in real-time. They created a full-stack experience around a particular service, optimizing for liquidity in one category. For these transactions, quality and success were more or less binary–either the service was fulfilled or it wasn’t–making them conducive to an on-demand model.

These platforms took on various functions to establish an end-to-end, seamless user experience: automatically matching supply and demand, setting prices, handling transactions, and establishing trust through guarantees and protections. They also often commoditized the underlying service provider (for instance, widespread variance on the driver side of rideshare marketplaces is distilled into Uber X, Uber Pool, Uber Black, Uber XL, etc.).

Unlike the previous generations of marketplaces, in which the provider ultimately owns the end customer relationship, these on-demand marketplaces became thought of as the service provider, e.g. “I ordered food from DoorDash” or “Let’s Uber there,” rather than the underlying person or business that actually rendered the service.

Over time, many startups in this category failed, and the ones that survived did so by focusing on and nailing a frequent use case, offering compelling value propositions to demand and supply (potentially removing the on-demand component, which wasn’t valuable for some services), and putting in place incentives and structures to promote liquidity, trust, safety, and reliability.

4. The Managed Marketplace Era (Mid-2010s)

In the last few years, we’ve seen a rise in the number of full-stack or managed marketplaces, or marketplaces that take on additional operational value-add in terms of intermediating the service delivery. While “Uber for X” models were well-suited to simple services, managed marketplaces evolved to better tackle services that were more complex, higher priced, and that required greater trust.

Managed marketplaces take on additional work of actually influencing or managing the service experience, and in doing so, create a step-function improvement in the customer experience. Rather than just enabling customers to discover and build trust with the end provider, these marketplaces take on the work of actually creating trust.

In the a16z portfolio, Honor is building a managed marketplace for in-home care, and interviews and screens every care professional before they are onboarded and provides new customers with a Care Advisor to design a personalized care plan. Opendoor is a managed marketplace that creates a radically different experience for buying and selling a home. When a customer wants to sell their home, Opendoor actually buys the home, performs maintenance, markets the home, and finds the next buyer. Contrast this with the traditional experience of selling a home, where there is the hassle of repairs, listing, showings, and potentially months of uncertainty.

Managed marketplaces like Honor and Opendoor take on steps of the value chain that platforms traditionally left to customers or providers, such as vetting supply. Customers place their trust in the platform, rather than the counterparty of the transaction. To compensate for heavier operational costs, it's common for managed marketplaces to actually dictate pricing for services and charge a higher take rate than less-managed marketplace models.

Managed marketplaces are a tactic to solve a broader problem around accessing high-quality supply, especially for services that require greater trust and/or entail high transaction value. If we zoom out further, there's many more categories of services that can benefit from managed models and other tactics to unlock supply.

What's next: The future of Service Marketplaces (2018-?)

We think the next era of service marketplaces have potential to unlock a huge swath of the 125 million service jobs in the US. These marketplaces will tackle the opportunities that have eluded previous eras of service marketplaces, and will bring the most difficult services categories online–in particular, services that are regulated. Regulated services–in which suppliers are licensed by a government agency or certified by a professional or industry organization–include engineering, accounting, teaching, law, and other professions that impact many people's lives directly to a large degree. In 2015, 26% of employed people had a certification or license.

Regulation of services was critical pre-internet, since it served to signify a certain level of skill or knowledge required to perform a job. But digital platforms mitigate the need for licensing by exposing relevant information about providers and by establishing trust through reviews, managed models, guarantees, platform requirements, and other mechanisms. For instance, most of us were taught since childhood never to get into cars with strangers; with Lyft and Uber, consumers are comfortable doing exactly that, millions of times per day, as a direct result of the trust those platforms have built.

Licensing of service professions create an important standard, but also severely constrains supply. The time and money associated with getting licensed or certified can lock out otherwise qualified suppliers (for instance, some states require a license to braid hair or to be a florist), and often translates into higher fees, long waitlists, and difficulty accessing the service. The criteria involved in getting licensed also do not always map to what consumers actually value, and can hinder the discovery and access of otherwise suitable supply.

Above: Bureau of Labor Statistics. (11/9/18)

Five strategies to unlock regulated industries

We're starting to see a number of startups tackling regulated services industries. As with each wave of previous service marketplaces, these new approaches bring more value-add to unlock the market, with variations in models that are well-suited to different categories.

The major approaches in unlocking supply in these regulated industries include:

1. **Making discovery of licensed providers easier**
2. **Hiring and managing existing providers to maintain quality**
3. **Expanding or augmenting the licensed supply pool**
4. **Utilizing unlicensed supply**
5. **Automation and AI**

1) Making discovery of licensed providers easier

Some startups are tackling verticals that lack good discovery of licensed providers. Examples include Houzz, which enables users to search for and contact licensed home improvement professionals, and StyleSeat, which helps users find and book beauty appointments with licensed cosmetologists.

2) Hiring and managing existing providers to maintain quality

Companies can raise the quality of service by hiring and managing providers themselves, and by managing the end-to-end customer experience. Examples are Honor and Trusted, managed marketplaces for elder care and childcare, respectively, which employ caregivers as W-2 employees and provide them with training and tools. In the real estate world, Redfin agents are employees whose compensation is tied to customer satisfaction, unlike most real estate agents who are independent contractors working on commission.

3) Expanding or augmenting the licensed supply pool

Expanding the licensed supply pool can take the form of leveraging geographic arbitrage to access supply that's not located near demand. Decorist, Havenly, Laurel & Wolf, and other online interior design companies enable interior designers around the world to provide design services to consumers without physically visiting their homes (yes, in many parts of the US interior design requires a license!). With improvements in real-time video, richer telepresence technologies, and better visualization technologies, more synchronous services are also shifting from being delivered in-person to online. Outschool and Lambda School are examples of de-localizing instruction, enabling teachers and students to participate remotely while preserving real-time interaction.

Another approach is to help suppliers navigate the certification process. A16z portfolio company Wonderschool makes it easier for individuals to get licensed and operate in-home daycares.

Lastly, there's the approach of augmenting certified providers so they can serve more customers. Fuzzy, an in-home veterinary service, uses AI and vet technicians to augment the productivity of licensed veterinarians; and a16z portfolio company Atrium builds automation and workflow management to provide efficiency gains in the legal industry.

4) Utilizing unlicensed supply

Some companies utilize unlicensed supply–notably Lyft, Uber, and other peer-to-peer rideshare networks. Another example is Basis, a managed marketplace for guided conversations with trained but unlicensed specialists to help people with anxiety, depression and other mild to moderate mental health issues.

In the pet space, Good Dog is a marketplace that brings together responsible pet breeders and consumers looking for a dog. Going beyond existing breeder licensing, which the company felt didn't map to what consumers valued, Good Dog established its own higher set of standards and screening process in conjunction with veterinary and academic experts.

5) Automation and AI

Other startups automate away the need for a licensed service provider altogether. These include MDAcne, which uses computer vision to diagnose and treat acne; and Ike Robotics and other autonomous trucking startups which remove the need for a licensed truck driver.

Opportunities for companies addressing regulated services

The last twenty years saw the explosion of a number of services coming online, from transportation to food delivery to home services, as well as an evolution of marketplace models from listings to full-stack, managed marketplaces. The next twenty years will be about the harder opportunities that software hasn't yet infiltrated–those filled with technological, operational, and regulatory hurdles–where there is room to have massive impact on the quality and convenience of consumers' everyday lives.

The services sector represents two-thirds of US consumer spending and employs 80% of the workforce. The companies that reinvent various service categories can improve both consumers' and professionals' lives–by creating more jobs and income, providing more flexible work arrangements, and improving consumer access and lowering cost.

The companies mentioned in this essay just scratch the surface of regulated industries. You can imagine a marketplace for every service that is regulated, with unique features and attributes designed to optimize for the customer and provider needs for that industry. (A full list of regulated professions in the US can be found here.) We fully expect more Airbnb- and rideshare-sized outcomes in the service economy.

If you're a founder who is looking to take on the challenge of tackling more complex services and bringing them online, we'd love to hear from you.

Why "Uber for X" startups failed: The supply side is king

THE REASON WHY "UBER FOR X" STARTUPS FAILED

TLDR: THE SUPPLY SIDE IS KING

Remember all the "Uber for x" startups?

A few years ago a ton of "Uber for x" startups got funded, but very few of them – maybe none? – worked out. It sounds good but ultimately most failed on the supply side. Let's explore why.

Rideshare has better economics, at the same acquisition cost

Rideshare is special. Acquiring a broad base of labor for driving is expensive, often $300+. But then they can get requests all day. You can work 20 hours and even 50 hours a week if you want. You continually need the driver app to find new customers

Where a lot of "Uber for x" companies fall down – valet parking, car washing, massages, etc – is that demand is often infrequent and there's spikes at a few points in the day. What's your supply side supposed to do the rest of the time?

In other words, "Uber for x" cos often have the same cost of acquisition and cost of labor as rideshare, but can't fill their time with work as smoothly / profitably

Marketplace outcomes are sensitive to unit economics

Rideshare networks are fickle and require a long period of being unit economic negative before they can break even, with enough scale/density. But a lot of "Uber for x" cos can never dig out of that hole, and stay unprofitable forever

This is one of the reasons why I'm bearish on food delivery as a stand-alone business in the long run. Uber can tap into their supply side and augment with food delivery earnings. Pure food companies have to get the same drivers but can't pay as well

The key is to go for a different pool of workers

So what kind of "Uber for x" ideas can work? Ultimately the ones that go for a completely different pool of labor. Folks who prefer to work from home. People who don't live near a city with rideshare. People who don't own cars. Etc.

If you can find a different pool of labor, they still have the same motivations around flexible schedules and easy earning potential. You can use the same techniques as Uber – simple UX, transparent pricing, etc – and apply them to these marketplace opportunities

In that way, the lessons from "Uber for x" are a subset of best practices you can learn from marketplaces. You need a strong strategy to get the supply/demand flywheel going. A big market with a defensible moat. Fragmentation that can be solved w transparency and aggregation

Don't emulate – approach from first principles, starting from the workers' POV

IMHO "Uber for x" cos failed to become a thing because they sought to emulate ridesharing when they should have just approached their particular market from first principles. There's still a ton of marketplace opportunities out there and am excited to see what people do!

Because all these marketplaces tend towards supply constrained, you should evaluate each opportunity/company from the POV of the supply side. Does it work for them? Can they do it 40 hours/week and stay sticky? When can you pull away subsidies? These are the key questions

The key lesson!

Supply side is 👑 .

If you're interested in more reading about Uber and marketplaces,
[I collected my favorite 20 links here]

The Power User Curve: The best way to understand your most engaged users

[Today we have an essay on one of the common frameworks we use to analyze investments at Andreessen Horowitz: The Power User Curve. I worked closely with Li Jin, a partner on the investing team, to collect our ideas into this essay which she wrote. You can follow @ljin18 on Twitter for more thoughts. -Andrew]

The importance of power users

Power users drive some of the most successful companies — people who love their product, are highly engaged, and contribute a ton of value to the network. In ecommerce marketplaces it's power sellers, in ridesharing platforms it's power riders, and in social networks it's influencers.

All companies want more power users, but you need to measure them before you can find (and retain) them. While DAU/MAU — dividing daily active users (DAUs) by monthly active users (MAUs or monthly actives) — is a common metric for measuring engagement, it has its shortcomings.

Since companies need a richer and more nuanced way to understand user engagement, we're going to introduce what we'll call the "Power User Curve" — also commonly called the activity histogram or the "L30" (coined by the Facebook growth team). It's a histogram of users' engagement by the total number of days they were active in a month, from 1 day out of the month to all 30 (or 28, or 31) days. While typically reflecting top-level activity like app opens or logins, it can be customized for whatever action you decide is important to measure for your product.

The Power User Curve has a number of advantages over DAU/MAU:

- ***It shows if you have a hardcore, engaged segment that's coming back every day.***
- ***It shows the variability among your users: some are slightly engaged, whereas others are power users. Contrast this with DAU/MAU: it's a single number and so blurs this variance.***
- ***When mapped to cohorts, Power User Curves let you see if your engagement is getting better over time — which in turn helps assess product launches and performance of other feature changes.***
- ***Power User Curves can be shown for different user actions, not just app opens. This matters if the core activity that matters for your product is deeper in the funnel.***

In other words, while the DAU/MAU gives you a single number, the Power User Curve gives entrepreneurs several avenues of analysis to assess their product's engagement to the most addicted users — in a single snapshot, over time, and also in relation to monetization. This is useful. So how does it work?

The Power User Curve will "smile" when things are good

The shape of the Power User Curve can be left-leaning or smile-like, all of which means different things. Here's a smile:

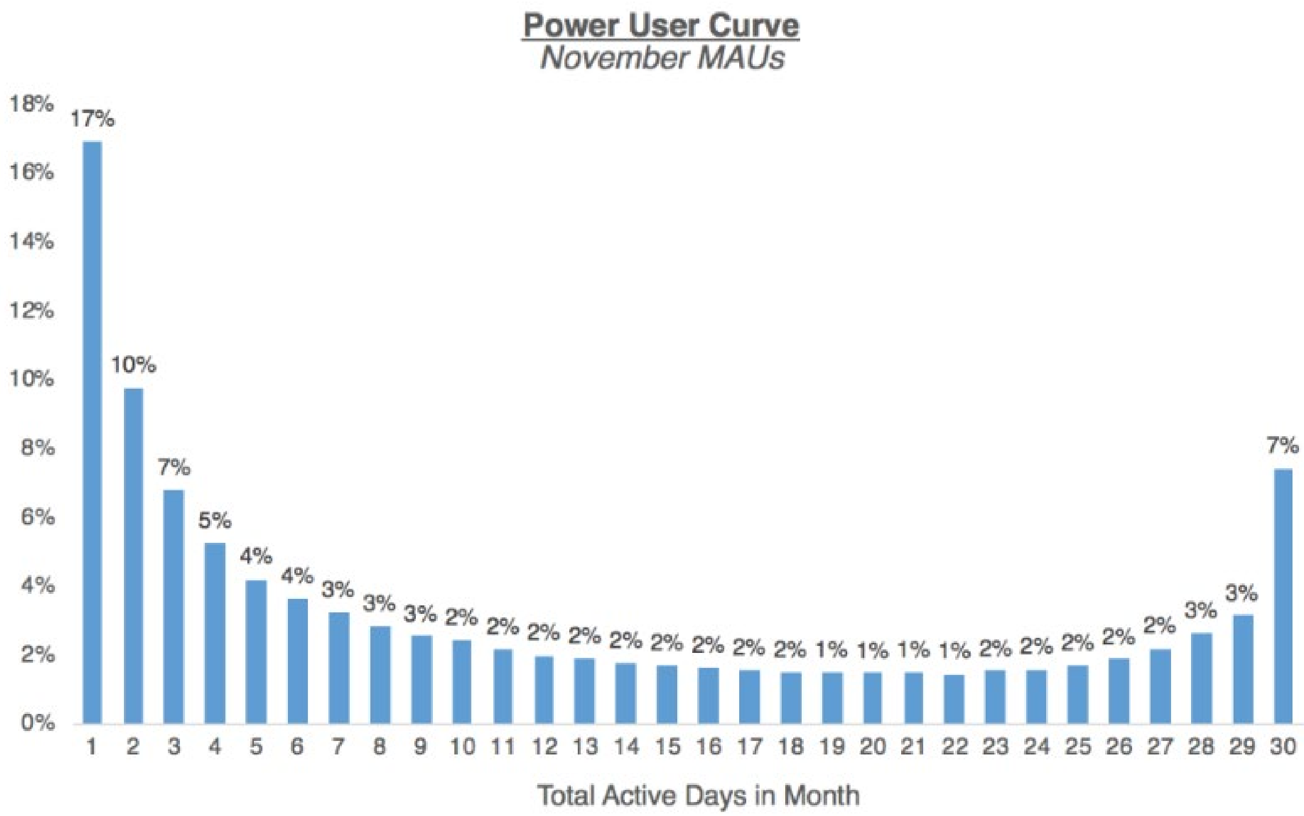

The Power User Curve above is for a social product, and shows the characteristic smile shape that indicates there's a group of highly engaged users using the app daily or nearly daily. Social products with frequent user engagement like this lend themselves well to monetization via ads—there's enough users returning frequently that the impressions can support an ad business. Remember that Facebook would have a very right-leaning smile, with 60%+ of its MAUs coming back daily.

What matters is that, over time, the platform is able to retain and grow its power users: successive Power User Curves should ideally show users shifting over more to the right side of the smile. As the density of the network grow, and with stronger network effects, it's expected that there's more reason for users to return on a daily basis.

The Power User Curve can show when strong monetization is needed

Let's look a different example, which doesn't smile:

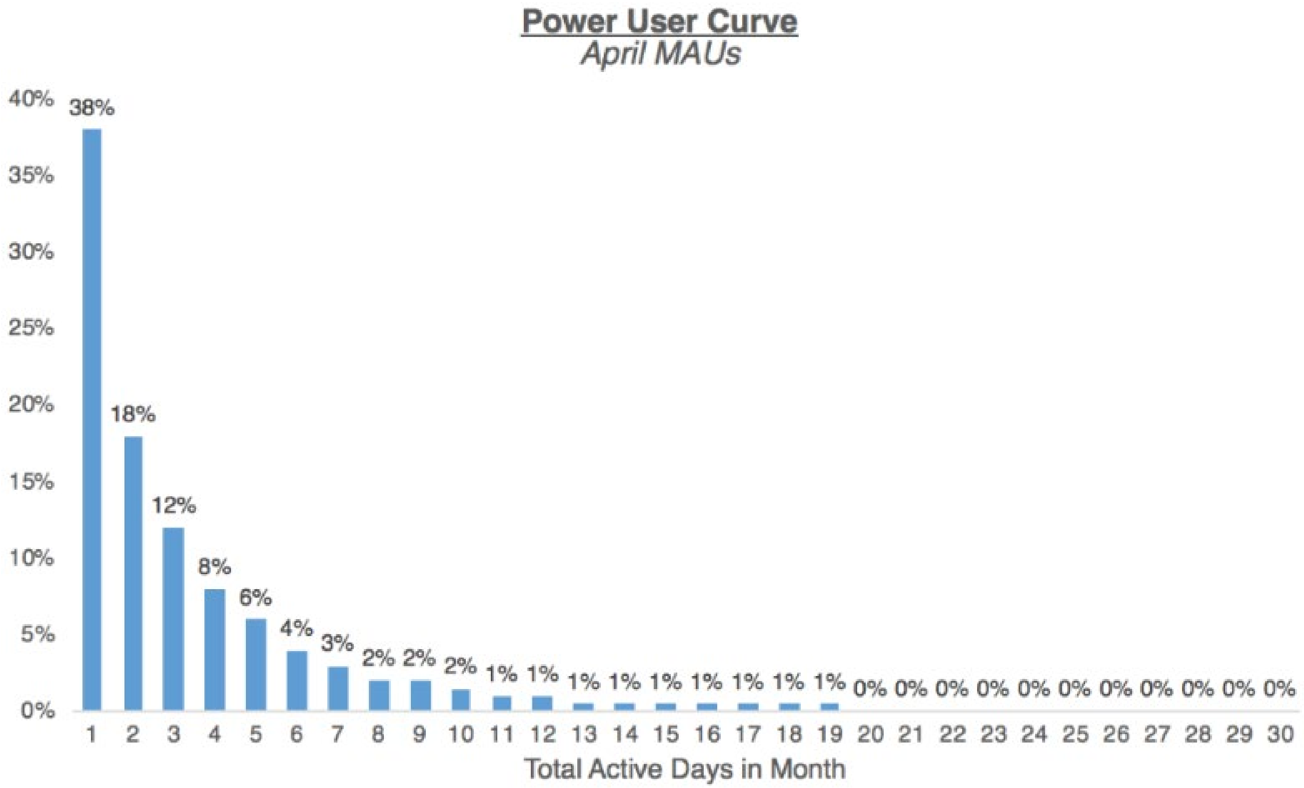

This Power User Curve of a professional networking product looks quite different than that of a social product. It's left-weighted with a mode of just 1 day of activity per month, and decays rapidly after those few days. There's no power users. But this light engagement can be okay — not every company needs to have a smile-shaped Power User Curve, just as not every product category necessarily lends itself to an ultra-high DAU/MAU.

When there's low engagement, what matters is that the company has a way to extract enough value from users when they *are* engaged. Think about an investing product like Wealthfront or networks like LinkedIn — few users are likely to actively check it on a daily basis, but that's ok, since they have business models that aren't tied to daily usage.

CEOs of such companies should therefore,think about: Is there a way to create revenue streams where the business can still monetize effectively despite users' infrequent engagement? Or, who are the users using this product more frequently, and how can I get more of them? Is there something about the product — e.g. onboarding, the core experience, etc. — where a significant chunk of the user base isn't experiencing the 'aha moment' that makes them "get" the product, and therefore not getting value from it right now (and if so how to get there)?

Some products should be analyzed in a 7 day timeframe – like SaaS/productivity – and others on 30 days

Another flavor of the Power User Curve is a histogram of users' engagement for a 7-day period, also commonly called L7. The 7 day Power User Curve shows weekly actives, not monthly actives. Plotting this version can make sense if your product naturally follows a weekly cycle, for instance, if it's a productivity/work-related product that users engage with Monday through Friday. B2B SaaS products will often find it useful to show this version, as they want to drive usage during the work week.

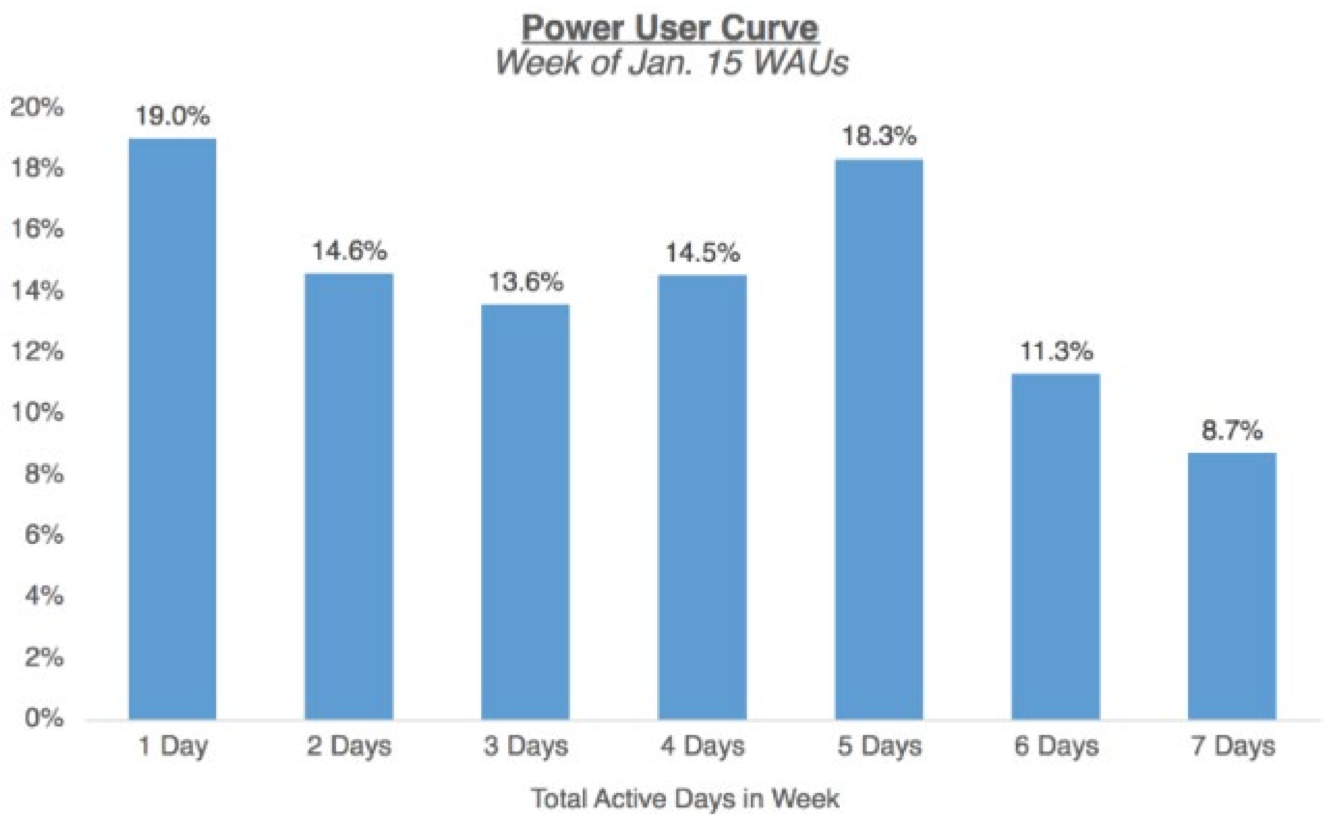

Note that using DAU/MAU wouldn't be the appropriate metric for this product as it's not designed to be a daily use product. You can also see there's actually a smile curve through 5 days, but fewer users are using it 6-7 days, which makes sense for the power users of a workweek product like this.

CEOs of such product companies should therefore want to understand: Who are the users engaging just 1 or 2 days each week? Are there certain teams or functions within an organization that are getting more value, and how can I build out features to capture the teams with less engagement? Or, if the product is really driving a lot of value for specific departments — how can I understand their needs better and make sure we continue building in a direction that supports their daily workflow (and that we can upsell new features)?

The trend of over time can show if the product is getting more engaging over time

Plotting the Power User Curve for different WAU or MAU cohorts can also be very insightful. Over time, you can see if more of your user base are becoming power users, by seeing the shift towards higher-frequency engagement.

Here's an example:

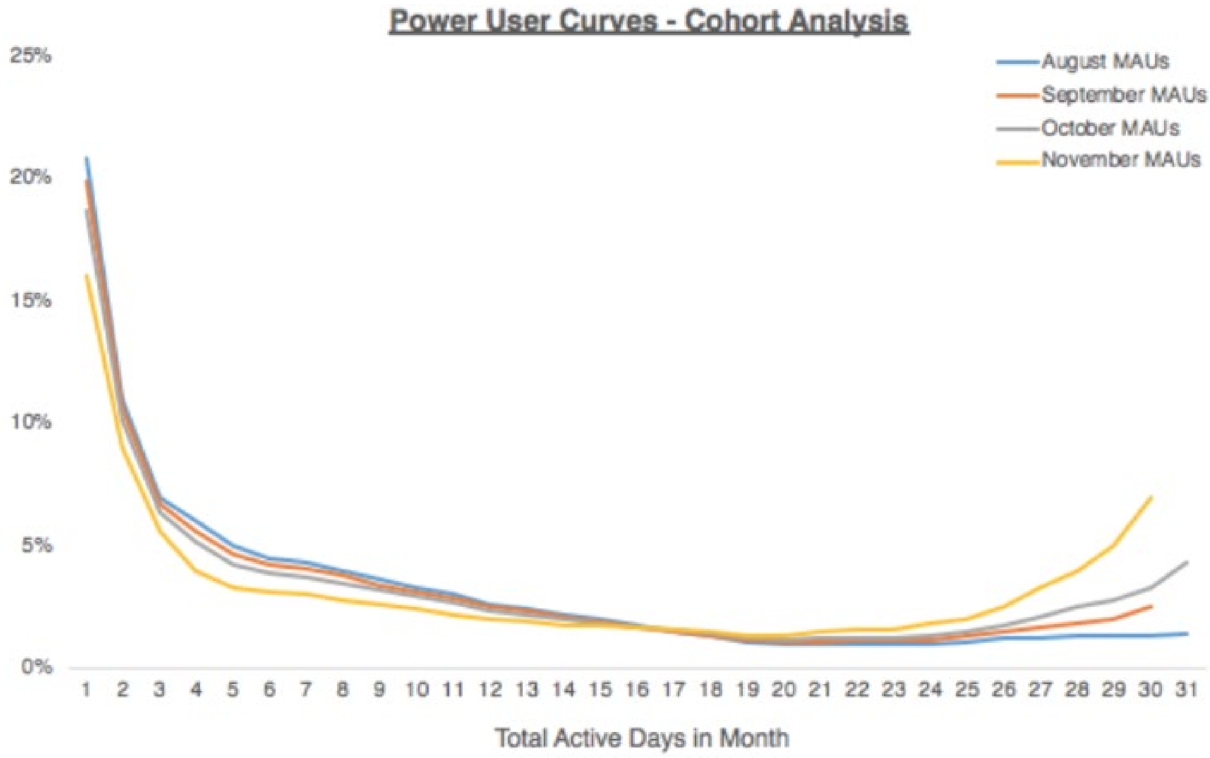

The Power User Curve for MAU cohorts from August through November shows a positive shift in user engagement, where a larger segment of the population is becoming active on a daily basis, and there's more of a smile curve.

You can see when the line starts to inflect in order to see when a critical product release or marketing effort might have started to bend the curve. This might be a place to double down, to increase engagement. For a network effects product, you might expect to see newer cohorts gradually improve as you achieve network density/liquidity.

On an ongoing basis, you can measure the success of product changes or new releases by looking at different cohorts' Power User Curves. If a product unblocks a bunch of features for power users, you might see a gradual increase in power users.

The Power User Curve can be based on core activity, not just app opens or logins

The frequency histogram can be keyed on actions beyond the visit — did someone show up or not — you can also go with deeper user actions. For instance, you may want to plot the core activity that maps closely to how your business is monetized… or that better represents whether users are getting value from your product. This is important because it forces you to think about what really matters to measure.

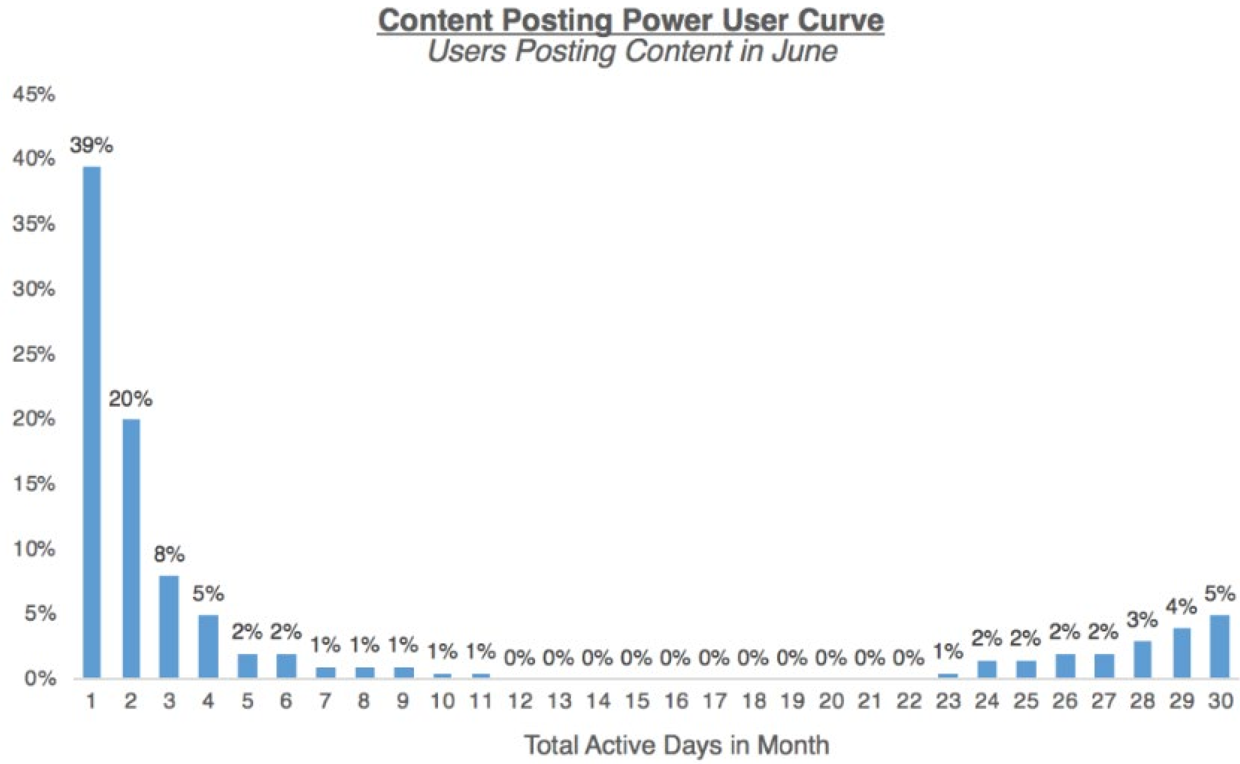

The above chart for a content publishing platform shows the total number of days in the month users posted content. A lot of products have smile-shaped core activity Power User Curves, because while most people tend to contribute lightly, there is a small contingent of users who are power users. Think of the distribution of Youtube creators, or Ebay sellers, or even how often you post on Facebook.

As the CEO or product owner of a platform like this, it's important to design the platform such that the everyone has a chance to succeed. On Facebook, the news feed algorithm makes sure that if you feel strong affinity to a person or organization, you'll still see their posts even if the sheer volume of other content (for instance, from more prolific media companies) would otherwise drown it out. On OfferUp, even if I seldom sell items, when I do list something, their algorithm makes sure that it's surfaced to the relevant potential buyers.

Why does this all matter?

Not everything is a daily use product, and that's okay.

Power user analysis allows you to get a better understanding of how users are engaging with your product, and make more informed decisions using that data. That might mean choosing an appropriate business model that works for your pattern of engagement, or designing better re-engagement loops for lower-engaged user segments, or doubling down on use cases that your high-engagement user base is already getting value out of.

The beauty of the Power User Curve over DAU/MAU is that it shows heterogeneity among your user base, reflecting the nuances of different user segments (and therefore what drives each of those segments). Creating versions of Power User Curve by various user segments can also be particularly insightful. For instance, for a business with local network effects (like Uber or Thumbtack), showing Power User Curves by market can reveal which geographies are developing density and strong network effects.

Power User Curves show if your product is hitting a nerve among a super engaged core group of users, even if perhaps the overall blended DAU/MAU is low. It also doesn't have to just reflect app opens or logins — you can hone in on an action that maps closely to users getting specific value out of your specific product and plot the Power User Curve for that action. The key for founders is to know that there isn't a single silver bullet to measure perfect engagement — rather, the goal is to find the set of metrics that are appropriate for their businesses. Comparing the Power User Curve of a social app vs. a work collaboration app doesn't make sense, but looking at your own Power User Curve over time, or finding benchmarks for your product category, can tell you what's working… and what's not.

[Thanks again to Li Jin for pulling together this essay!
Another plug for her Twitter account here. -A]

DAU/MAU is an important metric to measure engagement, but here’s where it fails

How DAU/MAU got popular

DAU/MAU is a popular metric for user engagement – it’s the ratio of your daily active users over your monthly active users, expressed as a percentage. Usually apps over 20% are said to be good, and 50%+ is world class.

How did this metric come into use? DAU/MAU has been a popular metric because of Facebook, which popularized the metric. As a result, as they began to talk about it, other consumer apps came to often be judged by the same KPIs. I first encountered DAU/MAU as a ratio during the Facebook Platform days, when it was used to evaluate apps on their platform.

Are you up to date?

Get new updates, usually once a week – featuring long-form essays with what’s going on in tech.

This metric was always impressive for Facebook because it’s always been high. It’s historically been >50%. In fact, I was curious at one point whether or not it’s always been that good. And it has! I found this from a Facebook 2004 media kit showing crazy high numbers even with a small base of 70k users:

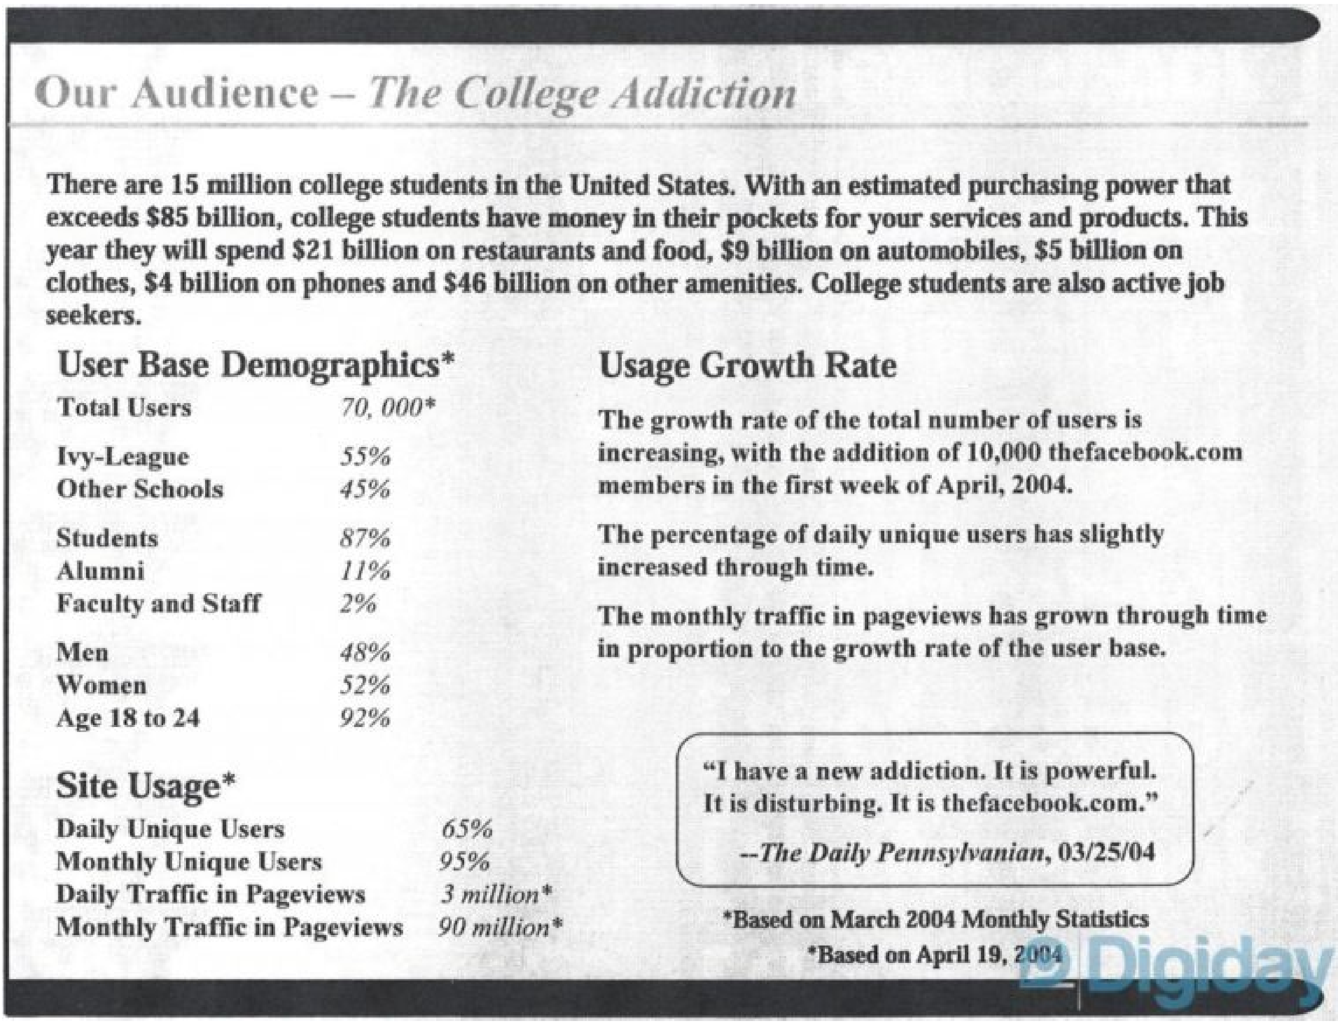

Our Audience – *The College Addiction*

There are 15 million college students in the United States. With an estimated purchasing power that exceeds $85 billion, college students have money in their pockets for your services and products. This year they will spend $21 billion on restaurants and food, $9 billion on automobiles, $5 billion on clothes, $4 billion on phones and $46 billion on other amenities. College students are also active job seekers.

User Base Demographics*

Total Users	*70, 000**
Ivy-League	*55%*
Other Schools	*45%*
Students	*87%*
Alumni	*11%*
Faculty and Staff	*2%*
Men	*48%*
Women	*52%*
Age 18 to 24	*92%*

Site Usage*

Daily Unique Users	*65%*
Monthly Unique Users	*95%*
Daily Traffic in Pageviews	*3 million**
Monthly Traffic in Pageviews	*90 million**

Usage Growth Rate

The growth rate of the total number of users is increasing, with the addition of 10,000 thefacebook.com members in the first week of April, 2004.

The percentage of daily unique users has slightly increased through time.

The monthly traffic in pageviews has grown through time in proportion to the growth rate of the user base.

"I have a new addiction. It is powerful. It is disturbing. It is thefacebook.com."

--*The Daily Pennsylvanian*, 03/25/04

*Based on March 2004 Monthly Statistics
*Based on April 19, 2004

Assessing product/market fit with DAU/MAU

It's an important metric, to be sure, but it's often misused to say that "XYZ isn't working" when in fact, there's a slightly less frequent usage pattern that's still equally valuable.

For consumer and bottoms up SaaS products, this metric is super useful, but seems to mostly exclude everything besides messaging/social products that are daily use. These are valuable products, but not the only ones.

Products that aren't daily, but still hugely valuable

Not everything has to be daily use to be valuable. On the other side of the spectrum are products where the usage is episodic but each interaction is high value. DAU/MAU isn't the right metric there.

- ***At Uber, our most profitable rides are to airports, via Black Car for a special night out, business travel, etc. These don't happen every day, and although there are folks using us to commute, that's not the average use case. So our DAU/MAU wasn't >50%. The driver side has clusters of "power drivers" who are active >30hrs/week, but as it's been widely published, our average driver is actually part-time. (Pareto Principle!)***
- ***Linkedin is another interesting example which is low frequency – only recruiters and people looking for jobs use it in daily spurts – but it throws off so much unique data that you can build a bunch of vertical SaaS companies on top of this virally growing database.***
- ***Products in travel, like Airbnb and Booking, are only used a few times per year by consumers. The average consumer only travels ~2x/year. Yet there are multi deca-billion dollar companies built in this space.***
- ***In fact, for SaaS, it seems to be the exception not the rule. While email and business chat can be nearly daily use, a lot of super important tools like Workday, Google Analytics, Dropbox, Salesforce, etc. might only be used 1-2x/week at most.***
- ***Much of e-commerce looks like this too, of course. You buy mattresses, new sunglasses, watches, etc fairly infrequently. Yet there are $1B+ wins in the category.***

You may notice a pattern here. If you're low-frequency/episodic, then you have to generate enough dollars or data that it's valuable. If you're high-frequency, you have a higher chance of growing virally and building an audience business that monetizes using ads.

Nature versus nurture

To extend this idea further, you can argue that messaging/social products with high DAU/MAU is actually the extreme case, and in fact most product categories don't index highly. A few years back I [shared this interesting diagram from Flurry] which compared different app categories and their retention versus frequency of use:

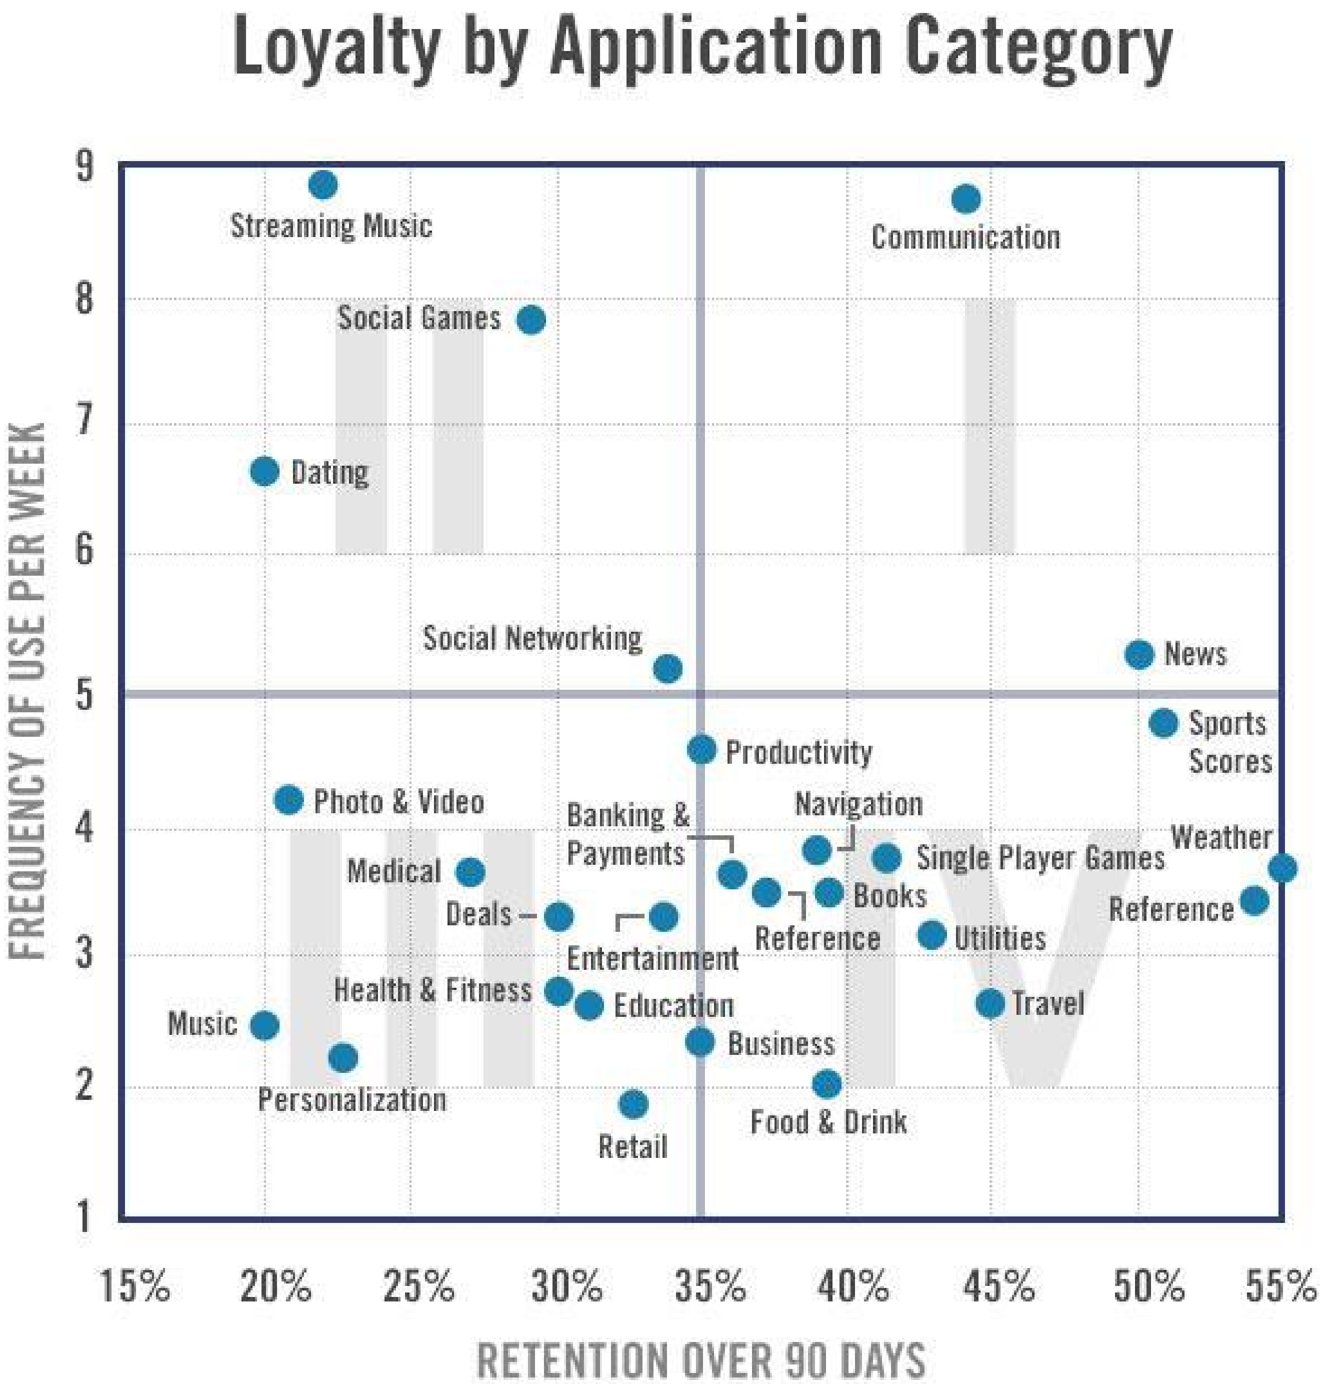

In this chart, a couple categories jump out:

- ***Social games have high frequency ("I'm getting addicted!") but once you burn through the content, you tend to churn***
- ***Weather is interesting too – you don't often check, maybe only on cloudy days, but you will have a need to check throughout your entire life- so it maxes out on highest retention rate over 90 days***
- ***Communication, for all the reasons discussed before, is both high frequency and high retention. That's awesome!***

What I'd love to see on this chart would be another overlay, monetization. There, I bet Travel, Dating, and Gaming would tend to stand out for different reasons. Travel because each transaction is big, and Dating/Gaming because it's frequency combined with a focus on monetization because you won't have the user for long.

So you want to increase DAU/MAU? It's hard

So let's say that you want your DAU/MAU to increase – so what do you do? Funny enough, a lot of people seem to implement emails and push notifications thinking it'll help. My experience is that it tends to increase casual numbers (the MAU) but not the daily users. In other words, it'll actually lower your DAU/MAU to focus on notifications because you'll grow your MAUs more highly than your DAUs.

I've also not seen a 10% DAU/MAU product, through sheer effort, become 40% DAU/MAU. There seems to be a natural cadence to the usage of these product categories that doesn't change much over time.

Increase, measure your hardcore users, network effects, monetization

If your DAU/MAU isn't super high, this is what I like to see instead: Show me your hardcore userbase. What % of your users are active every day last week? What are they doing? How are you going to produce more of them? Showing this group exists goes a long way.

Similarly, show how the freq of use increases in correlation to something. Perhaps size of their network – showing network effects – or how much content they've produced or saved. Then make the argument that by increasing that variable, DAU/MAU will rise in cohorts over time.

Finally, maybe DAU/MAU is just not for you. Sometimes you don't have to be a foreground app to be successful. Maybe you just need to build something awesome that does something valuable for people, makes enough money, and they use it twice a year! Also great.

DAU/MAU is useful, but has its limits

In conclusion, if your product is a high-frequency, high-retention product that's ultimately going to be ads supported, DAU/MAU should be your guiding light. But if you can monetize well, develop network effects, or quite frankly, your natural cadence isn't going to be high – then just measure something else! It's impossible to battle nature… just find the right metric for you that's telling you that your product is providing value to your users

Conservation of Intent: The hidden reason why A/B tests aren’t as effective as they look

When a +10% isn’t really a +10%

OK, this is an infuriating startup experience: You ship an experiment that’s +10% in your conversion funnel. Then your revenue/installs/whatever goes up by +10% right?

Wrong :(

Turns out usually it goes up a little bit, or maybe not at all.

Why is that? Let’s call this the ***“Conservation of Intent”*** (Inspired by the Law of the Conservation of Momentum 😊)

The difference between high- and low-intent users

For all your users coming in, only some of them are high-intent. It's hard to increase that intent just by making a couple steps easier – that'll just grow your low-intent users. Doing tactical things like moving buttons above the fold, optimizing headlines, removing form fields – those are great, but the increases won't directly drop to your bottom line.

In other words, the total amount of intent in your system is fixed. Thus the law of the conservation of intent!

This is why you can't add up your A/B test results

If you're at a company that A/B tests everything and then announces the great results – that's wonderful, of course, but just run the thought experiment of summing together all of those A/B tests. And then look at your top-line results. Rarely does it match.

The most obvious way to see this is to test something high up on a funnel, for example maybe the landing page where a new user hits, or an email that a re-engaged users opens – you can see that a big lift on the top of the funnel flows down unevenly. Each step of friction burns off the low-intent users that are flowing step-by-step.

Be skeptical of internal results, but more importantly, external case studies too

If you're at a big company and another team publishes a test result, make sure you agree on the actual final metric you're trying to impact – whether that's revenue, highly engaged users, or something else. Make sure you always review that.

Similarly, this is a reason to be skeptical of vendors and 3rd parties who have case studies that'll increase your revenue by X just because they increase their ad conversion rate (or whatever) by X. In these kinds of misleading case studies – often presented at conferences – not only do vendors have the ability to only cherry pick the best examples that reinforce their case, but also the metric that's highest impacted! Be skeptical and don't be fooled.

Unlock increases to the bottom line

First, understand what's really blocking your high-intent users. Those are the ones who'd like to flow all the way through the funnel, but can't, for whatever reason. For Uber, that was things like payment methods, app quality (for Android especially!), the forgot password flow, etc. If you can't pay or can't get back into your account, then even if you use the app every day, you might switch to a different app that's less of a pain in the ass.

Also, you can focus your experiments. You obviously get real net incremental increases on conversion the further down the funnel you go. By that point, the low-intent folks have burned off. You're closer to the bottom line. Look the steps right around your transaction flow – for ecommerce sites that might be the process to review your cart and add your shipping info, or the request invoice flow for SaaS products, etc. Think about high-intent scenarios, for example when you hit a paywall or run out of credits/disk space/resources/etc. All of these can be optimized and it'll hit the bottom line quickly.

Make sure your roadmap reflects reality

When it comes to your product roadmapping, yes you can definitely brainstorm and ship a bunch of +10% increases, but you need to add a discount factor to your spreadsheets to reflect reality. Can't just add up all your results.

When you focus on low-intent folks, you'll have to get creative to build their intent quickly. Things like being able to try out the product, having their friends into the product – these are the "activation" steps that generate intent. Here's a great place to start – a highly relevant essay on getting users more psych'd, guest written by Darius Contractor from the Dropbox growth team.

Conservation of Intent

Many of you have directly experienced the "Conservation of Intent" but now you have a name for it! It's tricky.

This is really a reflection of how working on product growth is really a combo of psychology and data-driven product. You can't just look at this stuff in a spreadsheet and assume that a lift in one place automatically cascades into the rest of the model.

Conservation of Intent: The hidden reason why A/B tests aren't as effective as they look

Brand marketing is mostly useless for consumer startups.

Startups build a great brand by being successful, finding product market fit and scaling traction, etc. But it's not a real lever. Let's not mix up correlation with causation!

If this seems contrarian to you, it's because there's a **vast ecosystem of consultants, agencies, and other middlemen** who are highly incentivized to have you spend $ and effort on non-ROI/ non-performant activities. Early startups should opt out of all of this

It's easy to confuse correlation and causation: If you're starting a consumer startup, you see successful late stage cos with fawning media coverage, amazing conference speaking slots, celebrities on the cap table, etc., and think that's what caused their success: Great brand.

But great brand is the **lagging indicator of success.** The buzz is created by the hard work that the entrepreneurs put in: Finding product/market fit, hiring a great core team, finding acquisition channels that scale. Brand marketing is great, but it should be layered on later.

The greatest consumer products in recent years slogged through years of obscurity. The overnight success of Uber, Airbnb, Instagram, etc were actually multi-year successes driven by hard work and multiple pivots.

Working on press mentions, conferences, etc can be a good way to get an initial hit of traffic. It's great! **But it's not enough**. Here's an article from a few years back: After the TechCrunch bump, there's life in the trough of sorrow.

Anyone who's been on the homepage of TechCrunch, AngelList, Hacker News, or even in the NYTimes knows that it's a increase to your dopamine but not so much your customer acquisition :) It's great for the early days, but you need a lot more to scale.

Furthermore, the metrics-driven argument is obvious. Ultimately, the engagement in every product can be deconstructed into a series of user cohorts that join and decay over time. How does brand help these cohorts? My observation: **They don't help much.**

One argument is that brand marketing can create buzz and word of mouth. OK if that's the case, why does every brand-driven commerce company have **>60% of their customer acquisition** happen through paid marketing? Why do they have to buy all their customers?

If brand marketing helps make acquisition ultimately cheaper, then why does every startup's paid acquisition become less efficient over time, even as the company becomes more well known? The same arguments apply to startups' re-engagement efforts.

It's true that a strong brand can confer defensibility in a noisy space – but it's brittle, hard to create, and hard to sustain. Hard to bet on that in the early days of a startup.

Where brand marketing does matter, especially outside of consumer: **Recruiting a great team. Raising money. Partnerships.** These are all small targeted audiences where you can reach them with more touchy feely efforts, and it can work! So put your emphasis there.

For early consumer startup efforts, it's better to focus on the basics. Understand your users, deliver a great product to the market that grows by itself, built moats, monetize in a user-aligned way. Grow your team, work with the best advisors/investors/etc. The basics.

Do all that, and your product's brand will take care of itself – and then you can layer on more brand marketing efforts to 10x the effect. Just don't do the steps out of order!

The Scooter Platform Play: Why scooter startups are important and strategic to the future of transportation

(📷 *lime)*

The scooter startups are way more important than you think, or in emoji-speak:

🛴+📱=🤖🚗🚁. Let me explain.

Right now, scooters are a lot of things – fun, cute, adventurous – but here's a couple words I rarely hear about them: Strategic. Important. Platform play. And yet they are.

Chris Dixon has written that the next big thing will start out by looking like a toy. Scooters are literally derived from kids' toys. It's the perfect example. (Btw, here's the perfect moment to re-read his essay [arguing that the next big thing will start out looking like a toy](#))

Like a toy, a scooter seems underpowered vs other transportation options. It only takes you on short trips – a few blocks at a time. It's cheap and makes less money than a highly profitable Uber trip to the airport. They are placed all over the place in cities, annoying many

These all seem like weaknesses, but in fact they're strengths. Because scooters are cheap, short-range, and ubiquitous, it means consumers are adopting them as an alternative to walking

SCOOTERS COMPETE WITH WALKING! What's the market size on that?? :)

As a result, the scooter apps are being downloaded in the millions by consumers – the adoption has been incredible. But now we have another starting point to capture the intent to go from Point A to Point B. That intent is valuable

These scooter trips are short, frequent, and cheap, driving high engagement in the app. In fact, if you live in SF they become a home screen app. You might check it all the time, before you walk a couple blocks

Combine those factors – millions of consumers, high frequency, and strong intent – and all of a sudden it's obvious why this is a big deal

When you're the first look and the highest frequency place to start your trip, it's the pole position in consumers' minds. Everything else is downstream

Google has one of the best business models ever. It's the starting point. It has a search box, maps user intent to URLs, and charges everyone downstream if they want to be promoted in any way

Scooter cos like Lime are also the starting point. High frequency and high intent. It has a search box for where you want to go, and maps user intent to a trip

Scooter apps could be the starting point for a lot of kinds of trips. Alongside Apple/Google Maps, rideshare, etc – the place where you'd go to book your autonomous vehicle rides. Or your VTOL / flying car trips. It could even upsell rideshare trips from Uber and others

Scooters look like a toy, but in fact they are something else:

Strategic. Important. Platform play.

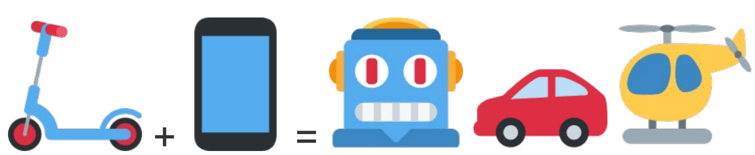

The IRL channel: Offline to online, Online to offline

(📷 dmagazine)

We've heard about Facebook ads, Google adwords – but today let's talk about the **"IRL Channel."**

The IRL channel is an underappreciated advantage of companies that exist in the real world – Amazon Echo, Envoy, Lime, Uber, etc – that use constant in-real-life reminders to try out and use the product

The viral acquisition benefits are pretty obvious. If you've never seen/tried a product, but you see them swarming around your city (or your workplace), then naturally you'll want to try it out

More importantly, some product usage patterns are naturally viral. Couple examples:

- ***Transportation fits into this bucket, which is why Uber's rider acquisition mostly viral/WOM Traveling and going out are social activities. You bring your friends and loved ones in the car with you, to share the costs. Even the fully utilitarian version – going from point A to point B – can be social, since there's often a person on the other side.***
- ***The new scooter/bike trend is another obvious example. Lots of brightly painted Lime scooters all over SF makes a splash. Put some pricing and instructions on the actual hardware, and riders who have big smiles on their faces, and you have a natural acquisition channel.***
- ***With Amazon Echo, the physical presence gets you an retention/ engagement benefit. Sitting on your kitchen counter naturally encourages you to use it. The newest one, the Show, has a display which invites you to interact. Sometimes the Amazon Echo thinks you're talking to it when you're not. I've always thought that Amazon is unlikely to ever fix this since it probably increases engagement when it occasionally gets things wrong :)***
- ***Envoy is a B2B example. I've signed in with the system at the lobbies of dozens of companies, which means if I ever have to make a purchasing decision in the category, they'll be the natural choice.***

There's not a ton of entrepreneurs who are brave enough to build new consumer hardware cos, but if you are, I think this has to be a key consideration!

The IRL channel is about a physical experience that drives you into a digital one. But the other way around is pretty profound as well.

When you see your social feeds populated with photos from highly instagrammable retail experiences like Boba Guys or the Museum of Ice Cream, like below…

(📷 dBoba Guys fb page)

(📷 laweekly)

… you can't help but pull up Yelp to figure out the closest place to go!

The other mega trend here is esports and the fandom community, of course. One day you're playing League of Legends and reading Star Trek fan fiction, and the next day you're going to esports arenas and checking out vidcon. It's a thing.

The IRL channel is real. It helps you with acquisition, retention, and more. It's starting to go both ways – from online to offline, which has been a force in retail for the past few years – but also offline to online, where IRL products remind you to interact with their digital sides.

Super fascinating, and I'm excited to see where this will all go!

How startups die from their addiction to paid marketing

Many of the biggest implosions in recent history – especially ecommerce – have been due to startups getting addicted to paid marketing while fooling themselves on Customer Acqusition Costs. As spend scales, it always gets more expensive and harder to track – never less.

A familiar story: New product launches. Nice spike, but it dies down. The product is low freq – gotta spend to grow. Marketing spend increases, it's profitable! More is spent, more money is raised via VCs. OMG this is working! Party!

Suddenly top line hits a ceiling. Payback period goes from 9 months to 12, then more. Unit economic profitable, but not with staff + HQ. Without top line growth, more investment dollars can't be raised. Budgets get slashed, then layoffs.

Even slower growth means a pivot is in order. Try something else, also powered by paid marketing. Maybe subscription? Premium? Try another thing. Then another. Irrelevance – or maybe bankrupcy.

This happens enough that y'all should be nodding your heads now – it's tough, but there's a pattern. This is the Paid Marketing Local Max.

The key insight here is that Paid Marketing is tricky to grow, at scale, as the primary channel. It's highly dependent on both against external forces – competition and platform – as well as the leadership team's psychology when things get unsustainable.

The first mistake is to start by thinking of everything as Blended CAC – dividing all your acquisition against dollars – as opposed to understanding CAC of each channel (Facebook, Google display, Google AdWords, etc.). The former is misleading.

Because your initial organic users are your biggest fans, your Blended CAC and per-channel CAC can often by off by 2-5X. As you scale your paid, your organic won't follow 1:1. So as you grow, your Blended will approach your dominant channel's CAC.

Scale effects mostly work against you in paid marketing. The longer your campaigns run, the less effective they become – people start seeing your ads too often. The messaging becomes stale, and novelty effects are real. Market performance has a reversion to the mean.

Saturation is also a thing. As you buy up your core demographic, the extra volume comes from non-core, who are less responsive. The first US-based ad impression on a property is the most responsive, but you eventually run out of those.

Competitive dynamics are real. They'll come in to copy not just your product, but also ad messaging and creative. It's not hard to fast follow, especially if you can start the test just with a experiments on millennial-friendly ad copy and landing pages.

Contrast that to viral channels, folder sharing in Dropbox or team channel creation for Slack – these are highly situational and only a few folks can copy. Whereas in ads you're competing with everyone going after your same demographic.

Addiction to paid marketing can get you into a local maximum. It's much harder to fix the underlying issues – creating real moats, product differentiation, doing deeper adtech integrations. Easier to just spend more and push the LTV window from 9 months to 12 to 18.

There's a few scenarios where paid marketing is justified, but it's situational. If your product has network effects that kick in after an activation point and really scale, you can use paid to help bootstrap that. Facebook uses paid to build out new regions, for example.

If you are really going to invest a ton of time from engineering/growth to integrate with all the APIs, try out a ton of things algorithmically, then you can develop a lasting edge. I've heard Wish does this well, but it's not common.

The new generation of ad platforms makes it possible to scale revenue to new heights, but without profitability. Make sure you don't get addicted. Build out new channels. Fix churn and frequency. Don't congratulate yourself too early. And calculate LTV/CAC correctly :)

So what do you do about it? One of the best case studies of this is from @drewhouston's Dropbox presentation from the early days. Lots of great stuff in this deck and it's worth paging through, now nearly 10 years later. Here it is.

On slide 18, Drew talks about early experiments they did on paid search. They executed the industry best practices at the time – go to trial-based pricing, hide the free option, optimize landing pages. Slide:

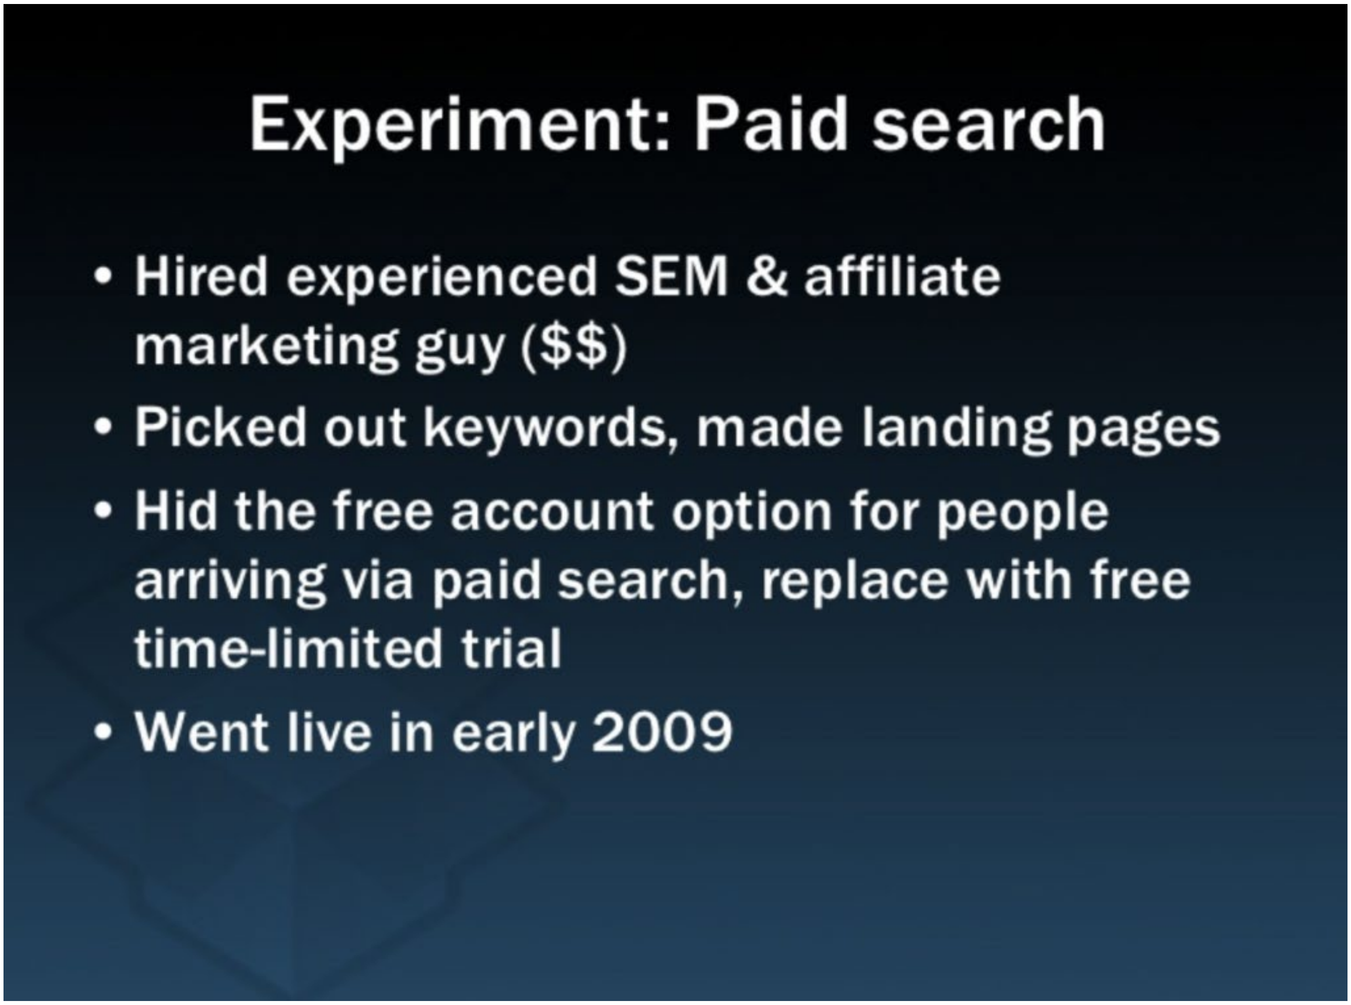

What they learned was that, in the mature market for cloud storage, there was already a lot of competition. All the paid marketing channels were unprofitable. Hiding the free option wasn't user aligned. Etc etc.

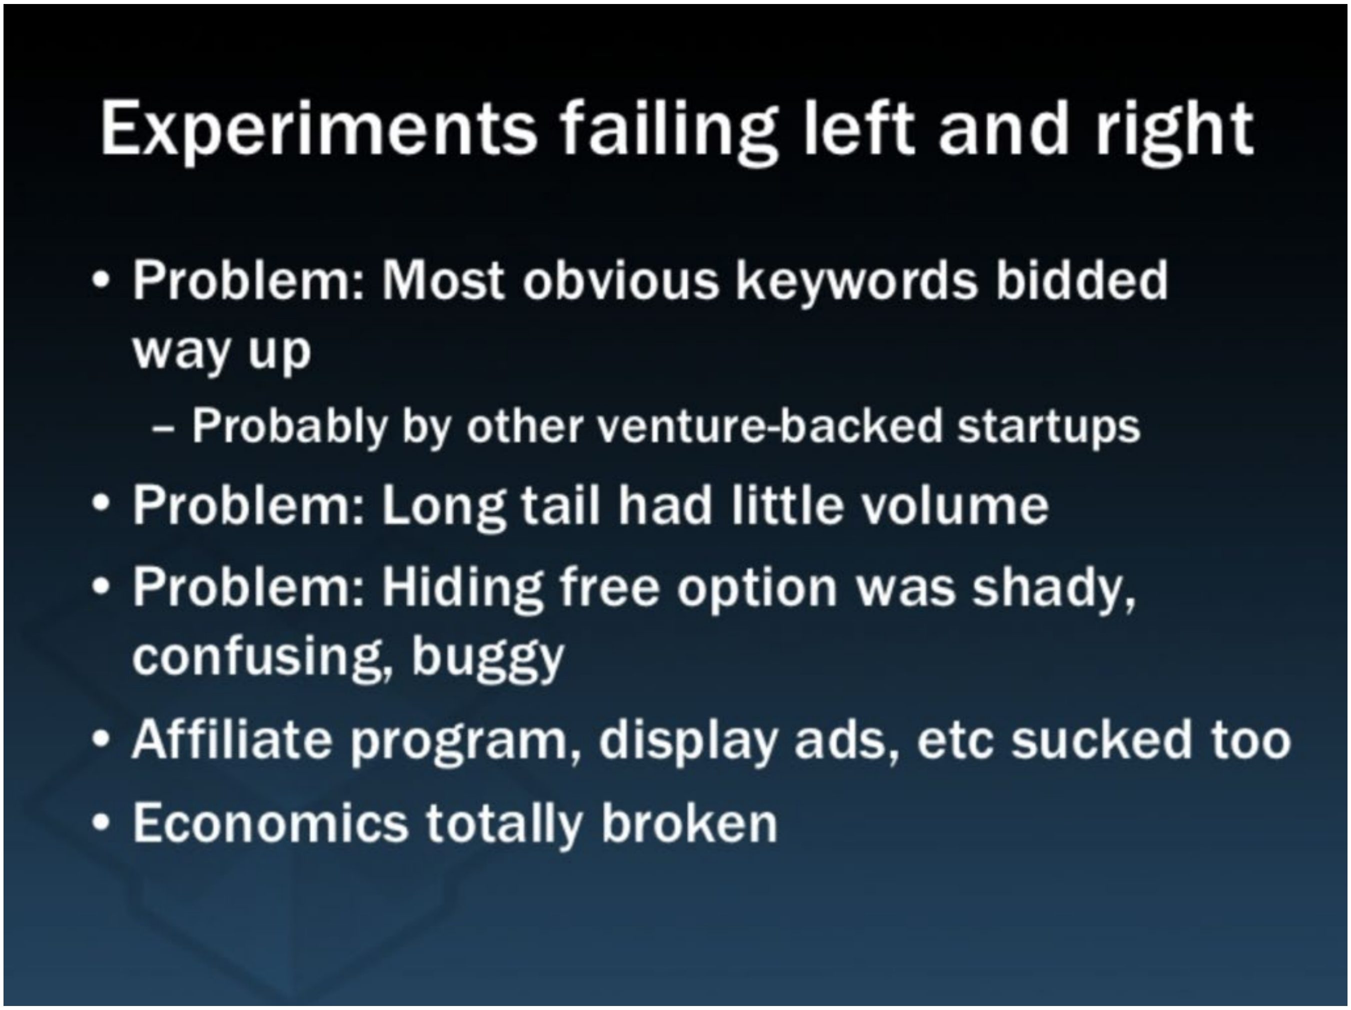

The obvious move would have been to continue to grind on the problem! Tweak pricing, optimize more ads/funnel/landing pages, etc. And many would have been tempted to do that, because it's worked for others

The interesting thing, and you can see in the deck, is that grew virally instead – via folder sharing, the give/get disk space program, etc. It seems obvious now, remember that back in the day, "cloud storage" was the space, and it's not clear that you can go viral there.

Dropbox has done well since then, of course!

Where we are now

500M+ Registered users
180+ Countries
$1.1B+ GAAP revenue
11M+ Total paying users
400B+ Pieces of content added to Dropbox

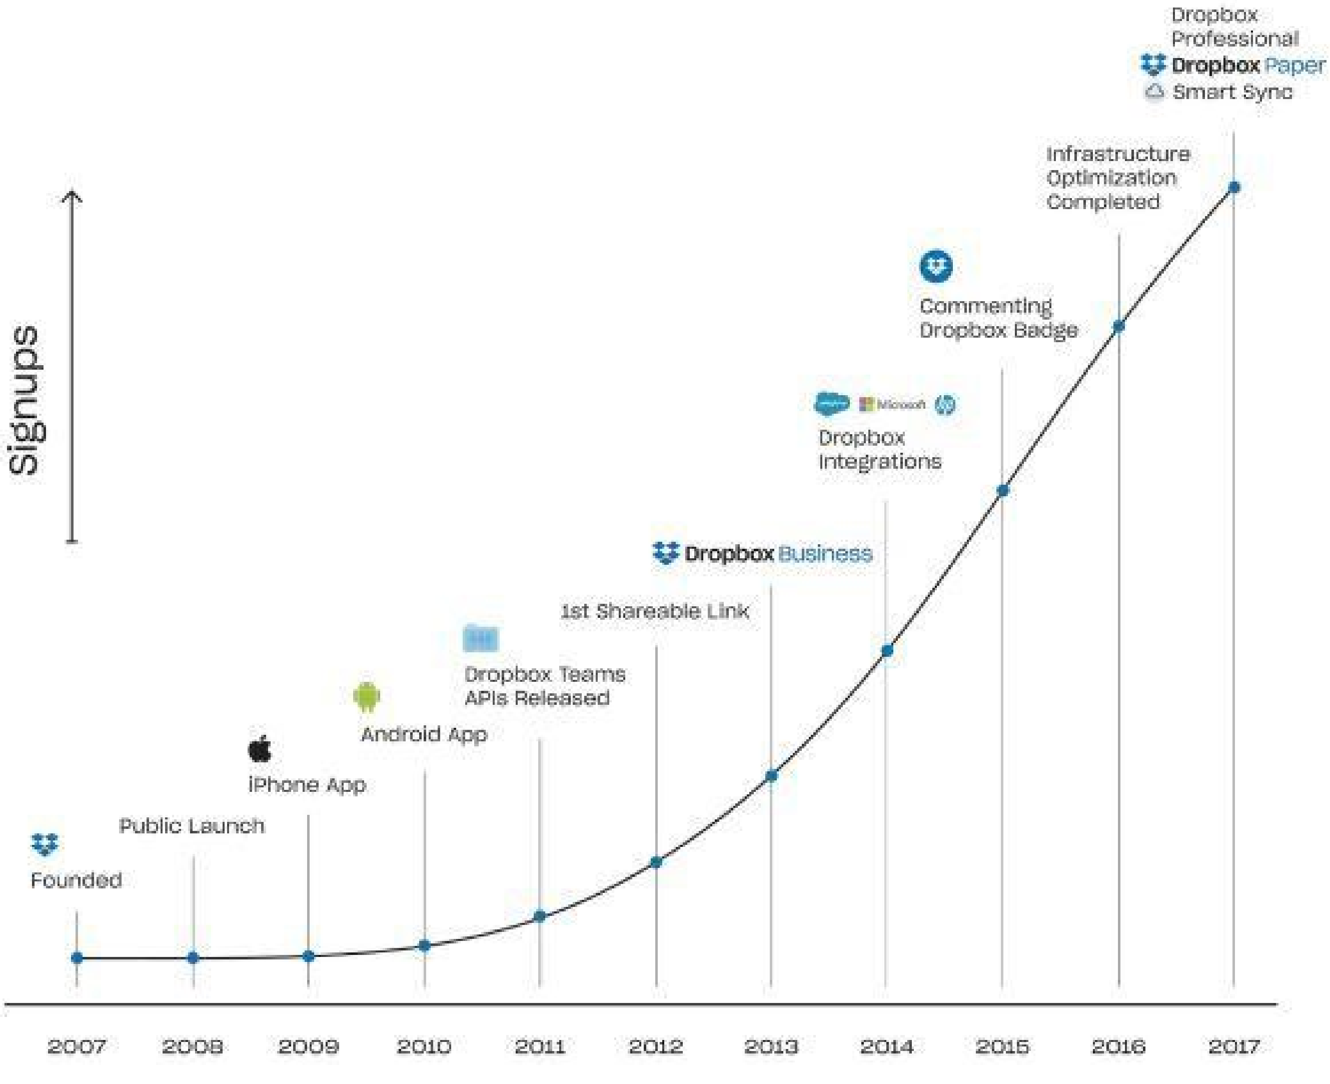

As an aside, isn't it interesting that exponential growth curves always look linear instead? Here's Slack's as well:

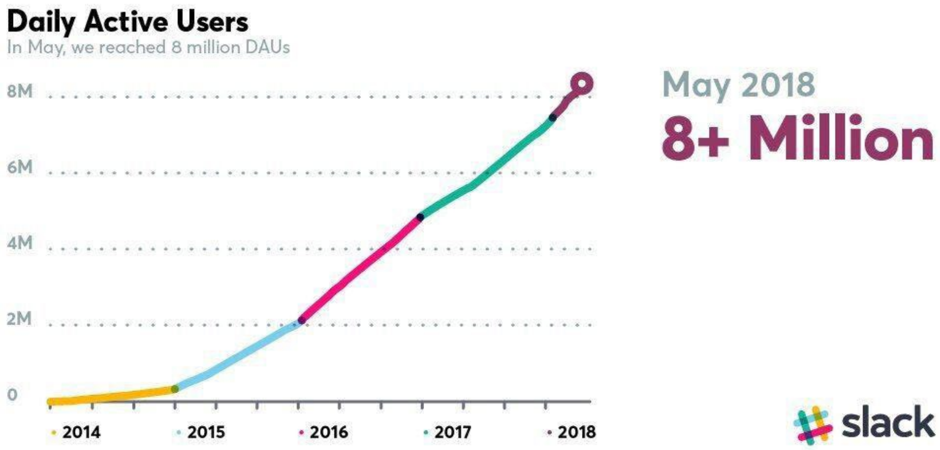

In some ways, you could argue that Dropbox is lucky that their initial forays into paid marketing didn't work. That made it easier for them to stop their efforts there, and to focus on the viral channels that are now their bread and butter.

On the other hand, it takes a lot of insight and reflection to go away from the current industry "best practices" – even if they erode profitability, cause shark fins, etc.

So for those of you who are thinking about going all-in on paid marketing, I challenge you to go deeper on that strategy. Perhaps cap your paid acquisition at 30-40% of TOF. Instead, where can you innovate?

In addition to Dropbox, I sometimes use the story of @Barkbox, which created a whole media property, Barkpost (http://barkpost.com) as a viral content sharing engine that can cross-sell the subscription product.

Or at Uber, although they never became significant channels, we were keen to work on sharing viral sharing features like Share ETA, Fare Split, and Location Sharing to potentially drive acquisition.

The point is, knowing that Paid Marketing is highly addictive and hard to scale down, all of us in the industry should always be thinking about the 2nd or 3rd channel, in addition to organic/WOM, to give us a way to wean off an ever-increasing ad budget.

To do that, you'll need empower your creative team to attack the problem from all angles- new viral product features, really investing in your referral program, building out your content/SEO strategy even though it'll take years. It's worth the investment!

Update: I'm joining Andreessen Horowitz!

Hi readers,

Big update: I'm joining Andreessen Horowitz as a general partner!

Starting in April, I'm returning to my roots to invest in and help grow the next generation of startups. I'll be focused on consumer startups, bottoms up SaaS, marketplaces, and more – utilizing my expertise in growth to launch and scale new companies. Incredibly excited.

How this came together tells you a lot about Marc and Ben, and how Silicon Valley works. I moved to the Bay Area in 2007, as a first time founder with a lot of energy and a lot of questions. I spent the first year meeting everyone I could, reading everything about tech, and writing down all that I was learning. A few months in, I was shocked to get a cold email from Marc introducing himself. Who knew that sort of thing happened? My blog was pretty much anonymous and I could be anyone – but he reached out to talk ideas, which made a big impression. I learned a lot about Silicon Valley that day.

hi

Marc Andreessen 8/11/07 Share this email

to me

Hi Andrew -- just wanted to drop you a note and say I really enjoy reading your blog -- let me know any time you have a free half hour or so and want to grab a cup of coffee!

Best,
Marc

Marc soon introduced me to Ben, and together, they provided a regular stream of advice/ideas/frameworks over breakfasts at the Creamery, Hobee's, Stacks, and other assorted Palo Alto diners. I was a first-time founder, and the real-life entrepreneurial experiences they relayed – on fundraising, finding product/market fit, hiring, and much more – proved to be insanely helpful. My startup ultimately didn't work out and the team soft-landed at Uber, but I always remembered the incredible support from Ben and Marc.

Andreessen Horowitz is a firm built on the same core values I saw first-hand. The investing team has a deep empathy for entrepreneurs that reflects their extensive operating experience. The partners on the operating team are incredible and enable the firm's unmatched support of founders and their companies. For all of these reasons and more, I'm thrilled to join the a16z team.

A decade ago, I learned how impactful it can be when a couple experienced entrepreneurs reach out to a new founder. I can't wait to close the loop by doing the same – working with new founders and their startups, and helping build the next generation of tech companies.

As excited as I am about the next step, I'm also sad to leave an extraordinary team and experience behind at Uber. I have nothing but admiration for the talented, passionate people who are working hard to ship all the amazing innovations that are coming down the queue. To all my friends and colleagues at Uber, thank you for the amazing two and a half years.

Finally, to the readers of this blog: I'll be writing much more! The new job will let me put down a lot of what I've learned over the past few years. I'm excited to share ideas and stories from across companies and industries with all of you! Looking forward to it.

Onwards!

Andrew

San Francisco, CA

About

Andrew Chen is a general partner at Andreessen Horowitz, a Silicon Valley venture capital firm, where he invests in new consumer startups

Previously, he led Rider Growth at Uber, advised/invested in dozens of startups, and has been writing for over a decade! Now residing in the San Francisco Bay Area

Follow @andrewchen