
Running Your Company

Patrick Collison | CEO of Stripe | October 18, 2018



Running Your Company with Patrick Collison

This is a lightly edited transcript of a conversation that took place between Stripe CEO Patrick Collison and YC Partner Adora Cheung. You can find a [video of the talk here](#).

Adora Cheung: So, Patrick, welcome. Patrick is the co-founder and CEO of Stripe, he launched the startup or now pretty big company in 2010, correct? With his brother John.

Patrick Collison: Well, actually we started working on it full time in 2010 but actually to your comment just there about company's launching, it took us quite a while to launch because we had all these banking partnerships in place and so on. So we didn't launch until September 2011. We had been working on it for almost two years at that point, and every time we saw PG or really anyone else from YC, all they would ask us is why we had not launched yet. So some things don't change.

That's interesting, so two years until you had to-

It was I think one year and 11 months from first line of code to public launch. Which to be clear, I don't recommend, that is not a good idea. Just we had to because we had to get all these other things in place, and because it took us so long to be able to publicly launch. We tried to be very disciplined about gradually expanding the number of users every month, and so even though we weren't publicly available, we got our

first user, our first production user just three months in. And then every month we tried to add at least a handful of users. So by the time we publicly launched, we did have about 100 users, which back then to us that seemed like a big deal and something very, very large.

Speaking of, actually when I was preparing for this interview, I was trying to jog my memories and I remember specifically because your office was near here in Palo Alto. I remember back then people would always talk about the Collison brothers running around, going to people's offices and installing their API into their web apps, and intrude, do things, don't scale fashion. I assume you were not only trying to make sure they installed it, but also get user feedback. I haven't so much that actually ... I don't know if you know, Paul Graham, PG, now calls it the Collison installation, and this is actually something we told the founders to go do this, do the Collison installation because obviously in hindsight it seems so obvious to do.

Well, it served two purposes and one is your point. It served as a really good way to get to do user research and to get UX feedback. If you've done this, I'm sure you had the experience where you design what you're absolutely certain is the most streamlined, user-friendly, straightforward, frictionless way to do whatever it is the product does. And then you present it and put it in front of a user and you're just going to ask them to do whatever it is, and they find it completely impenetrable and they're clicking all the wrong links and they can't find the next button even though the next button is there blinking in green and stuff.

Running Your Company with Patrick Collison

It's invariably incredibly painful that the nothing is so sobering as watching somebody use the first version of some new product. But the other reason for us was we had suggested to somebody they use Stripe or they switch to Stripe or whatever, and invariably the response is, "Oh, yeah, sure, that sounds awesome."

But it can be that they'll postpone it and they'll postpone it again, and it's like there's never a moment where it's like, "Okay, this is the evening when I'm going to switch to Stripe." Us going and accosting to the person helped ... He'll talk about it in sales there's always like you know why you and why this and why now, and these kinds of questions. Going in person created a why now moment, it's like well we're here to house.

Did you just show up or how did you ... Like what was that?

I don't think we ever actually just showed up although perhaps we should have but no, we tried to be at least semi-invited.

Stripe now today, you've come a long way since back then, it's really been a decade, not even. But today well you have 1300 employees across nine offices across the world, you're doing ... I have a list here, you manage 200 million API requests a day, you process billions a year for millions of companies. In your latest round of funding, Stripe is now worth \$20 billion. Anyway, the list goes on, I'll stop there otherwise people are going to think I'm your PR agent. But, anyway, you've clearly done something right and so I want to spend a lot of I guess the time today talking about running

your startup from the perspective of the startup CEO, you yourself. It's like zoom, what do you think about from zooming in on the day to day operations to zooming out to the long-term strategic decisions? Maybe to help us ease into the discussion, one thing is when you start off, from the very beginning, a lot of friends get together and they come up with ideas and they're super excited and they start working on it. And then at some point they need to decide, "We need a CEO for this company." Some people aren't meant to be CEOs but for you and John, I've met both of you, you're very smart, ambitious people with great qualities and attributes that correlate to becoming a great leader. How did you and John decide you would be the CEO?

Well, I think Stripe is unusual where John and I are obviously brothers and we've known each other for a long time. Because of the relationship, we've also placed a lot of trust in each other and we really do run the company together. There's no major decision that Stripe has made that we've not both been a part of. It's not always the case that despite being the CEO that I'm the person who ... like in the case of disagreement, it's not always the case that I prevail. Our dispute resolution framework is much more around which of us cares more than which of us hold this title or that? John's title is President and so there's some ... I don't know, both are significant roles. In that regard, I think we're an anomaly sort of the fact that I became CEO was honestly semi-random. I would say I think because we're brothers we're able to get to this unusual situation where we really do run it together.

Running Your Company with Patrick Collison

Is there a certain-

Sorry, it may not be a helpful answer because perhaps you're trying to encourage all these companies, it was like, "Shit, one of you guys has got to be the CEO but ..."

Well, do you think there's a rubric for this of, "Here are five questions you should answer and maybe then you decide."

That's a good question.

Well, it wasn't one thing which is that I think it is important to just have an efficient mechanism for reaching a decision and it can't be simply orange around a consensus. If there are three co-founders and none of you quite want to ... Nobody is a clear CEO and you don't have some efficient mechanism for having decisions get made, I think that is a recipe for failure. Even, I don't know, doing some quasi democratic voting is probably not a great idea either.

So for myself and John, we both have areas we respect that we specialize in, and so that doesn't mean we have absolute autonomy and authority there but we bias in that direction. He spends more time, for example, working externally with users, I spend more of my time working internally on product and engineering things.

Not to say that he doesn't make decisions there or I don't here, but again there is a bias in that direction. And then second we have those additional aspect where in the case of it being a major decision and we respectively disagree, then we really do try to make based on which of us is just more passionate about it. Because

that will correlate with the outcome. If one of us really wants to do something or thinks that X or Y is the right thing to do, simply wanting it to be so passionately is more ... That can be a self-fulfilling prophesy and so I think it's not irrational to have it be a consideration.

Yeah, I also see the best teams that work well together are the ones in which everyone, they all want the best idea to win, not their idea to win. So there's a stepping back in an unselfish way to get to that conclusion.

Definitely, and I think that something that we're lucky about that I think is very ... Well, exactly to your point, I think is definitely present in the most successful co-founding relationships that I've seen is some degree of dispassion and disagreement. For us, it goes it's easier to get to it because we've been disagreeing with each other for 20 years, and so it will [inaudible 00:09:13] of the emotional charge.

But I think that finding other mechanisms whereby you can get there such as it's not this, I don't know ... That the notion of disagreeing strongly is not a scary phenomenon and both parties or multiple parties, if they are more than two, are suppressing their feelings for fear of there being this new divergence.

How many more siblings do you have? You have one more brother?

One more sibling, yes, three of us in total.

Would he join the ... Or is it just you and John are the thing that's the special match there?

Running Your Company with Patrick Collison

Tommy, my youngest sibling, he's quite a bit younger than myself and John, and so John is almost or approximately two years younger than I am. When we started Stripe, I guess I was about 21 and I think I guess therefore John was 19, and Tommy was still definitely midway through high school and so just wasn't quite practical at the time.

Yes, finish high school.

Now I think if you asked him he'd probably say he'd never throw his load in with miscreants like us.

Yeah. Cool, so in terms of the role of CEO, often people say there's a threshold and time in which and it's related to product/market fit where you have a role as a pre-product/market fit CEO and then which is completely different in your role as a post-product/market fit to you. So I want to spend most of our time talking about pre-product/market fit but just to calibrate those questions, in terms for Stripe, what was product/market fit for you? Like how did you define it? Were there metrics to it? What number of employees you were at when you reached it, so on and so forth?

Yeah, that's a really good question and I think you're exactly right to divide things into those ... the story of a startup is two stories, it's a story of getting to product/market fit and then the story of what happens subsequently. Obviously, there isn't like a totally definitive binary line between them but I think it's helpful to frame the narrative in that regard.

Obviously, for Stripe, actually around the time we launched publicly, I think is basically when we had it. In that, when we launched publicly in September 2011, we rebuilt significant components of Stripe multiple times in response to user feedback. We're on the third version of our dashboard and the second or third depending on how you count major version of our API.

So we've gone through a lot of iteration response to the evidenced challenges that users had or the deficiencies the product seem to possess. When we launched, we were basically immediately bottlenecked on being able to serve user demand rather than generating user demand. I think directionally that's the inversion that happens versus the early days instead of really trying to figure out, "Well, okay, condition on or given some user, how do I make sure that it does what they need." They end up a happy, retained user in a sufficient fraction of cases, whatever. And then at some point it floats ... I think it's very different, do you take 100 users and some fraction of them onboard and some smaller fraction remain engaged or whatever and so 100 users the curve asymptotes downwards.

And then you take 100 users and again a meaningful fraction end up engage and they actually tell more people about it or they invite people. I don't mean strictly just in a viral sense, but just generally speaking. Just that lead to and generate more demand such as things simply in some super linear fashion growing.

Running Your Company with Patrick Collison

I think like being less than one is just very different to being more than one, and it seems again relaunched. That didn't generate that many users, I don't exactly remember but let's just pretend that 500 businesses signed up on the day we launched. Which immediately those 500 businesses told other businesses and people heard about it and all the res. So the next day we had a lot of businesses as well and so forth. Such that from the day we launched, the challenge became keeping up with demand rather than tweaking the product to somehow serve the market.

When you launched, how many people were you at point with employees?

We were about 10.

10, okay. So I guess before you launched then, well your day to day sounds like it was just, like we were talking about earlier, just running around talking with users and fixing issues. Was that literally every single day was it like that? Or were you doing every day?

Okay, so not literally every single day but I would say we really tried very hard to understand in very granular detail what exactly it was that people were doing, where they were tripping up and so on. So just to give you some examples, we had a chatroom for providing support when people were integrating Stripe and it was actually a public chatroom which had some downsides. Because if we had broken something or if somebody had some embarrassing issues, everyone would see us.

We thought that was good because it would actually put more pressure on us to have the product be good. We had literally every time for the first, I don't know, call it 10 years of Stripe, every time somebody sent an API request to Stripe, it sent an email to us.

We were looking at every single API request, and if we saw users do something weird, we all go, "Why did they do that?" Why were our docs confusing or whatever? I'd get a dinner in the evening or something ... Well, again it seemed like a lot of the time I would have maybe 10 emails because Stripe was not handling a lot of the requests back then. But you're literally looking at each individual action, and actually Stripe we just celebrated or at least passed, we're not very good at celebrating but we passed our seventh anniversary just a few days ago on September 29th.

So I was looking at our ... We had an hourly stacks email that we used to send and so the day we launched we handled 22 payments in the previous hour which again seemed huge to us back then.

But I was noticing in that email I actually had forgotten this that we had things like that. In the email, we literally had a list of businesses that had submitted three or five or something API requests but had instead had gone live, they had not launched the businesses. We literally just had the emails of all these businesses.

Running Your Company with Patrick Collison

The idea was that we were then literally individually reach out to them as like, “Hey, you seem to use Stripe a little bit but you didn’t go live. Did we screw up or was the product somehow inadequate or whatever.” We did things like every time anyone ever hit an error of any sort, that generated like a high priority email to us. So we would try to immediately to go solve it and that also generated I think those pleasant user experience for it.

It’s frustrating to hit an error in some service but we could then often 15 minutes later reach out to them and say, “Hey, we saw that you had this problem, it’s all fixed now.” Sometimes where the user made a silly mistake, they mistyped an API parameter or something and we just reached out to them and are like, “Hey, noticed you had a typo in your code.” Which perhaps for some of them it was a little bit unsettling but at least helped us increase the throughput product feedback. But these are all examples I’d say of the general pattern of really trying to be hyper-attentive to all the micro details of what people were doing with the product and iterating rapidly in response to it.

Generally speaking, I think that pre-product/market fit metrics are actually relatively unhelpful. I think you really want to buy is very strongly towards as much inspection and high throughput quality to feedback as possible. Because probably not that many people are using your product and so if it’s 20 users, you can in some sense afford to just look at everything they’re doing and try to understand what’s working, what isn’t.

In some sense it’s how much do they ... You can tell subjectively like how much they love your product or how much are they going to probably be really upset if it just disappears.

Totally, and actually on this point we had a thing. At the bottom of every webpage, we had a little text box anchored to the bottom of the brass or frame and one line high text input but we had placeholder text in it to try to prime people to tell us things. It had my favorite part of Stripe is ellipsis and people just fill it out and can hit enter. But also of course has most of the prompts are negative. It was like the worst thing about Stripe is or the worse thing at this page is or I really hate the way Stripe does or whatever.

Again, at that stage you have to be masochistic and so again we’d be, “Oh, I was waking up to all these emails telling us all the terrible things about Stripe. But that it was helpful to do lest for the day ahead.

How did you stay happy if ... Like in the early days [crosstalk 00:18:41]-

What makes you think we did? I would say it was well happy is such a squishy concept, right?

True.

Because there are lots of things that we ... I guess when I look back, and look maybe this is the rationalization I tell myself. But when I look back at through life of the things that I’m most glad I did, I wasn’t exactly happy with while doing them.

Running Your Company with Patrick Collison

Like often I was very stressed out or I had to work really hard or whatever but they are the things that [inaudible 00:19:16] brought the most fulfillment. Well, part of it is there is a rich literature here and so I won't dive too deeply into it, but I do think happiness is a tricky concept to pin down because is it happiness in the moment? Is it the sense that you have a year later looking back, and then so on.

Language is squishy and I thought completely specifically defined, but I think generally speaking better utility function and gradient to try to climb is that of fulfillment. I would say the experience of doing Stripe was not especially happy because we were constantly incredibly aware of all of the ways in which the product was severely deficient and all the challenges we faced

[...]

Error of all the ways in which the product was severely deficient and all the challenges we faced and, I mean, it seemed like there was no Fintech category back then. It was just, like, two teenagers ... just, so we're teenagers trying to compete with PayPal, which many people told us was not an especially promising avenue to pursue. So, not especially happy per se but it did feel kinda fulfilling. I enjoyed working with John and the kind of people we subsequently hired. It was really fun working with the kind of customers we were serving. They were just businesses doing all these kinds of wonderful things and they were really smart people, and it felt like that, if it worked, it could be consequential in the world.

And, so, I would say it's sort of felt like ... I mean, ... Well, you know, I don't know what it feels like to be sort of a scientist or something, but I'm guessing when you're pursuing all ... you have this kind of big question and you're pursuing all these kinds of avenues to sort of try to better understand or something, guessing day to day is not especially happy because most of your experiments don't work, or whatever, but perhaps there's some analog there in terms of it still feels, in some sense, meaningful.

Yeah, in some sense, day to day, like you said, you're just running around with your head cut off, but maybe on a weekly, or at least a monthly basis, it's just taking a step back and seeing what progress have I actually made, 'cause week to week it's one percent here, one percent here doesn't seem much, but on a monthly basis or maybe longer than that it seems like there's movement.

Yeah, and I think that's true. And I think it's almost certainly a bad idea to work on a company or to work on anything if you're truly unhappy working on it, right? And so there's clearly some balancing act there. And, I think, I mean a lot of ... Although there's so many different sort of difficult judgment calls to try to make on a start-up but, you know, of course, part of it is if something is making you unhappy or if it just doesn't seem like an especially promising avenue or it's not really working or whatever, your time has relatively high opportunity costs.

Running Your Company With Patrick Collison

You don't get to start that many start-ups in your life, right? And, so, knowing when to call it quits, I think, is actually ... I mean in start-updom we really extol and uphold the virtues of determination at all costs, you know? Never quit and so on, and that's clearly not the right answer. And sometimes you should quit.

And, so, I think what you're getting at is true and right for ... I think there does need to be some happiness, satisfaction, fulfillment. There's a good book of that name 'Satisfaction', which I recommend. Kind of week to week, month to month if you're on the right trajectory.

Oh man! I wish I'd known that. I could have ... I remember when I was trying to ... Stripe myself, I had these bad memories of trying to implement PayPal, and the stacks of documentation and taking five days to figure it all out. And so I sat down ... 'cause you guys are saying, "Oh, it'll take you 10 minutes", and I was like, "No way!" And it probably took me actually five minutes at that point, and I was super happy. Maybe I should have sent you a review.

Yeah, I mean, that would have been great especially if the review had criticism for us. ... But, no, it's definitely helpful to have competitors with not very good products.

Yes. Okay. So, in the early days you were doing a lot of stuff yourself. At some point ... I guess, two part question: one is: What are the things that you weren't good at? And then, two: At what point did you hire someone to, what I assume is, to do those things?

Well, I mean, I'm not that good at most things, right? And I say this in a non-false modesty, or I don't mean to be artificially self-deprecating, but for almost anything that Stripe has to do there are people who are better at that than I am, and so, I think, to some degree, the story of, say product market fit and to say our current stage is implementing the recognition of that in more and more areas.

That said, I think maybe embedded in that question is, okay, well, acknowledging that you're probably not the world's expert in any single thing, it's kind of, in what order do you, sort of, replace yourself? And, for Stripe, the important thing that we were most obviously not that great at, was partnerships and working with the various banks that we had to get on board and so on. And, so, we made a ... In fact, Geoff Ralstone, who is here in the room today, I think was some culmination [inaudible 00:25:10] coronation and, sort of, very supportive, and perhaps, in the back of his mind, slightly pitying where he saw us trying to get these partnerships with these big banks in place and we really weren't getting anywhere very fast.

And, so, he told us that we had to hire this guy, Billy Alvarado, and he told us: "Just don't ask me questions. Just hire this guy". And, at the time, everyone at Stripe was an engineer and so we couldn't quite understand what a non-engineer would do.

If you're not writing code, what else is there? And, so, we were kinda suspicious of this idea and we went back to Geoff ... We met Billy and he seemed like a wonderful guy.

Running Your Company With Patrick Collison

We really liked him, but we couldn't quite get past this ... Okay, how does this work out in practice? And, so, Geoff told us that we should hire Billy and if it didn't work for the first few months or whatever, Geoff would go back and pay his salary. So it was kinda, "Your risk move for us". And, so, we didn't have a whole lot of money and the time, so, that was not insignificant.

So, we hired Billy and that was, sort of, an immediately trajectory changing event where, previously when we'd gone and talked to a bank, I mean, I don't know if they are literally doing this, but it certainly felt that they were kinda pushing the security button under the desk. It was like: "Why are these guys in my office?" And, then suddenly things started to go kinda much more smoothly. And Billy is still with Stripe today and has been an enormously pivotal presence. And, so, that turned out to be very helpful advice.

That is interesting. Geoff as an inflection point for Stripe. Never knew. Cool, so it sounds like you had hired actually people before Billy. I guess your very first hire, third person with you and John, was an engineer, I'm assuming?

Well, he actually wasn't an engineer but he joined Stripe and he started to become an engineer. So, he'd written a little bit of code previously, but he joined and there was a lot of code that had to be written. It was a guy I'd known from college. He's actually also Irish.

We had a lot of code to be written and, so, he was the kind of person who would, sort of, survey the landscape and just do whatever it was that was most important and required and, at the time, it was writing code. And, so, off he went.

So, kind of the transition from pre-product market fit to post-product market fit, all I see here is when I think back, this is the one point in which they think, "Oh, I could have done that better". So, how did you grade your success of that transition? 'Cause a lot of it is just taking stuff, delegating stuff better, doing other functions of the business. So, what exactly was your transition like and how could I, if I was in your shoes, tell that I was doing it well?

Yeah, so I don't think I did it especially well, and I think fortunately it worked out. But I think if I was doing it all over again I think we could have accelerated ourselves by a year or two if we'd gone about it in a somewhat more disciplined and self-aware fashion. I think one of the most pernicious mental models you can have here is you are on some growth curve and it's your job to sustain or marginally inflect upwards or somehow ... You know, curling, the sport, where you're kind of wiping the ice while the weight proceeds down along ... Sorry, I'm not Canadian. But somehow you're kind of making these very small interventions and perturbations on some underlying growth curve. I think that's an easy mental model to have and, I think, it's actively kind of fallacious and mistaken.

Running Your Company With Patrick Collison

I think a much better mental model to have is you're serving some market and that market is relatively finite in size and then you can always change the progeny and increase the market size, but you can think of it as being finite.

And then there's sort of the percentage of the market that you're serving, and then whatever percentage you are not serving is kind of, well, you just haven't built the sort of go to market functions and organization that's capable or, at least, has yet, brought the product to those market segments. And the growth curve or adoption curve is just kind of a function of that go to market apparatus, but basically it's not some kind of cosmic trajectory. It's something kind of very much of your creation and under your control

And, so, what I probably did not do, but I wish we did, is maybe just after we launched there's a whole bunch of immediate scaling work we had to do. Let's say, six months after launch, we should have mapped out the kind of concentric circles of our market, like maybe, those kind of, you know, very early stage start-ups that were kind of our initial constituency.

Those kinds of all technical start-ups but not necessarily very early stage and you keep going in these successive increments until eventually you get to, say, all companies handling online payments or something, right? And gone and sort of figured it out. Okay, what's the size of this market? What fraction are we recurrently serving? What would it take to serve more? And so on.

And I think it's quite striking when you see repeat founders go and start companies they almost invariably are willing to build the organization post-product market fit.

But what's invariably willing to build the organization ahead of where things are today which, I think, is exactly the right thing to do because they're thinking, "Okay. I have the right product, now let's work backwards from ... Okay, what would the organization look like that was serving the entire market, and let's just start building that organization because, again, the growth curve is under my control". And, of course, it's not 100 percent under your control but I think it's much more under your control than some people might think.

There's also, of course, the difference here, a major difference, when I think about, sort of, consumer versus, kind of, B2B use cases. For a consumer it's a bit more, I mean, people don't necessarily know exactly what they want. What's the market size for a news reading app or a dating app or something? It, kind of, depends a lot on the product.

Maybe, where more people should be reading news or something. But for some service or product that serves a discreet, logical, concrete function, that, sort of, set up businesses or entities definitely have or don't have, I think it's, kind of, much more rational and much more mappable.

Running Your Company With Patrick Collison

I partly had this epiphany when Aaron Levie, who's the CEO of Box ... John and I eventually ... we started down at Pal Alto, we moved up to San Francisco ... Well, Aaron Levie had Facebook messaged John and we'd never heard of him and we hadn't heard of most people in Silicon Valley, but he Facebook messaged us very early on asking to invest in Stripe and, I think, John didn't know who's this random guy? And, so, we never replied.

But then we heard of Box and we heard of Aaron, and we read his funny Tweets and all the rest. And so, we moved to San Francisco and we invited him to our housewarming. I think it was the first time we'd ever met him. So, we're not very good at partying and by 11:00 p.m. or midnight or something everyone was going home, but Aaron was still there. And Aaron stayed until 2:00 a.m.

I still remember him sitting in our front room telling us how much better it was to be building B2B software than consumer software for this reason, where with consumer software it's so hard to predict what people want. They don't even know themselves what they want. It's such an amorphous space. Whereas when you're selling software to businesses, businesses are mostly rational, are mostly, I don't know, the opposite of inscrutable is scrutable, I guess, and you can work backwards in a way that it's just much more comprehensible.

And, so, I still have this kind of, it's like the angel and demon on your shoulder. I still have, sort of, at 2:00 a.m. Aaron Levie sitting on my shoulder extolling the virtues of building software for entities that know what they want.

You've made the right choice. Yeah, in some sense there's a lot less variables or in consumer there's a lot more variables to consider and they're quite unknown. I wanna take a step back. You talked a little bit about thinking ahead about what your team is or what your company structure is gonna look like. How do you ... maybe this is too big of a question, so maybe we can whittle it down a little bit. How do you think about that? If I'm starting a start-up today I am close to maybe product market fit. But before that stage, am I thinking about this sort of thing? What should my team look like? What should my culture be, and stuff like that?

I think that pre-product market fit, the goal really is just to [inaudible 00:34:21] your product market fit. So I actually wouldn't bother thinking too much about all these kinds of distribution mechanics questions. Now, you can get to product market fits relatively quickly and so that pre-product market fit stage might only last six months.

You're not necessarily, like Stripe, where you're, kind of, in it for various reasons for, say, multiple years. And, so, it may not last very long. However, until that point I really think you should just be thinking about, "Okay. How do I get there?"

The main thing, that I think, companies screw up at the pre-product market fit stage, is sort of speed of iteration. I think, that, okay, what determines or ... Speed of fruitful iteration, and if you're completely ... If you're repeatedly rebuilding your product but [inaudible 00:34:59] responses or feedback, I mean, that's far

Running Your Company With Patrick Collison

less likely to be bringing you in a productive direction.

But, if you have some kind of meaningful, albeit perhaps small set of initial set of users and you're rapidly iterating a response to their feedback and observed behavior and so on, then I think that's a really good spot to be in. And, I think, again, at that juncture pre-product market fit, you should be doing everything you can to tighten that feedback cycle.

There's a fighter pilot who revolutionized airborne combat in the US in the second half of the Twentieth Century, sort of Korean War onwards, called Boyd, and he had this concept of the OODA Loop, which was a similar notion as in a previously built [inaudible 00:35:48] he wanted the fastest aircraft or most sophisticated weaponry and so on.

Where he was all about, "No. You actually want your aircraft and pilots and training that are really oriented around the fastest responsiveness and iteration to the particulars of the situation".

In a way that subsequently went on to really inform modern aircraft design. And, so, I think you wanna be one of these modern fighters where you really optimized to respond as quickly as possible to rapidly evolving situations.

Again, in this, kind of, pre-product market fit stage, instead of from a hiring standpoint I think it should really be about, "Okay. Well, what's gonna get you there faster?" And, I mean, I think, at an early stage it's most likely, people who help you build the product.

But, of course, not too many because, I mean, at some point, you might be able to do more but you might actually be less responsive 'cause you've a bigger team to manage or something, right?

And so, as an empirical matter, teams that are somewhere between two and ten, depending on exactly what you're building, is kind of the optimally responsive size.

But, I think, it really is about that time observed necessity or deficiency or characteristic of your users' behavior to executed fix or improvement. And for whatever it is that minimizes that.

You say two to ten, which is helpful, but is there ... are there observations or evidence in which you have hit the peak? Like, you should not add an extra ... This extra person is gonna be negative value add to the whole operation?

Yeah, I mean, every person ... Well, start-ups in general involve all these impossibly difficult judgment calls in this high dimensional possibility space, and so it would be great if you could collapse it down to a fairly simple trade off, but I don't think you can.

However, I think, in principle, the calculations that you're trying to run is: Okay, each successive person takes time to hire and so slows us down in that regard; takes time to unmoor, slows us down in that regard; takes time to just, kind of, meld with the, sort of, culture and learn to stack all the stuff.

Running Your Company With Patrick Collison

That also takes cost and time. Then involves subsequent ongoing costs of coordination and alignment and the organizations that distribute it across more neurons until there's that kind of persistent tax and that's not just necessarily a linear cost, but it could be, sort of, quadrat [inaudible 00:38:22] or something giving you the multi-way communication problem.

So you've all these costs in additional persons so the question is: Okay, but this person can also help us get more shit done; they can write more code or talk to my users or whatever, and so it's, kind of, is that fixed benefit to that additional person worth all these other costs?

Again, with, I think the ultimate arbiter being, will we be more responsive to learning about our users given the presence of this additional person? And I, think, whether or not you will be depends on the complexity of the product and the complexity of the market and all that stuff. But, I think, that it's in principle the calculation you wanna be running.

Speaking of hiring people, I've heard you say that the key qualities that you look for in a future Stripe employee is intellectually honest, cares a great deal and just loves getting things done. Which are great attributes because some people just don't fit in those categories, so it's good to have the separation there. When you meet someone, how do you figure out this is the right person?

It's very hard and I wish I had a more definitive rubric and a particular set of questions. Although if I had a particular set of questions maybe they'd stop working because people would learn

how to game them or something. And, so, it's hard to fake just being smart, right? And, so, that one's not as hard to discern. It's oddly hard to fake being intellectually honest and, sort of, the character

Intellectually honest and sort of the characteristic of kind of being able to see multiple sides of a debate or of an argument. There are so many complicated questions where sort of the only thing that I'm really skeptical of is kind of certainty in either direction, just because the questions are intrinsically sort of involve very contingent trade offs. It's also, it's not impossible, but it's hard-ish to fake just being nice.

Like, something we hear a lot about here at Stripe is just sort of people who are pleasant and warm instead of just making others happier as a result of their presence. Well, perhaps if you can fake that perfectly forever, and we have time, right? But, yeah, they're all sort of actually there's no clear rubric for them and not sure that a clear rubric could even work.

Running Your Company With Patrick Collison

Yeah. So, you talk about, like, get people who help move the organization faster, if at all possible, and I think there are two issues usually that start slowing down the organization as you start scaling, just because each person adds just more complexity to the organization, but also, I think, one is asymmetric information, just people, like, are never on the same page. And I think, you've talked a lot about this, I think, so everyone can Google around for lectures you've given on this on, you've done email transparency. Obviously, having metrics transparent to the whole organization, but the second problem to that is, if you fix that, there's also this asymmetric interpretation problem, which is everyone's, there's this, like, black box function that how people interpret, even if all the inputs are the same, the outputs are all different and it's, especially at your scale now, it's nearly impossible to figure out everybody's function. So, when you're thinking about creating your organization and building it out, like, how do you reduce that noise and try to get everyone on the same page?

Well, the first thing is, I wish we had actually between say, five and 50 people. I think we were much too consensus oriented.

We, of course, weren't completely consensus oriented. I mean, we couldn't have gotten anything done if we were, but I think we kind of biased too much in that direction. It's just not that efficient, right? And it's necessarily not the case, and so, kind of, like, you consider, perhaps maintain some kind of fiction for more or less time, but ultimately speaking, there is no way of sort of sustaining that.

And I think that's a relatively common mistake, right?

Because you almost invariably come from some kind of end musketeers for some small end and sort of kind of orientation where instead of there's no particular need for sort of formal decision making mechanisms or sort of subsequent communication and set decisions or whatever, right?

But then you kind of hit 15 people and neither is. I think companies don't adjust quickly enough to sort that new necessity. Again very much us included. Until actually we did roughly well and the kind of ambient availability of certain context and information, so on. But actually kind of in some sense poorly at sort of deliberate, explicit communication of decisions being actual, kind of tactical decisions. Bigger decisions like what are our priorities for the next six months, or things of that nature.

And so the harder piece of advice, and perhaps over-extrapolating from our personal experience, but the harder advice just reflecting back would be to kind of - it's kind of like pre- and post- [inaudible 00:43:58] pride market fit thing, sort of actually about the size of the organization, where kind of, when you hit a certain size, I'm going to say 10 people for the sake of simplicity, it's a bit before, it's a bit after.

I think you need to adjust more deliberately to this kind of explicit communication model of being sort of quite firmly non-consensus based. Nobody likes the idea of sort of being hierarchical. That sounds pejorative, but in some sense hierarchical.

Running Your Company With Patrick Collison

Through the top down, it does move things faster. Maybe some people don't like that, but it does move faster.

That's right. And the reason is this kind of delicate balancing act. You want to sort of orchestrate where on the one hand you want to sort of really prioritize speed and agility which involves being somewhat hierarchical because those are sort of efficient, sort of symmetry breaking mechanisms and ways to sort of have shots get called.

But on the other hand you really do want people to have this kind of strong, ownership mentality and real sense that they can cause things to change or identify problems or sort of inject new ideas in unrelated areas.

So it's this delicate act where you definitely can be excessively hierarchical. How do you facilitate enough autonomy and agency but also not have things devolve into this, sort of uniform morass of [inaudible 00:45:33] motion?

I think that's hard and requires all of these ongoing tweaks and nudges, and is dependent on the personalities involved and all the rest.

Again, I really wish there was sort of the [inaudible 00:45:48] simple, well just do X, Y and Z and you'll be good. If such X, Y's and Z's exist no one's told me yet.

If only there was a formula for everything! So riding on the same theme, as Stripe grows, you strike me as a company that does the opposite of most companies, which as they get bigger they just slow down and they get

less innovative. For you guys, and it's hard to even keep up, like you're just pumping out new products, which seem to be successful to me like Stripe Ballast, Stripe Press, Stripe Terminals and so forth. In the early days where there was just 9 or 10 of you if someone had a new idea it's probably really easy to get to you and then you guys would decide. How has that changed over time to make sure that someone who's 6 degrees away from you with a good idea actually gets to you? How is that decision made to actually execute on it?

Well it's very nice to hear you say that you feel we're getting faster as we get bigger because we're constantly self-flagellating over how frickin' slow we are. Like paranoid about degenerating into, I don't know, some sort of immobile stupor. To whatever extent we do get things done or appear fast, I think it's largely because we're very paranoid in this dimension.

I think the default outcome of companies as they scale is to become more ossified and rigid and closed to new ideas and directions and things that contradict their prevailing orthodoxies and all of that. As to how we can avoid that, there's lots of kind of obvious things. Partly it's sort of having leaders who care about the rate of progress and just love seeing things happen quickly, and lots of things of that nature.

I think maybe, sort of deeper rooted, is we try to be a sort of, "Yes, and...?" Culture, in that I personally love ideas and potential things we could do. No matter how good the idea is, we should not pursue most ideas, right?

Running Your Company With Patrick Collison

Even if independently it could be a super successful company but there is a fairly finite number of things we can do, and on the one hand need to recognize that intellectually and not pursue most things while on the other hand, enjoying the exercise of contemplating what it would look like if we did it.

So one thing we do every year, for example, we have this kind of crazy ideas process where we literally send a document to the whole company, a hack [inaudible 00:48:38] document that's open to everyone and people can add ideas to it that they think are probably a bad idea but if they worked might be a great idea! And it's very important that they have to be probably a bad idea because if they're probably a good idea then it's not that risky, and kind of by definition is we should probably do it, and so whatever, maybe we should just do it.

People really self-center a lot in most organizations because they don't want to look stupid and they don't want to be associated with having these whacky, bad ideas and so on. So we try to be fairly firm that if you add ideas to this it has to be probably a bad idea. Actually a lot of cool things we've subsequently gone on and done first came from that process, Stripe Atlas being one of them.

We helped Stellar get off the ground, this cryptocurrency, back a few years ago. That also came from this process. That's one of the mechanisms by which we try to have that, "Yes, and ...?" culture. There aren't that many, just as an empirical matter, there aren't that many large organizations that really do this successfully, but I think that the most

successful, larger organizations somehow do succeed in sort of this iterative, repeated process of augmentation. Amazon and Google being the most prominent examples.

Obviously when Google launched there was no Gmail or YouTube or Android or Google Maps. Amazon is sort of even a more conspicuous example in some sense. The kind of repeated attach of successful ancillary businesses. There is a very natural temptation as you grow to become increasingly close to this. I think it's very important to avoid.

Cool! I have a couple more questions and then maybe we'll have time to take questions from the audience. One is, looking back, is there something ... Did you have a strong opinion of how startups should be run as a CEO that have just completely reversed because you are now the CEO?

That's a good question.

Another way to potentially put it is what are things that you did where you're like I know for sure this should have been done and they just turn out to be a mistake?

Well, I already mentioned the consensus orientation. One that we're going through right now, which is quite a significant divergence is we were fairly centralized up until relatively recently. We had some remote employees from a pretty early stage but by and large Stripe was concentrated in San Francisco.

Running Your Company With Patrick Collison

Last week we announced our 4th global hub and basically what we've decided to do is we're going to have major product and engineering efforts and teams and functions and all the rest in San Francisco, and also in Seattle, and also in Dublin, and also in Singapore.

These non-San Francisco locations are not going to be operations offices or satellite offices or places where we do some kind of localization and local market adaptation. We want to have a kind of completely Denovo new products that become super successful started out of these offices.

That's not the standard Silicon Valley pattern, where Apple and Amazon and Facebook and Google and all these companies tend to be highly concentrated in these monolithic corporate headquarters. Our thesis is several-fold. Partly that kind of availability of talent is becoming far more geographically dispersed, and partly the Bay area is becoming increasingly, untenably expensive location to locate.

We really want Stripe to be globally infrastructured, that works just as well in Asian markets or Latin American markets as it does for businesses in the US, and the era of the internet as being predominantly North American or North American/Western European or whatever. The days of being such a phenomenon are over, and so I think that's a fairly substantial break with descriptive best practice of the past. We obviously think it will work. We wouldn't do it if we didn't, but it is on some level risky. We don't have good prior examples to point to.

There's some great companies in Singapore like Grab and Carousell but there aren't really any examples yet of an American company establishing major [inaudible 00:53:43] and engineering hubs there. We're pretty optimistic it can be done, but if it works we'll be the first, or one of the first.

Last question. So, in a hundred years from now what is Stripe going to be? What do you imagine it to be?

Well we're only 7 years old, so that's a difficult question to answer! We're trying to build this economic infrastructure for the international.

This sort of platform for globalization and technological progress in the sense that it ought to be just as easy to start a company in Nigeria as it is in New York, and it should be just as easy for somebody in Brazil to buy something from any of these companies as it is for somebody in the western world.

It just seems so crazy to me that that just hasn't happened yet. It sort of feels like Stripe should have happened ten or twenty years before it did. I don't think that we're pursuing this unprecedented, inconceivable idea so much as correcting a deficiency, a rip in the fabric of internet infrastructure.

I think we still have at least 5 years to go in correcting this inadequacy, so what happens after that I'm not sure.

Running Your Company With Patrick Collison

In some sense, in the future all transactions should be digital and they could very well just be all going through Stripe. Right now in the US, someone was telling me like 80% of all Americans have done some transaction through Stripe.

That's right. In the last twelve months more than 80% of American adults have bought something from a Stripe business, at least one thing from a Stripe business, which is cool. And it's not just the case in the US. In Singapore, where I was last week, that figure is about 70%. I think of it more in terms of the things that are possible or got started. We think about the rate of new firm creation, what companies are getting started, how successful are those companies, which markets can they serve and so on. Every now and again we look back and we

look at those market coverage stats, but I guess that's really what motivates us. Its more sort of are these businesses should exist and/or should be able to offer their products and services in these places where they currently can't.

That's the kind of thinking we use to inform the product and is the core locomotion day to day, and maybe the outcome of that is all these people get to benefit from it.

All right! Thank you so much Patrick!

Read what you're missing:

[@farnamstreet](#)

<https://fs.blog>