

THE KNOWLEDGE PROJECT #32

Patrick Collison



Welcome to The Knowledge Project with Shane Parrish, the curator of Farnam Street Blog (<https://fs.blog>)

In [The Knowledge Project Podcast](#), I talk with interesting people to uncover the frameworks you can use to learn more in less time, make better decisions, and live a happier and more meaningful life.

On this episode, I have Patrick Collison, the co-founder of Stripe, which he started with his younger brother John in 2011. While Stripe started as a company to make online payments easier, it's morphed into an internet infrastructure company. Patrick is one of the most well-read and thoughtful people I've ever met. After listening to this conversation, you'll realize [that] his success is less about luck and more about thought. I'm pleased to have Patrick Collison on the show.

Shane Parrish: Patrick, I'm so happy to get to talk to you.

Patrick Collison: Thanks very much for having me.

You have the unique background of having dropped out of high school and dropped out of university. Can you explain what went through your mind dropping out of high school?

Well, I didn't technically speaking, drop out, although I sort of practically speaking did. But given my lack of education credentials elsewhere, I should, for the sake of my parents, insist that I did, in fact, formally speaking, graduate from high school. But I guess what happened is that I became very interested in programming, and I wanted to spend as much time on it as possible. Ireland actually has this interesting thing called "transition year," this year between two major exams of high school or at least Ireland's high school equivalent.

Transition year is a formally designated year that's optional, where you can go and pursue things that you might not otherwise naturally tend to pursue, and the school tends to be much more permissive of going and spending three months abroad or going and doing some work experience in this area or whatever the case may be. And so, in that year, I basically decided to spend as much of it as possible programming, and so I did that.

And then I returned to school for the latter half of Ireland's high school system, and it felt so much slower and less fun. As part of the programming, I had visited the U.S. for the first time. I had gone to Stanford for the 2005 International LISP conference, and it was a fairly small conference, but it was very eye-opening for me. I remember walking around Stanford and thinking, "Man, American colleges seem great."

Back in high school in Ireland, I decided to see if there was some way that I could just go to college in the U.S. the subsequent year. It was a long story, but I eventually figured out that I could not do it if I ... followed the standard Irish education path, but that I could do it if I did the British terminal examination. And so, I resumed my self-education, except instead of programming, I was now studying for these British exams, and did that for the subsequent year and ended up starting at MIT the next fall.

How do we get from MIT to where we are today, which is Stripe's offices in San Francisco?

It's sort of a long and torturous story, and I'll spare you most of the less interesting details. I guess the overarching thing is, while people in the U.S. tend to have grown up in an environment in which college attendance is really prioritized from an early age, and where you're optimizing your extracurricular activities from the time you're 14, and where you're choosing your kindergarten on the basis of what the downstream college acceptance rates look like, and all that kind of stuff. Growing up in Ireland, that wasn't part of the culture or discourse or environment at all.

And so by the time I got to MIT and to college in general, it didn't feel like that big a deal; it didn't feel like this was the terminal state that I had spent my entire childhood and adolescence trying to reach. And so as other things and other ideas and opportunities crossed the transom, I think I was more open to them than my peers, not because of, I think, any differences in me, but just because of differences in the cultural environment I came from.

My brother John and I—John being a little bit younger; he was now in this transition year in Ireland—we decided to start a company six months after I got to MIT. I'd really just started, and I felt that I had some time to spare because I'd started college a year younger than most of my peers. That company worked okay, and it's kind of a long story. It ended up becoming a small acquisition.

I went back to MIT because, when I'd started there, I had been very interested in math and physics and had an interest in this idea of potentially becoming, or at least attempting to become, some academic, and of course at a place like MIT, that's sort of the default around you. Everyone is trying to get a PhD or to become a professor or whatever. And so I think that environment had some effect on me. And so I went back because I felt that I hadn't really properly rejected the hypothesis that maybe I should try to become a professor; maybe physics is what I should be at least attempting to spend my career on. And after a year back at MIT, I decided that that was not the case.

Progress in physics really felt like it had slowed down pretty substantially compared to the 1910s, '20s, '30s. The period in which so much of what we were learning about ... that sort of broader period of discovery, it felt like the period in which we existed in, say, 2010 was ... just not the same rate of progress. And so there was a little bit of that, and then also some amount of appreciation of myself that I think I just enjoyed programming and software and technology more than I did math and physics, even though to some degree, it was a little bit painful to realize that.

I want to explore a little more about the cultural differences between Ireland and the U.S. and how [they affect] you as the CEO of Stripe.

I think that there are maybe a couple of things.

Ireland is very outward looking—necessarily so, in that Ireland's improbable rise from poverty over the latter half of the 20th century was very significantly enabled, maybe almost wholly enabled, by exports, by importing American multinational companies, having them set up factories and bases and hubs of different sorts in Ireland.

One of the world's first special economic zones was created in Shannon, which was very close to ... 10, 15 miles from where I was born. Deng Xiaoping visited us and found this quite inspiring, and so decided to set up special economic zones in China. And so Shenzhen and the Pearl River Delta—that special economic zone was in some ways directly inspired by what he saw in western Ireland. I think the fact that there's such a very visceral link between betterment and progress and economic development and this kind of outward-looking sense that the possibilities of the rest of the world are much greater than those internally—that's very pervasive in Ireland.

And I think that's certainly influenced Stripe in the sense that we really are all trying to emphasize the imperative and the potential of globalization. And while maybe in the mid-'90s, that was something that was uniformly accepted, at least in elite circles, now we see that's something that perhaps is being questioned somewhat more, but I guess the Irish experience is very much one of seeing it as an almost wholly unalloyed good. And again, I think that greatly influenced us here, certainly me.

It's interesting, too, from a cultural standpoint, where Ireland has had very high rates of immigration, particularly post the expansion of the EU in 2004; a very large number of Eastern European immigrants moved to Ireland when those countries acceded to the EU. That was really not accompanied by any material social strife or conflict or a lot of the challenges that we've seen in other parts of the world. And so again, I think that an appreciation for borders that are more open, or more openness to immigrants, more facilitation of opportunity, things like that, again, I think that really is the Irish experience.

And of course there's the reverse version, where so many Irish people themselves have benefited enormously from being able to go and pursue lives in the UK and Australia and the U.S. and Canada and so on, and that's

“[Organizations] at least inadvertently, if not deliberately, prefer cohesion over correctness.”

again, just really been a part of the national ethos.

And then maybe more softly, I guess, Irish culture places a lot of importance on just a kind of warmth and... There's a particular tenor to interpersonal dynamics and trying to have other people enjoy themselves and be at ease and have a good conversation and whatever else. I think maybe that's something that's influenced us somewhat at Stripe, where we want Stripe to be a warm place. We play music at reception and in the kitchen to just try to put people at ease and to create enough soft noise around them where they feel comfortable having just a good conversation. Maybe that's because of entirely unrelated reasons, or maybe again, in some way we were influenced by the kind of environment we grew up in in Ireland.

How would you describe the culture at Stripe? What do you actively try to achieve with that?

I'll answer that with a caveat, and the caveat is that I'm pretty sure the answer I would have given to this would have differed in some material ways two or three years ago. That's in part because I think we're coming to realize things that we just hadn't really appreciated or seen the significance of two or three years ago. Also, in part, because literally what it is that we need today is just different from what we needed two or three years ago. And so I think there's a double contingency in the answer, where it's a function of just what we've realized at this point, but also what it is that the organization and the company needs and the challenges that we currently face.

With that caveat, I think there are a few things that we really prize and try to seek in the people we hire. First, a kind of rigor and clarity of thought. So many organizations prize smoothness, smoothness in interactions and trying to reduce or minimize the number of ruffled feathers. And they at least inadvertently, if not deliberately, prefer cohesion over correctness, and we really try to identify people who are seeking correctness and who don't mind being wrong and who are willing to at least contemplate things that seem improbable or surprising if true or really divergent [from] that which is the generally accepted status quo.

And that's hard to find. I don't think most of the educational institutions that we all tend to have attended do a great job of teaching that. We look for that combination of openness and rigor.

Next, I don't exactly know what the right word is, but a determination and competitiveness, and I guess willfulness, in that ... just doing anything of significance is hard. Anyone who's tried to do anything that they themselves considered significant knows that very viscerally. And especially for a startup, the default outcome is your relatively near-term non-existence.

The default outcome is that you do not survive; and to survive over the medium or, even with more difficulty, over the long term, that is like an unnatural act. And so you need to find people who not just are willing to push against the expected trajectory of non-existence, but people who actually

“I think the three that really stand out to me are this rigor and clarity of thought, this hunger, appetite, willfulness, determination, and this ... warmth and desire to make people around them better off.”

enjoy that, who want that. Because if they're merely willing to do it but they don't actually enjoy it, then the work is probably going to be less fulfilling for them over the medium term.

And I really don't think that is for everyone and I don't think that's a bad thing. The cliché, of course, is that startups are extraordinarily hard, and they just are. You want somebody who is at a stage in their life where that's the kind of challenge they want, where the fact that the particular area in which they're going to be working is undefined or significantly under-built-out or significantly broken or whatever the case might be, that that's what they're looking for.

And then we try to find people who just have—again, to return to these words—interpersonal warmth and a desire to make others around them better and just a degree of caring for others and a desire to be, nice is kind of an anodyne word, but to be nice to them and to make them better off.

We really try to find people who you just actively enjoy spending time with. You spend such a large fraction of your life inside the walls and under the roof of whatever organization [or] institution you're working at, and so, given that I really think it's worth prioritizing this and I think, I of course don't know for sure, but I think we go to some greater lengths to find these people than other organizations tend to do. There's other things as well. It almost goes without saying, but we really care a great deal about ethics and integrity and people, but I think so too do a lot of other organizations.

I think the three that really stand out to me are this rigor and clarity of thought, this hunger, appetite, willfulness, determination, and this ... warmth and desire to make people around them better off. Those are three that really stand out to me.

Take me back to the early days of Stripe and the struggles you were having; maybe walk me through some of the things that you've learned since then or some of the mistakes that you had made.

Sure. The background context here is that under almost every ostensibly sane analysis, Stripe looked like a bad idea. This was a crowded market, there were tons of existing incumbents, there were significant regulatory and partnership/institutional barriers to entry. We had no experience in the domain. We were very young. We weren't even U.S. citizens, and an ecosystem that again, just because of the regulatory dynamics, that sort of adds to the complication. We had no obvious mechanism for gaining significant distribution, and we were not a naturally viral product or one with an organic adoption the way maybe a social network or a consumer product might have.

And so for all those reasons, I think a lot of people very reasonably thought that Stripe was a bad idea or us pursuing Stripe was a bad idea, and they certainly didn't hesitate to tell us that. To be clear, I think they were doing something reasonable by doing so. They were giving us their honest and ...

“If there is such an obvious idea and such a surprising absence of a solution, maybe there’s some kind of latent force that we’re not seeing that actually makes solving it impossible.”

reasonably justified assessment. So, it all started within the background context of that.

I think the thing that primarily gave us the confidence to actually attempt it was [that] it just seemed so strange that something with Stripe’s character didn’t exist—in that we really looked for a Stripe before we started it. It felt that it must be the case that there was some service, some company, somewhere offering infrastructure and APIs and payments and economic tools that are straightforward to use for a developer. This is one of the top needs that any business operating on the internet has. Arguably by definition in being a business on the internet, [it] must have access to these tools. And there are tens of millions of developers operating on the internet. Given the magnitude of that market and the obviousness of the business model, it really felt like it had to exist. And so we had forlornly Googled for it with different permutations of keywords, and then after a couple of months, became somewhat resigned to the fact that no, it did not, in fact, ... probably exist.

Its non-existence was so strange to us that it actually initially kind of discouraged us at first. If there is such an obvious idea and such a surprising absence of a solution, maybe there’s some kind of latent force that we’re not seeing that actually makes solving it impossible.

For example, we were also interested at the same time in why consumer banks were so bad, in that they weren’t really keeping abreast with technology and the fees were really high and they were getting fined by the CFPB, etc., etc., etc.

And as we looked into it, it became apparent that actually there were at least two good reasons as to why the problem had not been solved. One, the banks are subject to such onerous regulation such that it’s very difficult for them to do anything themselves. For example, the differences in a checking account and a savings account, which might seem quite unfriendly from a consumer’s standpoint, that’s actually essentially mandated by law, and so it’s not all the bank’s fault. The second reason is, the office of the Comptroller of the Currency, which is the entity that issues federal banking charters, had basically stopped issuing new banking charters due to the financial crisis, and so, if you came along, [and] you’re like, “Well, I’m going to go solve all these problems in consumer banking,” you were essentially blocked from doing so by the regulatory apparatus.

And so we wondered, in a similar vein, is there some force like that in online payments? Not necessarily regulatory, but were there some constraints that we just weren’t observing.

After a couple of months of investigation, we decided that no, there didn’t appear to be. Of course that you can never definitely reject it—but we really couldn’t find one. And so we decided to build a prototype. The prototype was going to build up on top of existing payment systems, and so it didn’t do anything overly ambitious. It was just enough to get a sense for what... It was almost like a concept rendering of what a solution could look like rather than

“Too many times when you don’t see what the problems are, it’s usually because there’s kind of subjective blindness rather than because it’s actually difficult to see the problem.”

a solution itself. But it was sufficient to get just a couple of our friends started using it.

I think the particular thing we realized that caused us to really take it a little bit more seriously and concretely to drop out of college was the realization that this problem that we perceived, and developers like us, people building some little side project or with this very nascent startup or something like that, that the problems we perceived for that segment of the market were actually the problems that larger companies had as well. That what we thought initially might be a little lake of opportunity was more akin to an ocean, and when we talked to companies with hundreds of millions or billions in revenue or companies in other countries and so on, and we just asked them to recount their problems and what they wished existed and everything else, they basically gave us the same roster of features.

We thought about it and just looked to the macro figures; we saw that about, at the time, 2 percent of all consumer spending in the world happened on the internet, even though we were 20 years into the web’s evolution and even though we’d all engaged in lots of e-commerce and so on. When you looked at it on a macro basis, it was apparent that we were still barely off the starting blocks. I think the combination of those things, where we decided that there didn’t appear to be some dark energy preventing a solution, and that the set of problems we could see actually seems very pervasive rather than just a microcosm, and then thirdly, that actually this whole market and environment was still actually at a sort of surprisingly nascent stage when you looked at the full picture, then we had to drop out.

You guys went from two employees—you and your brother as co-founders—to 800, 900 now?

About a thousand now.

A thousand employees. And what have you learned from scaling the business?

I think on some level, scaling a business is both relatively straightforward and extremely hard. It’s relatively straightforward in the sense that it’s usually not that difficult to see what the problems are, and too many times when you don’t see what the problems are, it’s usually because there’s kind of subjective blindness rather than because it’s actually difficult to see the problem. And so it’s more of a question of “What are you oblivious to because of your own biases?” rather than “What is particularly difficult to observe and what are your corrective mechanisms to account for that?”

So I think straightforward in that sense, and I guess also straightforward in the sense that usually solving the problems is not outlandishly difficult. It’s not easy, but you have to hire someone in this role, you need to figure out how to raise this capital, you need to build this system, whatever the case might be. None of those are easy things, but they’re also not scientific breakthroughs. There are other companies that have done it. There are generally playbooks that exist, and while your particular strategy might need

“You don’t get to really choose the clock cycle and the time horizons.”

some correction or refinement and you might hit some walls along the way, it’s rarely unprecedented.

And then I think it’s extremely difficult in the sense that you don’t get to really choose the clock cycle and the time horizons. There is a category of flash games, desktop tower defense games, where you have all these little critters scampering across the board trying to break into your fortress or whatever—I start to feel a little bit like that, where you fundamentally don’t control the rate of problem appearance. You just control the other variable, of the rate at which you are building defensive or mitigatory mechanisms to deal with those problems.

And sometimes the rate of the problem creation can outstrip the rate at which you can solve them, even though in principle any one of them is relatively manageable. I think that really adds a lot of difficulty. And then, while on this very abstract level, dealing with the problems is tractable, the character of having problems materialize at every level of the organization and at every level of abstraction and at every magnitude, and so on, that’s just a kind of unnatural thing that I think is, just on a psychological and emotional level, difficult to deal with.

And so while you might recognize on some contemplative, stoic level that this is how it goes, that’s not necessarily how it feels in the moment. Some days you almost have to smile at the unreasonableness of the swath of problems and challenges that have materialized on your desk or in your inbox. In the same way that you see the constellations in the stars, the problems look so implausible and so unreasonable that, like, someone must secretly be screwing with you. And so there’s that kind of emotional self-management.

Then, of course, there’s the challenge of dealing with uncertainty where it’s, I guess... You’re operating in sort of the weird zone where you’re often making decisions that have significant long-term impact, or that are at least difficult to reverse or to course-correct, in the face of great uncertainty. And the uncertainty is often unnecessary in the sense that you could in principle go and significantly reduce the uncertainty. You could go and study the question more, you could go and obtain more information, you could go and run an experiment. It’s not like cosmic uncertainty, where it’s true Knightian unknowability. I think when it is true, deep, unmitigatable uncertainty, then I think it’s not too hard to say, “Well, we’re just gonna choose something and make the best decision we can.” I think it’s a more frustrating kind of uncertainty, where it’s actually not necessary, but the thing that’s limited is essentially the cost of obtaining further information, reducing that uncertainty. And so you’re left in the dissatisfying situation where you have to make a highly consequential decision [and] there’s a lot of uncertainty. “We could have less uncertainty, we could take steps to mitigate that, but we just don’t have time to.” Making a lot of decisions in that zone is somewhat dissatisfying, and I think correctly so, when you’re correctly reacting to the fact that it could be otherwise.

“The thing that’s limited is essentially the cost of obtaining further information, reducing that uncertainty.”

And then lastly, maybe you’re playing this multi-armed bandit problem where you’re constantly trying to balance exploration and exploitation, or optimization of that which already exists, and doing it better and better while trying to figure out “What are the things that we aren’t doing or that we don’t know or we haven’t even considered or, if we were doing, would make this other part of the organization vastly more effective?” and so on. It’s very hard to know what the optimal rate of exploring those things is while also operating inside the system, or optimizing outside the system and optimizing inside the system.

Again, I think a lot of the challenge of scaling the organization is finding, at each kind of moment, the right way to balance those things. But without ever having sat down before to try to distill this into any unified theory, I think that a lot of the experience of scaling an organization is specific versions or specific applications of those dynamics. And just figuring out how you yourself, or how your organization, or how your peers and colleagues deal with that, and what kind of structural mechanisms for doing so are.

And then maybe very lastly—those are all kinds of structural ones—I think there is just also a personal version, where you certainly don’t start out being well adapted to or, at least in my case, particularly skilled in organizational management and leadership. Depending on the rate of growth of the company, you need to acquire those skills on ... a timeline that’s largely out of your control, and depending on the rate of growth of the organization, that might be a pretty difficult thing. And so, certainly in my case, I think I’ve just had to accept my managerial inadequacy relative to what either is required in the moment or will in the near-term, impending future be required, and to just figure out strategies to try to acquire those skills and abilities as rapidly as possible.

I want to go back to the explore/exploit kind of comment that you made, which we can probably just relate to focus. How do you think about focusing on one thing and being exceptional at that? Or doing a variety of things and trying to be exceptional on all of them?

You mean in the organization or personally?

Oh, in the organization and maybe personally, if that’s different.

I don’t know the better answer other than using coarse heuristics and then being willing to revisit or make an exception if something seems particularly promising. Roughly speaking, we invest most of our effort—I don’t have the precise number at hand; let’s just say 70 or 80 percent—in optimizing that which we already have, that which we already know is producing returns, that which there is a relatively clear line of sight from the input, the work, the optimization, whatever, to the output to improvement. And then some fraction of the work and the distribution of bets skew... but some fraction of the work, let’s call it 20 percent, into things that are more speculative.

“It’s very easy to fly the company into the side of a hill.”

I think it’s necessarily the case that, say 70 or 80 percent, is devoted towards optimization of that which already exists. If we did not do that, then this default non-existence we just discussed would be guaranteed. It’s very easy to fly the company into the side of a hill. And so I think the question is just, “Do you spend 20 percent of your time on things that are more speculative or do you spend 0 percent?” And then maybe, secondly, to what degree do you allow those answers to be different at different levels of the company in different places? And how much of a uniform answer, and how much heterogeneity, do you permit or do you design for?

As we’ve grown, we’ve tried to shift into a model where it is somewhat less uniform. In certain teams, less optimization of what already exists is going to be required; it’s going to require more exploration. And in other parts of the company it could be tilted in the reverse direction. And I think that kind of recursive decomposition ... is really required to avoid the diseconomies of scale that otherwise set in as you grow.

How do you decide which speculative projects to take on? Are they based on disrupting your business? Or, “These are things that I want to do or I want Stripe to do or ...?”

I don’t know that there’s a better answer beyond, given all of the axes of constraints and returns, which ones seem like a good idea? I think it’s like investing, when you’re asking, “What are the criteria for investing in a company?” When you normalize down from this really high-dimensional space of market and founders and ideas, all these things, you normalize all that down into “What do you think the return profile looks like?” Well, you invest when the return profile looks good enough.

I think similarly, when you decide which ideas to pursue, of course, on each axis, there are many things you prefer, or you don’t want, or whatever—for example, something that requires less effort rather than more, or entails less downside risk rather than more.... Those are all good things, but I think where it all nets out is, well, when you take account of all those factors, the question is just: which things just seem like a good bet?

And so, just to give a concrete example, Atlas, the service we launched for helping new founders incorporate companies, in particular without the geographic restrictions that tended to exist before — it’s essentially open to founders anywhere in the world. There was no ... one reason as to why that was a good bet. You can’t just measure that on any one axis. But when you look at it overall and you see that, “Well, if it doesn’t work, it’s hard to see how it could cause that much downside for Stripe. It’s not going to require an enormous fixed-cost investment in order to learn as to at least whether it’s initially working. If it did work, it seems like it could produce quite significant returns. The kinds of things we’ll have to do for it directly are things that are probably valuable for us in other parts of the business, so we’ll learn interesting new capabilities and skills in the course of doing it”, etc., etc. I think the reason there aren’t more good bets made in the world is because making good bets is difficult. And again, I think in different areas...

“Newer things tend to operate on faster clock cycles.”

Difficult in terms of recognizing them or difficult in terms of acting and executing on them? Or what do you mean by difficult?

Both. I think most organizations are institutionally resistant to bets because most people are necessarily optimizing things that already exist. And again, that's correct. Things that are not optimized along the way, especially things that are not being fixed and patched up and corrected as they burgeon, those are going to break. And so, the optimization is critically important. I don't mean to sound remotely dismissive towards it, but bets are a very different character. And this sort of a continuum of betfulness and riskiness....

Large institutions and incumbent organizations dislike bets, structurally speaking, and find them difficult to understand and difficult to interact with. People in startups are less worried about the risk of failure, whereas people in existing systems must worry quite a bit about the risk of failure.

Newer things tend to operate on faster clock cycles. Dijkstra talked about the idea of the Buxton index and the time horizon upon which an organization makes its decisions. And so, maybe a university makes its decisions with over a decade-long time horizon, whereas maybe a company makes decisions on a quarterly time horizon; maybe an individual makes decisions on a weekly or monthly time horizon, whatever. The observation was that organizations with very different Buxton indexes find it difficult to work together.

And if an organization with a really long time horizon is working with one that's rapidly updating and rethinking, it's just like a fundamental kind of impedance mismatch. And so I think that, to your question, instead of why it's hard and why there aren't more good ones in the world, I think there are lots of different kinds of impedance mismatches like that. It's not just the time horizon thing, but I think there's just ... a fundamental deep intrinsic difference between existing incumbent systems and the actions and the minds that are required to optimize them, and the exploration of figuring out that which is totally orthogonal different and new.

How do you keep the mentality... When Stripe started, the cost of failure was really low. Now you have a thousand employees, they all have families, you have a business, you have people who have invested a lot of money in the business. How do you maintain that ability to place massive bets?

It's really a question of how do we make sure that we can place bets that don't have excessive downsides or ... fatal downsides or cumulatively fatal downsides across a whole portfolio of bets? And I think that actually... I think the impediments to placing good... Again, I'll caveat all this by saying, it's not like Stripe has a long track record of making really good investment-backed decisions; we are far from being the Apples or the Berkshires, or whoever. A multi-decade track record...

We'll come back here in a decade and we'll re-evaluate.

If we're here in three decades, which, as established, would not be the default

“How do we make sure that we can place bets that don’t have excessive downsides.”

outcome, and we have a great portfolio of successful such decisions, then perhaps we can opine with some modicum of confidence. But it feels to me, and we’ll see if it’s right or not, it feels to me, that actually the reasons that organizations don’t tend to make more of these or make more good ones is, it’s more sociological, more institutional, and less that it’s fundamentally too costly.

Because in most cases, the downside cost is not that large, either in terms of direct financial cost or in terms of the potential broader damage to the organization, whatever form that might take. It’s much more that the mindset of improving that which already exists is just quite different to the mindset of, “Screw the old system, let’s do something that’s fundamentally new from scratch.” And so I think the challenge is, in significant part, how do you reconcile these two mindsets?

Stewart Brand talked about “pace layering” in buildings—different parts of the building have to change at different rates. How do you design for that?

And I think the analogous question for an organization is, how do you do organizational pace layering? How do you have parts of the organization that can try to do something fundamentally different [from], and hopefully superior to, that which already exists? And how do you have people who basically disagree with you trying to do something new, who think that “No, the way we’re currently doing it is in fact the right way; we’re just going to do it better and better.” And because these people fundamentally structurally disagree with each other and must have significant conviction in their respective approaches (otherwise, they won’t do great work)... how do you have those people at the end of the day have dinner together and fundamentally feel like they’re on the same team?

How do you do that?

Come back in 30 years.

I think I recall one of the interviews that I was watching as prep for this, where you talk about one of the first five or six people worked at Bridgewater?

No. One guy in particular did. And over time, we’ve hired more people who have, but I would not say we were particularly Bridgewater influenced.

Did you command the notion of thoughtful disagreement before that influence? And if so, how did you ...

Well, it’s hard to know exactly where to attribute it and it’s probably over-determined. And maybe there is some underlying personality traits that we each had come to in different parts of our lives, in somewhat coincidental ways. For a start, to your earlier question, Irish people are always disagreeing and always arguing. And so again, maybe it’s the cultural dimension to it. It’s not something that people tend to shy away from.

“I think one deep mindset difference in people is often [between] those who enjoy finding the limitations of arguments and beliefs and those who don’t.”

Because they didn’t see it as an attack on them?

Exactly. Right. I think there was just a common shared personality trait in a lot of the people who helped establish the culture of Stripe, where they enjoyed disagreement and trying to find the boundaries of an argument — the places where it doesn’t apply, and what the exceptions might be, and just trying to get a feel for the topology of that space. Stumbling in the dark, trying to construct a map of where different intuitions and heuristics apply and where they don’t and so on.

I think one deep mindset difference in people is often [between] those who enjoy finding the limitations of arguments and beliefs and those who don’t. Tyler Cowen talks about—I think it’s his second law—that there are no knockdown arguments and there are no arguments that are just uniformly completely true. There are always the limits to it; there’s always the other side. I think that’s very deeply true, but I think there’s just a question of affect and personality.... do you enjoy finding those limits and the exceptions and thinking about, “Well, maybe this is less true than I think” or “where is it less true than I think?” Or is that just ... a stressful process?

And I think that getting that rigor and clarity of thought requires a joy of discovery. Like, “Ah, this thing I believe, this rule that I thought existed, it’s actually not good in this place.” And having that be an enjoyable discovery rather than something stressful and threatening. I think globalization is a good example there, where, as we discussed, I think that globalization is a net—overall, for the world—fantastic thing and something that support is rising for on a global basis and [that] has propelled more people out of poverty than almost any other force ever. And yet, there are people like Dani Rodrik, [David] Autor at MIT, and others who are prodding at the edges of that, showing “well, but not in this place” or “not in this way.” “Maybe it has this underappreciated downside.” I think that’s great. I think those are important questions and really interesting work. ... Again, the underlying sentiment is interest in where the heuristics and the intuitions and the rules and the arguments are wrong.

I want to come back to some of that a little bit later. I think one of the questions that people want to hear from you is, what would you say is the biggest difference between the Patrick making decisions today and the Patrick making decisions maybe five years ago, in terms of how you actually make those decisions?

I think there are four big differences. The first is ... I now just place more value on decision speed. If you can make twice as many decisions at half the precision, that’s actually often better. And then, given the fact that the rate of improvement of decision making with additional time almost necessarily tends to flatten out, I think that most people—certainly that Patrick of five years ago and partly from Patrick of today included—should be operating earlier in that curve. Make more decisions with less confidence but in significantly less time. And just recognize that in most cases, you can course-correct and treat fast decisions as a kind of asset and capability in their own

“If you can make twice as many decisions at half the precision, that’s actually often better.”

right. It’s quite striking to me how some of the organizations that I hold in the highest regard tend to do this.

The second thing is to not treat all decisions uniformly. I think the most obvious axes to break them down on are degree of reversibility and magnitude. Things with low reversibility and great impact and magnitude, those ones you do want to really deliberate over and try to get right.

But I think it’s very easy, absent care, to have ... the mechanisms you put in place for those decisions to seep into decision making for the other categories. And really in the other three quadrants, you can afford to be much more flexible and much more fluid. And again, to really just prioritize speed — because obviously, if it’s very reversible, then, by definition, you can always correct it later, and if it’s of low import, then who cares? So, that’s the second one: just being cognizant ..., and before making the decision, try to categorize, “Well, what kind of decision is it?”

The third thing is, I now try to fairly deliberately just make fewer decisions. “Why am I making the decision?” And for some kinds of decisions, there are some good reasons for that, and there are some decisions the CEO ought to make and is fundamentally on the hook for, but there are some decisions where if I’m making it or if I have to make it, that probably suggests [that] something else organizationally or institutionally has broken.

And I think the need for a decision from anyone, not just for me, is often ... only an epiphenomenon. And there’s really some other underlying issue that’s causing you to have to make it in the first place. So, I try to think about that, and concretely, do more to push others to make decisions, and push them back to people who ought to be the domain experts.

And then fourth, when I realize that I would make a decision differently [to] how someone else would make it, not even really discussing the decision itself but trying to dig into, what is the difference in our models such that they want to make Decision A and I want to make Decision B?

One thing we’re currently spending time on here at Stripe is having different parts of the organization write down what they’re optimizing for essentially. Like what their mission is; what the long-term key metrics are for their part of the organization; who their customers are, either internally or externally. The persistent ongoing underlying nature. Hopefully, once there is agreement on those longer-term things, then maybe a difference on any particular decision might just be “well, we differ on what the most instrumentally effective way to achieve this outcome is, but we’re also both really unified on what the desired end state is.”

And there I think disagreement over instrumental efficacy, that’s really the much less problematic disagreement, because, if you’re right, then we’ll soon learn that, and if you’re wrong, reality will probably make that pretty clear in short order. I think the more troubling ones, and the ones that tend to cause more persistent friction in an organization, are where there is latent disagreement on what you’re actually optimizing for, but that’s never explicitly

“One thing we’re currently spending much time on here at Stripe is having different parts of the organization write down what they’re optimizing for.”

surfaced and uncovered. And so, now, in decision making, I place more importance on making sure that we have the right foundational agreement such that the disagreements that then tend to arise are of the essentially more superficial sort and where agreement is actually less important.

Part of culture is learning from the decisions the organization makes. What do you do at Stripe to make sure that people are learning, and what do you do personally to make sure that you’re learning from the decisions that you’ve made, both positive and perhaps ones that you, in retrospect, would have wished you could make differently?

I’m inclined to say — I don’t know if I actually believe this — but I’m inclined to say in response to that question that decision making in organizations is slightly overrated, in that organizations are not like investment entities or funds or managers. With investing, it’s fundamentally very binary. There is a moment at which you either buy or don’t or sell or don’t or whatever. And maybe it’s somewhat more continuous in the case of, say, public market investing and so on, but given constraints and decision-making time, I think you have to treat it as more binary. You assess the stock and you make a buy-or-not decision.

Whereas in organizations, everything is more fluid and continuous; it’s much more about designing the feedback mechanisms...

Or biological?

Yeah. Exactly. And there’s the famous [Phillips] water model of the economy, sort of the circulating fluids, and you can vary the interest rate or the inflation rate or whatever, and it helps you understand the underlying dynamics. And I think an organization is much more like that, and so I think the things to optimize are the incentive structures and the mindsets, and the definitions of the goals, and the feedback mechanisms, from the outcomes to the inputs, and the work, and the operations themselves, and all those things, and less the binary decisions.

I don’t completely dismiss, obviously, the importance of decision making and that there are times where we decide, “Well, are we going to launch the product or not?”, “Are we going to start that project or not?”, “Are we going to replace this system or not?”, and so on. There are of course real decisions, but I think it tends to be much more... I guess, maybe it doesn’t feel like the right unit of analysis to me. I think the right unit of analysis is that of the cell, and the question is, well, in an organization, what are the cells and what are the organs and how do they interact? What are the feedback mechanisms between them?

Let’s geek out a little bit on the feedback mechanisms here. What sort of feedback mechanisms do you try to make sure are in place? At what point in the process do you try to acknowledge what they are?

I really think that, and this is not to evade the question, but I really think it’s too early to answer that, in the sense that I can tell you what I think today and

“We spent a lot of our time scrutinizing other organizations, trying to find out and kind of reverse-engineer what works for them and why.”

the changes we've made over the last year and things like that. Stripe has been a thousand-person organization ... or has been a more-than-500-person organization for just over a year. We're beginners at this! Three years ago, Stripe was under 100 people. And I think either to opine as if or, to even more problematically, believe that we kind of have it figured out, would be ... hubris.

And so ... what we've been talking about, I think that's maybe some of where our, or my, thinking comes from, but I don't know what the right answers are yet. And we spend a lot of our time scrutinizing other organizations, trying to find out and kind of reverse-engineer what works for them and why. And I think that part of what's interesting in the tech industry is that, it's kind of pure knowledge work. We're still, I think, quite early in ... figuring out, in terms of how to optimally coordinate and collaborate on it, in that you can draw a lineage of HP, and Intel and Microsoft and Google and Facebook and so on, WhatsApp. And they're all these ... examples that I think suggest that we may not have it all figured out.

The fact that WhatsApp was such a minuscule team, and Instagram too of course, despite operating at such scale, or the fact that...

The weight of a new paradigm?

Yeah. And the way Facebook operates is very different [from] the way HP operates.

At Stripe, which company cultures do you admire the most? Not business models, but cultures, and why?

Well, I admire cultures that are strong, first off. Culture that, when you ask somebody “Can you describe it?”, they can expound on its merits for more than half an hour. And, in almost every case, describe at some length all of the things they don't like about it. Because if it's strong, it's improbable that every aspect of it is something every person really agrees with or feels an affinity for. And so whether it's The New Yorker or the military, a shared characteristic of those cultures is that they are strong. I think that's the first order of things. And I don't think that describes most organizational cultures. I think most organizational cultures are some kind of milquetoast common denominator.

Second, I admire cultures of perfection. Both The Economist and Apple have extraordinarily high standards for themselves. And really, in both cases, the work has a kind of primacy. Who designed the latest iPhone or who wrote that article? In both cases, that's anonymous because there's such a belief that the works speaks for itself. I have a lot of admiration for that.

And then cultures that have longevity and have really sustained success. Our biggest investor is Sequoia Capital. Sequoia has been the top firm or in the top three firms—it's a subjective ranking, but call it unquestionably a top-three firm—for essentially its entire existence. And there is basically no other VC firm that has been a top-three firm for multiple decades.

“It’s very striking to me how Warren and Charlie at Berkshire and how the folks at Koch industries are so into a kind of epistemology and structuring of doubt and accounting for biases and mechanisms for a clarity of thinking to a very striking degree.”

And I think the obvious question is, why is that? What’s different about Sequoia? There have been tons of VC firms, and a lot of different firms have had at any moment in time a strong claim to being a top-three firm, but what are the underlying institutional characteristics that enable that to be sustained? Of course, this applies to some of the other organizations we’ve mentioned, like, say, The Economist or The New Yorker. Or even—this is one that I’ve wanted to learn more about of late—Koch Industries.

Charles and David are most famous for their political activities, but if you just look at the company that has compounded from \$20 million in annual revenue to now, according to public estimates, \$100 billion over ... five decades, and there aren’t that many organizations that have compounded like that for that long without there being one driver of success. There’s no one thing that enabled that rise; they didn’t ... stumble upon some resource that they cornered. There was no iPhone for them; it’s truly something deeper and more institutional. And the fact that that’s been sustained for so long I think is interesting in its own right.

What is it that Sequoia Capital, Koch Industries, and The New Yorker share? I haven’t quite unpacked the answer to that yet.

Can you give us an example of what you’ve learned from studying Koch Industries?

It’s very striking to me how Warren and Charlie at Berkshire and how the folks at Koch Industries are so into a kind of epistemology, and structuring of doubt, and accounting for biases, and mechanisms for a clarity of thinking, to a very striking degree. Obviously, if you read the public writings, or if you go to Omaha and you listen to what Warren and especially Charlie talk about, it’s sort of half investing and half applied epistemology, half philosophy. And that’s been the case as well to a remarkable degree with Koch. And I don’t know them well enough by any means to opine in a deep sense. I’ve never been to one of their factories, I’ve never looked at one of their financial statements, and so I’m not qualified to assess in any kind of comprehensive way, but just in terms of what it seems that the leadership prioritizes, it’s interestingly consistent across two of the most successful multi-decadal institutions in the U.S.

There is something to be said, going back to your point earlier about learning from companies that have consistently demonstrated over a period of time, whether there are huge ... one-off hits that have caused most of that track record.

Right.

You’re a huge reader. Where did this love of books get started?

Well, we had crappy internet when I was growing up because our house was sort of alone. There was so much noise on the phone line, and we didn’t have

“We had crappy internet when I was growing up.”

internet for years, and then we got it; it was trickle slow and so on. I was fortunate; my parents were very willing to pursue all these hare-brained schemes, and so we eventually got an ISDN line, which was ferociously expensive, but that was the fiber of its day. At least as far as I was concerned, 7.6K a second was majestic. I could barely keep up with the speed.

And then we eventually got a satellite internet connection, which was really a game changer, but effectively meant that for the first 14-ish, 15 years in my life, there was no internet. We lived in a very rural part of Ireland. I was quite distant from even my friends at school. And so all there really was for us to do was play in the garden, which we did a lot of, and to read. It's funny; I often wonder about this in the context of “If I had kids or when I have kids, what's the optimal upbringing for them?” And, of course, you think, well, you kind of want them to grow up in a stimulating environment and have all these experiences and extracurriculars and everything else, but to me, that was not my upbringing. My upbringing was a kind of...

Get out of the house; go play?

That. Now, there was still plenty of stimulation around us. Our parents had lots of books, and so we could burrow our way sequentially through the shelves. It was pretty unfettered. And I think our parents had a... they followed our interests and supported them, but they didn't choose them. It felt like they pushed from behind rather than pulling in front. I think that's where the reading thing came from. I don't know. I run quite a bit, and I don't even run because I enjoy it that much. I enjoy it, but it's nothing kind of in the immediate moment; it's not like it's euphoric or anything close to that... it's pretty painful. There's the Greg LeMond quote about how—it's very dispiriting when you think about it, and it is very deeply true—how “it never gets easier; you just go faster”. And that's true of running. Like if I stay running for the rest of my life, it will never get easier; I will just—maybe—go faster. But it feels like something I ought to do, I vastly rather [prefer] having run than not having run. And so I continue to do it.

With reading, I don't feel like I'm weird; I feel like everyone else is weird, in that there's just ... so much stuff to know, and I guess I just feel stressed out by... like, it feels important, it's obviously important, and I don't know it. And so, shit, I better get to work. When I'm reading, I'm not in this ... especially blissful place. I enjoy it perfectly fine, but I think there are extremely important things that I really should know and I don't, and that feels problematic.

How do you filter what you read? There are millions of books. There's one of you, right?

Well, I discard a lot of books. I like the insight that there's a set of great books that are really worth reading, and there's a subset of those books that are really enjoyable to read. Maybe it's like 10 or 20 percent of them, say. And the subset, the intersection of really worth reading and really enjoyable to read, is

“I maybe start half the books I get, and I probably finish a third of the books I start.”

actually still more books than you can read in a lifetime. And so I've decided, “Well, I'll read all the books that are really worth reading and really enjoyable to read, and when I run out of those, then I'll go back to the books that are merely worth reading.” You know very quickly... you can decide if this is an enjoyable book to read or not. Or, “should I discard it?”

I think reading should be treated as a more active process. You should skim, you should skip, you should backtrack, you should discard and potentially return. You're not subject to the book; you're not a passive consumer. The book is there for you, you bought it, it's yours. And jump back and forwards, tear it in half if you want, annotate it wildly, use it.

I wholeheartedly agree.

I maybe start half the books I get, and I probably finish a third of the books I start. And that works out to finishing one to two books a week, but if I finish it, it's probably been recommended by somebody in the first place, and then it looked interesting enough, upon some very superficial skimming, to start. And then if I finish it, it was quite interesting. I actually like a lot of selection that happens along the way.

And then I think just the other thing worth pointing out is the line from Basho, the Japanese poet, that you shouldn't follow the people you most admire but you should follow what they admired. And I try to do that. I try to figure out, for the people who seem to be doing really great work or to have really interesting ideas, or just who I admire in whatever regard, I try to get, how do they get to who and what they are? What influenced them or what's upstream? And often it's quite obscure, but I try to disentangle that.

When do you typically read?

In the morning, in the evening, while walking. While walking is a good one, actually. Your peripheral vision is such that you can actually quite function to read a book while walking. And there're other people who have tried to do this and do it much more and faster than I do. You spend a lot of time walking and so being able to do that, I found to be quite valuable. Often while I'm eating.

You're sitting at home on your couch, it's after dinner, and you pick up a book for the first time; walk me through how you process that book, what you look at.

Normally, I'll jump sort of midway through it and I just start reading and ask “Would I like to have ended up here?” And, almost certainly, there'll be a bunch of the terms I won't recognize or the antecedent ideas I won't be familiar with or whatever, but, yeah, “do I want to be here or have gotten here?” And if for a couple of pages, it seems like the answer is yes, then I might backtrack to the start, pursuing it a bit more seriously.

“[My brother has this insight, that] at every moment you should be reading the best book you know of in the world.”

John has this insight that, and it's kind of related to the previous point, that at every moment you should be reading the best book you know of in the world. I don't mean ... the absolute best for everyone, but the best book for you. As soon as you discover something that seems more interesting or more important or whatever, you should absolutely discard your current book in favor of that, because any other algorithm necessarily results in you reading quote-unquote “worse stuff over time.”

Suboptimal?

Yeah. Exactly. And so I'll be reading the book on the couch and then maybe after 50 pages, I'll be in my room and I'll stumble across something else, and I might just switch rails. The other thing that I think is actually quite valuable is just leaving books out. When somebody recommends a book, I'll very often pick up a copy, ideally, a used hardback copy; the hardback books—they're more durable. And now, with Amazon, used hardbacks are really cheap. I'll leave it out. So there are books in the kitchen, there are books in my bedroom, there are books on my bed, just strewn everywhere.

And surprisingly commonly, either someone else will recommend the book or some aspect of the book ..., and it's still salient, it's still around you, and you're like, “Oh, yeah. I really should check out that thing.” Or something else triggers its relevance. You read an article; you just start appreciating a point or a question or something. Part of the reason that I still really value physical books is because ... for now at least, we still exist in physical space and it creates a sort of idea space for you that makes productive collisions more likely to happen.

What types of things do you typically mark up in a book and what does that look like?

I tend to just make notes in the margin. I tend to underline stuff, but in the margin, and I underline—I'm missing the term—I annotate it, mark it, highlight it in the margin because then you can flip through the book and jump quickly to the parts you marked. And then the other thing is, on the last pages, kind of the inside cover at the end, I tend to very quickly note page numbers for particularly interesting points, or things that jumped out, or whatever, so that I can easily go back to a book and I have a list of the 30 things that I found most interesting.

So you keep the book, a book that you completely read, that you like?

Yeah.

How often do you come back to that book?

If I want to make a particular point, or be reminded of a particular aspect, or whatever, maybe I will, but generally speaking, I don't. And I think part of the value of making the annotations is of course to imprint the points more firmly in your mind so that you don't need to come back as much, in some

“I think part of the value of making the annotations is of course to imprint them more firmly in your mind.”

sense. If it's really good -- I don't often do this -- but if it's really good, I might write a review for friends and just maybe share an email or a Google doc or something, or just share snippets to friends. That's valuable both because again, the act of summary or summarization aids as a kind of synthesis and better recollection, but also, of course, it triggers pointers and further suggestions from those friends.

And so, if you want to identify candidates in adjacent or ... If I don't perform the clustering and figure out what adjacent candidates might be interesting for further exploration, writing a review is a good place to start.

What books have you written reviews on for friends this year?

One that I really enjoyed was [A Culture of Growth](#) by Joel Mokyr. It's basically a book about, why did the enlightenment and the industrial revolution, ... really the industrial revolution, start when it did and where it did? He makes a specific case. And there're obviously, tons of different arguments that have been made for this. And because it only happened once, we can never know definitively --was it the abundance of coal in the UK? Was it something about the intellectual property system and patents? Was it the high cost of labor in the UK that made productivity-enhancing improvements more valuable? Was it something about trade? And so on and so forth.

And Mokyr basically makes the argument that it was primarily intellectual and more than quote-unquote “economic.” And secondly, that it was specifically a synthesis of the importance placed on scientific knowledge, where we realize that scientific progress and knowledge about the world exists and can be important, and that progress is possible and that we're not just imperfect imitators or receivers of the knowledge of the ancients. And so kind of a belief in scientific progress, and coupled with a belief in the practical importance of engineering and of the more prosaic aspects of industry and of practical pursuits.

Mokyr offers the example of Bacon, who both inspired the Royal Society—it was kind of one of his followers who created it—but also intended to catalog the practical knowledge of all of the craftspeople in the UK and the implicit functional knowledge they had. It's an interesting combination of the really high-minded and the very practical. Mokyr kind of teases through all these arguments and the republic of letters in this nascent rise of science on the continent and so forth, but all in service of this question of, why the industrial revolution then and there, and he talks about versions of it in China and so forth.

Anyway, I think it's a very important question and Mokyr's discussion of it is, I thought, particularly interesting. And so, yeah, I summarized it with my friends.

Awesome. Which book or books would you say have most influenced you?

I asked this question on Twitter ... a couple of weeks ago, and some of the responses I got were really interesting, and a lot of people responded, many

“The book is just the artifact upon which you project the characteristic that had already arisen or the belief that had already arisen.”

more than expected to. I didn't actually—embarrassingly, I feel guilty about this—I didn't post the response myself, and I thought about it and it's actually a very hard question to answer. Like I actually worry that it may not have been a good question because it's hard to know, did the book influence you or did you have an inkling or leaning? And then you read something that really resonated, but it's actually not... The book is just the artifact upon which you project the characteristic that had already arisen or the belief that had already arisen. And the book is not actually causal in and of itself.

Now, maybe it's so interesting to talk about the book as a symbol for the belief, but there is that kind of question. And then also, what I've often found is, I think the books that perhaps did in fact influence me the most in a causal sense are often not necessarily that good, and that maybe I'll read a book that triggers a realization, or some idea, or something. And the bolt kind of jolts me in some direction and then I'll go read better things about that question. And so it probably would've been better if I had just started with the better stuff, but in some truthful descriptive sense... It was like the worse one that actually influenced me. Maybe a better version of the question is like, “Which books do you wish you'd read sooner or something?”

Let's answer that question.

Yeah, I know. Hoist by my own petard. It's also about clusters of books. In that when I think about programming, for example... it would be hard for me to answer this question and not cite any programming books. Programming has been so influential in my mind and in my life, but I can't really point to any single programming book. I can name 10 that I think in aggregate work together, like [Paradigms of AI Programming](#) by [Peter] Norvig, and [Structure and Interpretation of Computer Programs](#), and [K and R C](#) [The C Programming Language by Brian Kernighan and Dennis Ritchie], books about operating systems, the [Tanenbaum book](#) [Modern Operating Systems], etc.

And in the aggregate, those hugely shaped me, but I don't think I can single out just one. Even two books about PHP, which were written by a guy who now works at Stripe—one of those books is the book that taught me to program, and so answering this question I can hardly not cite those, but it's kind of really the cluster. And you can give a similar cluster about science, or about economics, or sociology, or whatever, and so maybe I'll have to just get back to the better version of the question.

Switching gears a little bit, what's the smallest habit that you have that makes the biggest difference?

I reach out to people whose work I admire and tell them that, and often it leads to a dialogue. And in some cases, I've gotten to know them pretty well. I'm fortunate that Tyler Cowen, whom I mentioned, is a friend, but I was never introduced to him; I just randomly emailed him years ago, actually invited him to a Bitcoin Meetup that I held in 2011. And I did not, however, buy any bitcoin, but I invited him to that Meetup and he replied and apologized [that] he couldn't make it, but we ended up in a dialogue after that.

“I feel like it’s too important a question to answer simplistically, or too complicated a question to answer briefly.”

When you reach out to other people, half the time they don’t respond, but half the time they do. It’s asymmetric; it doesn’t really cost too much when they don’t, and it can be incredibly rewarding when they do. If I did not do that, I would have missed out on a huge amount.

How would you answer a question about what your personal values are?

It feels like too important a question. It’s kind of the book question. I feel like it’s too important a question to answer simplistically, or too complicated a question to answer briefly. And thereby, perhaps unsuited to something extemporaneous and I’m sure whatever answer I gave, when I’m thinking about it in an hour’s time, I’ll kick myself and realize I had left out this critically important dimension to it. I can cite some things I value, but the sense of giving a complete answer is very oppressive.

This is of course the value of Twitter, where because of the constraint there isn’t the same... because the system chooses when to cut you off rather than you choosing when to stop, that’s quite liberating. Maybe if you allowed me 20-second speaker values, I could do that, I could blame the constraint on anything I omitted.

We have 10 hours of recording left. What would you say is the most common mistake that you see people make over and over again that you wish you could correct? And you have 140 characters.

Maybe not having the right peer group or not having the right—mentors isn’t quite the right term, because mentor implies something quite active—but not striving to be more like the “right people”. Or not just being, in either case, deliberate enough about that. Of course, who the right peer group is for you is, that’s an entirely personal and subjective question, but whoever it is, it’s going to be massively formative and influential in determining where it is that you end up.

There’s a quote about how you end up the average of your five closest friends. I think there’s a very deep truth to that. But if you accept that, then of course who your five closest friends are, choosing that, and we do, though we may not think of it this way, we do choose those people. Like, you are choosing who you are. And of course that’s a bidirectional process where who you want to be is determined by who you’re around, which determines who you want to be around, and so on.

I think certainly my mental model when I was 18 is that “my five closest friends are people I ran into who liked me and where I liked them, and we’re cordial and close and all those things, but that it’s fundamentally mediated by happenstance.” I think people should invest more in the answer than they do. Once you’ve found the right people, you should really invest in the relationships. Because, if you accept [that] they can shape you -- and you think they’re the right people to shape you -- then, embrace that shaping.

“Who you want to be is determined by who you’re around, which determines who you want to be around, and so on.”

And then on the mentor point, on the latter one, I think almost all of us, at least subconsciously, have a set of people we hold in really high regard, or would like to be more like in at least some ways, and so on. I see people; in my opinion, they haven’t either found the right people or just like the right relationships and so on, and if they had someone who was steering them more or in better ways, [they] could just be much better off.

I want to talk a little bit about the future of e-commerce and maybe Silicon Valley culture. I know we’ve got to end soon, but talk to me about how payments... you foresee them changing from not only the customer perspective but from the merchant perspective over the next, you know...

Well, I think there’s two levels that maybe ... in that there’s all those ... basic mechanical software payments, where we forget just how much friction still exists and how many business models and transactions and businesses and everything are impeded for fundamentally stupid reasons. For example, because micro-transactions aren’t possible, both because the fixed costs are too high and because the friction is too high, then things that one would pay for with microtransactions just don’t exist. Not that they pursue a different monetization model, though in some cases they might. But as a general matter, a significant number of them just don’t exist. Or because maybe it’s hard to purchase things that are really expensive in a way where the risk of fraud is sufficiently low, then you don’t pay your rent online, say. And then I think maybe the most important dimension to that is the geographic balkanization and inefficiency that ensues where it’s extraordinarily difficult for somebody in Brazil to buy from somebody in Germany, or somebody in Germany to buy from somebody in India.

And so you get these ... unnatural sub-clusters existing, not because of deep, necessary limitations, but because of something much more arbitrary and contingent. Economists talk about the “gravity equation”, and the fact that the proclivity of any two countries to trade falls off with the square of their distance, and there’s always a big question like, well, is that about something fundamental in culture, or about surprising returns to proximity, or what have you.

Assuredly, there’s some of that stuff, but I think talking about the challenges, complexities, and hidden costs of paying methods, that doesn’t feel like a very deep thing. It doesn’t feel like something that is significant enough, on some level, to have such far-reaching and deep consequences. But I think a lot of these ostensibly “cosmic phenomena” are actually consequences of these very prosaic and straightforward limitations. And so, I really think that solving this aspect of commerce on the internet, like literally just making it easy for any two parties—a business and the consumer, in arbitrarily chosen countries—making it easy for those two entities to transact will have enormous consequence for the world.

That sounds like such a straightforward idea that it almost sounds clichéd. But the fact that it sounds clichéd should not blind us to the fact that it is still extraordinarily far from being the case today. We have had commerce on the internet for decades at this point, but it’s still... like, 90+ percent of

Brazilian credit cards do not work online outside of Brazil. Brazil is not some backwater! It's not some inconsequential country. Obviously, it's one of the top economies in the whole world, and yet Brazilian consumers basically cannot purchase outside of Brazil.

And so, it's difficult to overstate the magnitude of the limitations and inefficiencies that prevail today. So, that's kind of the current payments level. And then on top of that, I think there is—or beneath it, depending how you look at it—there's maybe just like a deeper question of what determines how many firms there are in the world. And what determines the character of those firms? Are they doing something innovative or novel, or are they doing something prosaic that has existed for a long time? What determines who starts and why and the probability of survival? What determines the growth trajectory and the expansion rate into other markets and other products and so on? And I think part of the Stripe hypothesis is that things like that, that ... one would think are very difficult to move are actually movable. And really macro measures, like the number of people who start a company or who start a technology company or again, the success rate of those companies...

Just to give some illustrative maybe intuition pumps here, when we survey companies started with Atlas, 60 percent of them tell us they would not exist if not for Atlas. I know they could be wrong. Maybe some of them actually secretly would, but maybe some of them are overstating their own resourcefulness or overstating... maybe they're underestimating the challenges they would have faced. I think that number could either be too high or could be too low. But let's be conservative and say that it's actually 40 percent. If Atlas is causing 40 percent of those founders to start companies where they otherwise would not have, and if the subsequent success rates look similar, that's a huge deal, especially if Atlas itself gets big. Over time, that can have real economic significance.

Or, if we can make it the case that businesses sell to twice as many global markets as they would otherwise sell to. Again, integrated over an entire portfolio, that's a really big deal. Nick Bloom at Stanford did this really interesting work, has done a whole bunch of interesting work about management practices. Do management practices matter? Is good management merely correlated or in fact causal in terms of leading to the advent of better outcomes? And they did an RCT, a proper trial in India where they taught better management practices to a quarter of firms and did not to a controlled group and saw double-digit percentages in revenue over a multi-year period. I don't recall exactly; I think it was 13 percent over three years or something like that.

That's incredible low-hanging fruit. All they did is teach better management practices: 13 percent more revenue, 13 percent more value provided by the company as assessed by their customers, just from better management practices. When we think about Stripe and what to do in the future and the possibilities that exist, and so on, it's much more I think about how do we perturb this overall system to move some of these macro outcome measures, like number of technology firms started, survival rate of these companies, expansion rate of these companies, magnitude of the value provided to the

“People who are driving fishing stocks to extinction care more about their ongoing right to do so than the rest of society cares about sustainable ecosystems.”

end users, consumers/customers and so on, and mediated by payments as this foundational layer because it's something every business necessarily has and because it gives us good understanding of the dynamics within the business and so on.

But it's, on some fundamental level, not about the payment, even though we think that, per the first point, the impact of just solving the payments will itself be enormous.

Do you think reducing friction across the board is a good thing? Or do you think friction in certain parts of it actually serves the system?

Well, serves it for who?

That's a good question.

Oh yeah, sure. Look, I think across society, I think so many of the things that looked like bugs are actually features from the perspective of somebody, of some constituency. And of course so much of politics is reconciliation of the countervailing interests of different constituencies. Of course the problem is that in so many cases, the incremental gain of the constituency is substantially outweighed by the social utility loss of the rest of society. And so, bad teachers do great in the U.S., but almost certainly that's a net bad trade for society, but the bad teachers care more about their ongoing employment than the rest of society cares, evidently, about correcting that.

The same thing applies to fishing policy, where...

The perspective makes all the difference.

Well, people who are driving fishing stocks to extinction care more about their ongoing right to do so than the rest of society cares about sustainable ecosystems. I think that's just the character of political economy. And so yeah, absolutely, I think, to return to our earlier example, it's not even clear that... One could look at the fact that essentially no new banking charters are being issued in the U.S. as a bug, or of course depending on your perspective, it's a wonderful feature. It's great for the regulators and it's great for the banks. Profits of consumer banks are higher than they've ever been.

Until they all get wiped out in the next crisis, right?

And then because they're even more systemically important than they were in the past, to the extent there was a systemic argument for bailing them out in '08, there'll be an even stronger argument in the future.

It's almost like when we were talking about this earlier, but when you get big, you have more loss aversion and see your goal as not necessarily to get better from your customers' perspective; it could be to prevent competition, prevent new entrants that might be a more—without moral judgment on it—it might actually be a more effective business strategy than innovating for your...

For sure. Absolutely. No question. I think we're very dissonant on this point as a society, where on the one hand, we decry a lack of innovation; on the other hand, in our collective action we do so much to ensure that it doesn't occur. And so on the one hand, we decry the state of the medical industrial complex and the 18.5 percent of our GDP that is spent on health care costs, and the plateau or even decline in life expectancy and the declining rate of drug discovery and so on. And yet on the other hand, we, through regulatory structures, make it harder and harder to engage in drug discovery.

You can't even start a hospital unless you've got a certificate of need. But if you observe that, well, hey, medical care in San Francisco doesn't seem so great, and it seems extraordinarily expensive, even though it seems like a very thankless undertaking, I'm going to try to do better. Well, first you'd better get approval for that. You can't just enter the market, and I think that kind of... And I'm not [passing] kind of a normative judgment. I have my personal preferences, but I'm not casting normative judgment as to what we ought to do as a society. The thing that I feel strongly is that we're inconsistent in our stated desires.

There's like a perpetual seesaw, if you will, where success sows the seeds of its own destruction. How would you make an argument right now that San Francisco or Silicon Valley is doing that?

The obvious one... Well, the two obvious ones, I guess, are in culture and in housing and costs in general. On cost, on the latter, everything is getting more expensive, and nobody seems to quite understand exactly what's going on. If you take health care again, for example, the case has been made that this is not in fact a bad thing, that what would you expect an enlightened society that has solved all of its other material needs to spend its money on? Healthcare, it's the last thing, it's the last frontier. Perhaps we are actually getting commensurate improvements, if you disaggregate appropriately and you analyze the right way. Or perhaps not.

How much of this is some Baumol's cost disease, where, as some things are getting more efficient, and the higher productivity and higher wages are causing cost increases elsewhere to pay for opportunity costs and all the rest? But I think specifically in Silicon Valley, and specifically on the cost of living and housing, Silicon Valley is the greatest concentration of wealth creation that I think has ever existed in the U.S. on a per-square-mile basis, potentially that has existed ever in the world. Facebook, Google, Apple, Intel, they're all based in a fairly small number of square miles. If you were to look at Seattle and the Bay Area together and look at that aggregate urban zone, separated as they are by a two-and-a-half-hour flight, then of course you can layer in Amazon and Microsoft as well. And obviously what you see is that their rise in success was enabled in part by cheap mobility and cheap expansion. And again, through just political economy and our collective decision making, that no longer exists. Cheap mobility no longer exists, and cheap expansion... And you can see it now in this latest generation of upstarts, be it Twitter or Uber or Airbnb or Lyft or whatever, who are facing these really significant structural headwinds.

“So much of the wealth that’s being created, this improbable fountain of wealth creation, is accruing to the lottery winners of the existing landowners rather than to the people who are actually doing the work.”

So much of the wealth that’s being created, this improbable fountain of wealth creation, is accruing to the lottery winners of the existing landowners rather than to the people who are actually doing the work. And because of that accrual, the barrier to entry for newcomers is getting progressively higher, and you see it in declining rates of mobility. And furthermore, the other people in the city [who are] not in the tech industry, who might otherwise benefit from it, are of course getting priced out. This is not necessary; you can look at places like... Tokyo has over the last couple of decades been in an improbable ... not especially improbable, but has been such an enormous economic success story, and you had the boom and the bust and the supposed stagnation of Japan in the early '90s. But broadly speaking, Tokyo has done really well, but because of vastly fewer limitations on housing supply, [they] have had very stable housing costs, have not had the same displacement.

The issues we face and we see here in San Francisco, where it’s getting ever... a 40 percent rise since we got to San Francisco in 2010, that’s not necessary, it’s not natural, and it’s a function of our collective decisions rather than some secular and unavoidable economic force. I find it dispiriting because it’s a negative sum in the sense that it’s not just that these gains go to these existing landowners, but actually there will be fewer future gains.

I think you should be mad about this, even if you don’t live in Silicon Valley and you don’t have the slightest interest in doing so, because it’s much less likely that the next cool technology that you’d like to take advantage of will exist. It’s a suffocation of future potential and future gains. And there aren’t many places... Well, if you believe in increasing returns to scale—this is Paul Romer’s work and others—that because of the collision of ideas and people in cities, [this] makes them more productive than if they’re elsewhere. If you believe that to be the case, and there’s ... pretty good empirical data that it is, then you can’t just move elsewhere. You can’t just move to Nevada or wherever, the south. You actually will be less productive in those zones. And so again, I think it’s a real loss in terms of spillover... to the rest of society in the service of not building six-story buildings in San Francisco.

What do you think your role as a large employer and thoughtful citizen of San Francisco is in this?

Well, I don’t make any secret of the injustice, and the moral injustice in terms of the displacement that’s occurring and the economic wrongheadedness of the prevailing policies. I’m a landowner in San Francisco, John and I own a house together, and I hope its value declines, in that... I think it’s impossible to answer what the price of land should be, but I think it is very clear that on a marginal basis, the social returns of cheaper land in the most productive region of the country would vastly outweigh the reduction ... wealth cut to existing landowners.

But going back to the banks, everybody has a system that they want to protect.

Totally. Of course you can try to estimate the magnitude here. Over at Berkeley, this guy, Enrico Moretti, has estimated that 50 percent of U.S.

GDP growth between 1964 and I think 2010 was left on the table, as it were, by inefficient land use and land allocation. And obviously 50 percent was a high number and quite speculative and it's very difficult to measure the counterfactual, but even just the idea that one can with a straight face hypothesize that it could be anything remotely in that vicinity, I think gives you a sense for how high the stakes here are.

And yes, we can decide that we place such an enormous premium on the aesthetic appearance of the San Francisco of today, recognizing that it is approximately a third of the density of even just Greenwich Village in New York. The other extreme is not Hong Kong; you can cripple San Francisco and get to Greenwich. We can decide that that's our preference, but sober estimates are measuring the cost of that in double-digit percentage points of aggregate national GDP.

And of course when you look at our revealed preferences in terms of where we like to take vacations to or where we dream of spending a summer someday and things like that, it's to European cities which tend to be of very significantly higher density—Paris, London—much, much higher density than San Francisco. Again, I'm hesitant to cast normative judgment, but I personally feel strongly.

I think that's a great place to leave this. This has been a phenomenal conversation. Thank you so much for coming on the show. Where can people find more about you?

Well, if they want to start a business, they should head to Stripe.com, but if they want to subject themselves to more of the particular detritus that I post, they can head to my Twitter account, which is just [@PatrickC](#).

Thank you so much.

Thank you.

Normally, podcast transcripts are only available to members of [the Learning Community](#). However, we release one or two a year publicly. If you'd like access to a bunch of extras like podcast transcripts, the FS Knowledge Library, and deep dive discussions with like-minded people on every topic under the sun, [join the FS Learning Community here](#).

fs

FS.BLOG
@FARNAMSTREET