Position Paper - A Framework for Deciding How to Specialize v1

Philip Morgan | September 2018

Contents

A Framework for Deciding How to Specialize

A strong market position comes from a decision to specialize followed by a few years of disciplined followthrough. There are 5 kinds of market positions professional services firms can cultivate, and 5 corresponding ways of specializing.

But how to decide where to specialize in the first place? Some of you will land on a good specialized focus. Your intuition or naturally assertive decision making style will take you there—perhaps with some iteration along the way—without any outside help.

This article is for the rest of you. Those that benefit from process and structure or at least an idea of how to approach a high-stakes decision you've never made before.

I have a framework I use when I advise clients on this process. I'll share that framework with you here. It's simple enough to tackle on your own, but if you want expert, empathetic guidance applying it you might consider one of my group workshops or a private coaching engagement.

First you'll need to understand the 3 ways in which you could decide, the missionary-mercenary spectrum, and risk profile. I elaborate on this in my books (The Positioning Manual for

Technical Firms and Specializing Without Failure), but this article should help you understand this framework enough to use it for your own business.

The 3 Decision Approaches

When deciding how to specialize, you are looking for an advantage of some kind. These advantages break down into 3 categories:

1. A head start of some kind
2. An interest or connection of some kind
3. An entrepreneurial theory

A Head Start

You might specialize in a way that corresponds with a head start in your business.

A generalist developer that has worked with 15 clients—5 of which are clustered in finance—might decide to specialize vertically in the finance industry because their experience working with finance clients represents a head start. They understand the industry jargon better than most generalists would. They understand how software projects in that vertical tend to be commissioned, funded, bought, and managed. They may understand exactly how business problems tend to be solved with custom software in the finance industry. All of this experience moves our generalist developer closer to a domain expert who can assertively lead the sales conversation and, perhaps, the resulting project. They might not be all the way there to that domain expert leadership position, but they have a head start.

A developer that has a deep understanding of working with regulated data on cloud platforms might specialize horizontally in that domain because it has value to lots of businesses, and few other developers find it worth learning about. Like our generalist developer who decides to specialize in finance, they have an expertise head start that may allow them to lead a sales conversation and add value to the project that few other developers can.

These are just two of a multitude of possible examples of building on a head start.

An Interest or Connection

Instead of identifying a head start, you might identify an interest or feeling of connection to a vertical or problem domain and specialize there.

A Short Sidebar:

I recommend you investigate all three decision approaches and choose the one that is the best fit for you. So ultimately you might pursue an interest or connection instead of a head start, but you would most likely do so after identifying and investigating all potential head starts, interests, and entrepreneurial theories.

An interest is exactly what you'd think; you find a market vertical or problem domain *interesting*.

A Longer Sidebar:

Most developers would define their primary interest as building software. I won't adequately support this argument here because it's tangential to the point of this article, but coders that have made the transition to consultant combine their interest in software with an interest in a market vertical or business problem.

You *can* specialize in a technology platform (defined here as: a language, framework, or platform). In fact, most developers default to platform specialization. It's the place their interest most easily leads them to.

Like the 4 other ways of specializing, platform specialization has its pros and cons. Most of my clients hire me to help them when the cons of platform specialization have become apparent—and painful—to them.

The primary drawback of platform specialization is fairly rapid commoditization of the skillset. For example, picture someone who became interested in iOS app development in 2010. The timing of their interest would have led them into that platform specialization at an ideal time. Demand was strong and the supply of iOS dev talent was relatively low, resulting in premium rates and lots of client-side urgency in the sales cycle. "We need an app like *yesterday*!"

Developers—heck, anybody—who faces this situation would have no compelling incentive to figure out how to do lead generation well or develop a strongly differentiated value proposition. You're busy makin' apps and billin' hours!

It's just too easy to enjoy a medium-term bubble that works in your favor. I view this as a significant drawback of platform specialization because a talent market like our iOS example here tends to quickly attract a greater supply of lower-cost talent, which commoditizes the market. Developers who could previously command premium rates for no reason other than having rare, valuable skills start facing rate pushback and longer, more difficult sales cycles.

Of course, our iOS developer can move on to the next hot platform before the air completely leaks out of the iOS bubble. In fact, there will almost certainly *always* be demand for iOS talent. I'm not saying that demand ever goes away, just that the bubble-like quality of the market for iOS skills cools off, and it becomes a commoditized market for talent, and in a commoditized market you don't have buyers dying for anybody with a pulse and an ability to write iOS apps.

Again, each of the 5 ways of specialization has pros and cons, but the eventual commoditization of most platform specializations is one good reason to consider specializing in a market vertical or problem domain, which means exploring other forms of interest and connection you may have. What verticals do you find interesting? Where do you feel some sort of connection to a type of client? What evergreen business problem do you have an interest in solving? Again, this interest in something outside the world of software development is a key prerequisite to making the coder —> consultant transition.

For example, imagine that you—like a client of mine—find the maritime shipping business oddly fascinating. This interest could pull you towards a unique specialization, one where you learn to build software specific to the maritime shipping industry.

Another example: Cory Quinn's interest in the financial and business implications of cloud billing—a business problem—led him to cultivate a horizontal specialization in just this area (for more on Cory's story: http://consultingpipelinepodcast.com/119)

As an example of *connection*, I think of Evan McBroom, whose connection to Christianity via his personal faith led him to focus his marketing and communications agency on churches. Evan elaborates on that here: http://consultingpipelinepodcast.com/100

A strong interest or feeling of connection can be a very powerful ally in your journey towards a desirable market position, and can compensate for lacking a head start. Naturally, there is a tradeoff. If you have no head start or you have a natural head start you choose not to pursue, your interest or connection can compensate, but at the cost of the effort and time required to build some kind of advantage in terms of expertise or something else of value to your clients.

Entrepreneurial Theory

Instead of identifying a head start or looking for interest/connection, you might pursue an entrepreneurial theory.

Matt Rogish of ReactiveOps was in essence pursing an entrepreneurial theory when he and his team scrapped a significant amount of custom tooling to move their "DevOps as a Service" offering to Kubernetes (more on Matt's story: http://consultingpipelinepodcast.com/120). They had no particular head start with Kubernetes because at the time ReactiveOps embraced it, it was early adopter-phase technology. Additionally, this specialization decision wasn't based on any particular interest/connection. Instead, it was a measured bet on technology that hadn't yet achieved mainstream status. The purpose of the bet was to achieve efficiencies that made ReactiveOps' core service offering better.

When you pursue an entrepreneurial theory, you are seeking advantage less from your personal characteristics or background and more from the strength of your business model or marketing.

Returning to the example of ReactiveOps, when Matt Rogish talks about the decision to forgo the sunk cost of years of custom tooling and go all in on Kubernetes, I think you'll hear in his voice a calculated but limited *personal interest* in Kubernetes itself. What motivated him was a vision for how Kubernets could make his *business model* stronger in the medium to long term. In fact, in that interview (again: http://consultingpipelinepodcast.com/120) you'll hear Matt talk about how in the short term, the switch to Kubernetes was costly because in addition to the retooling cost, there was a customer education cost he had to bear. This cost was worth it in light of the benefits to the business model, and the limited *personal* interest Kubernetes might have held for Matt was worth it in light of the business upside.

The Missionary-Mercenary Spectrum

I'm not sure there's any scientific basis for this, but it's useful to think about your motivation for being self-employed along a spectrum. On the one end are the "missionaries". I'm one of those. We're seeking meaning our lives, and our businesses need to reflect this search for meaning.

I'll let HBR take over for me in defining the missionary-mercenary spectrum:

> What's the difference between mercenaries and missionaries? As Doerr explained to an audience at Stanford Business School, mercenaries are "opportunistic." They're "all about the pitch and the deal" and are eager to sprint for short-term payoffs. Missionaries, on the other hand, are "strategic." They're all about "the big idea" and partnerships that last, and they understand that "this business of innovation is something that takes a long time" — it's a marathon, not a sprint. Mercenaries have "a lust for making money,"

> while missionaries have “a lust for making meaning.” Mercenaries obsess about the competition and fret over “financial statements,” while missionaries obsess about customers and fret over “values statements.” Mercenaries display an attitude of entitlement and revel in the “aristocracy of the founders,” while missionaries exude an attitude of contribution and welcome good ideas wherever they originate. Mercenaries strive for success; missionaries aspire to “success and significance.”

Source: https://hbr.org/2016/04/what-separates-high-impact-entrepreneurs-from-those-who-dont-make-a-big-difference

In my work with self-employed software developers, I see few people on the extreme mercenary end of this spectrum. I suspect the dramatic potential upside of digital products is more attractive to them than the relatively tame upside of a healthy services business, or they pursue an entrepreneurial theory and build a team like Matt Rogish.

I do see *lots* of “missionaries”, who you might overhear saying things like the following:

- "I love helping software teams improve their process and communication."
- "After I encountered Elixir, I never wanted to use any other language because of how effortless Elixir made stuff that's a PITA in other languages."
- "When I find out that a big commerce operation does crazy shit like deploying straight to production without a proper release process, I want to shake them and tell them how important—and simple—a good release process is!"

This is the language of the missionary. Software developer missionaries are less likely to specialize based on an entrepreneurial theory and most likely to choose based on a head start or interest/connection. Please remember that these are general patterns, not rules. What I'm presenting here is a framework to help you think through what's ideal for you, not an employee policy manual for what you can and cannot do while working for yourself as a specialist.

There's one more important factor that effects how you might choose to specialize, and that is your risk profile.

Risk Profile

Your decision to specialize should excite you rather than totally stress you out. Specializing makes things better, but it does not make all aspects of your business easier. You still have to deliver results for clients, sell with confidence, and find ways to connect and build trust with prospects. The energy and positive outlook required for these business functions would be drained away if you specialized in a highly risky way that made you uncomfortable and stressed out.

There is almost certainly a difference between how much risk you can actually handle and how much risk you can comfortably handle. In other words, most people can actually handle more risk than they are comfortable with. But how much risk you can *comfortably handle* is a threshold you should respect as you consider how to specialize.

How do you determine where this threshold is?

First, un-brainwash yourself. If you consume much press coverage of startups, you'll have been exposed to over-simplified tropes where risk-seeking founders bet it all on a moonshot and reap their reward when, against all odds, they “succeed”, meaning they get to take on a ton of debt and give up control of their company to investors. This is a distorted view of entrepreneurial risk-taking.

In most services businesses, the coupling between risk-taking and the payoff for that risk-taking is much closer. A courageous but risky decision can pay off in weeks or months.

Extreme levels of risk taking are not required to run a successful services business. To be sure, *some* level of risk taking is inherent in working for yourself, and some additional level of risk taking can pay off in the context of a services business. But in the world of services, reducing risk for your clients can create a ton of value. A somewhat risk-averse services business owner can leverage their own fluency with avoiding risk in a way that benefits their clients.

So before you assess your own relationship to risk, make sure you’re not looking at yourself with a distorted mirror.

Second, think of risk as having two facets. Your emotional response to risk is your risk tolerance, and your actual financial ability to sustain a loss is your risk capacity. You may have a high risk tolerance and low risk capacity, the exact inverse, or some other variation of those two factors. Risk tolerance and risk capacity both increase or decrease your ability to handle risk, and I combine them together using a simple proprietary algorithm into what I call your risk profile.

Third and finally, realistically asses your past behavior with regard to risk (ex: “What degree of risk have you taken with your financial decisions in the past?”). This is a better measure of your risk tolerance than hypothetical questions, although hypothetical questions derived from prospect theory are useful (ex: “In addition to whatever you own, you have been given $1,000. You are now asked to choose between a sure gain of $500 or a 50% chance to gain $1,000 and a 50% chance to gain nothing. Which do you choose?”).

After you’ve done your level best to realistically understand your risk profile, set a risk threshold. I have a way of doing this with my clients and workshop participants, but you could do a lot worse than assigning yourself a score on a simple 3-point scale. Later, when you apply this framework by identifying all your specialization options, you can score those options on the same 3-point scale and eliminate any that exceed your risk profile.

Intermission

These are the 3 core elements that I weave together when I help clients through the process of deciding how to specialize: finding your advantage, the missionary-mercenary spectrum, and your risk profile. The specialization decision becomes easier when you explore each of these 3 areas.

Now to be clear, the decision rarely becomes a “no-brainer”, or emotionally un-taxing. If you’ve been operating as a generalist, specializing always requires courage and then lots of discipline. If you happen to be *angry* or frustrated at how being a generalist has harmed your business, that emotional energy can fuel a swift specialization decision. I’ve seen this a number of times.

This article and the work I’ve done around the specialization decision may make it *easier*, but in general the decision is always at least a moderately difficult one. That’s not surprising given how important a decision it is.

The second part of this article is me turning this framework into a linear process you can follow. I hope this helps you apply the framework to your own specialization decision making.

Applying the Framework

To apply this framework, I recommend you go through the following process:

```
1. Inventory
2. Set a risk threshold
3. Narrow the inventory
4. Validate
5. Decide
```

This is the same process I take my clients through when they're deciding how to specialize. I'll offer a few details below on each step.

Inventory

My advice here is as old as the venerable "pros/cons" list: *write it down*. When facing a high stakes decision, our minds can play tricks on us. We can second-guess our previously clear thinking. That's why you should create a *written* inventory. Do it in a spreadsheet so you can easily sort things later.

You might ask: what goes on this inventory?

I'd suggest inventorying every project you've been involved in as an employee, freelancer, or consultant. You want to create an inventory that is 1) relevant and 2) high signal to noise ratio.

This inventory captures previous experiences that might be relevant to your new focus. Of course, you don't know what that new focus is, so your inventory should capture *everything* that might be relevant. However, if at some point you dramatically leveled up in your work, you can probably exclude that old pre-level-up experience from your inventory. If at some point you had one of those mid-career epiphanies that caused you to dramatically change direction, you cam probably exclude that pre-pivot experience.

The largest inventory I've ever seen captured about 150 projects. One this lengthy could take 6 or 8 hours to complete, but it's time well spent, so capture as much detail as possible. Each row in your inventory should include the following:

- Client name
- Vertical the client is in (ex: manufacturing)

- Brief project description (ex: new website, API for Foo, map UI)
- Impact of your work (ex: improved user morale, reduced time to enter claim by 300%, help client compete with Foo)
- Profitability (to you) of the work
- Interest: how much do you enjoy working with clients like this or solving problems like this? A 3-point scale is a good way to rate this.
- Access: how many contacts do you have in this client's vertical? Note that this does not have to be an absolute number. Instead, go with a 3-point scale representing: 1 = `not many` 2 = `a good amount` 3 = `a lot`.
- Credibility: what kind of credibility came out of the work? Again, a 3-point scale is a good way to rate this.

Extend Your Inventory

Next, you need to explore areas of interest where you have no previous experience, and you need to include any entrepreneurial theories.

Carve out an hour or two and make sure you're relaxed and uninterrupted. Review all the top and sub-level categories listed on https://www.naics.com/naics-drilldown-table/.

What seems interesting? What calls out to you in some overt or subtle way? Add them to your inventory with Credibility and Access scores of 1 if you're a complete outsider to the vertical.

Work in any entrepreneurial theories you have as best you can to your inventory. Naturally the risk score for these and verticals you've never worked in before will be quite high.

Set a Risk Threshold

Your completed inventory is a tool you can use to surface patterns. I know as you're reading this your reaction is probably, "Thanks, but I don't need a dumb spreadsheet to show me the patterns in my own project history". I hear you, but I've lost count of the number of workshop participants who have reported greater clarity and confidence about the head start(s) their past client work has set them up with. The inventory is an extremely simple tool, but a powerful one. Don't be too proud to try it.

About those patterns. If you sort by the Client Vertical column, for example, you'll easily see if there's a vertical where you have a large amount of experience. Or if you sort by the Interest column, you'll group projects where your interest is high. The same with Access or Credibility.

I have a proprietary algorithm I use for combining the numerical scores for Interest, Access, and Credibility into what I call a risk threshold. Actually, "algorithm" is a bit of a stretch. I'm just adding and subtracting according to the following simple heuristics:

- Higher credibility, access, and interest *reduce* the risk of pursing a given area of focus.
- Lower scores for credibility, access, and interest increase the risk of pursing an area of focus.

Here's a simplified version of the formula I use:

```
(A + I + C) / 3 = R
```

Set up a new column in your inventory to handle this math. As a result each project should have its own risk score. A project with an Access score of 1, an Interest score of 3, and a Credibility score of 2 would have a risk score of 2.

"Ah ha!", I hear you say. "I knew all along that this was a subjective, unscientific process! Those numbers are subjective, not objective absolutes!"

You're right, but trust me—this is a hell of a lot better than guessing or—as I've often jokingly suggested—throwing a dart at a printed out list of specialization options. This process really does deliver clarity if you give it a chance.

Onward!

Narrow the Inventory

I also have a proprietary risk profiling tool I use with clients. I combine the DiSC work personality assessment tool, questions about past behavior with money, and questions about financial decision-making derived from prospect theory. There is a partially subjective aspect to this tool, which is why I can't spell it out simply here the way I did with the risk threshold calculation above.

Nevertheless, you can at least roughly figure out your own risk profile. First two simple definitions.

Risk tolerance: your emotional response to risk.

Risk capacity: your financial ability to withstand a loss.

Risk tolerance and risk capacity are not necessarily correlated. You will find people who are very comfortable with significant amounts of risk (high risk tolerance) but who would be very negatively impacted by a bad month or two in their business (low risk capacity). The inverse can also be found in the wild.

Your risk profile is the combined effect of your risk tolerance and risk capacity on your decision-making and behavior.

To assess your own risk profile, do the following. It's a simplified version of what I use in my work with clients:

1. *Risk capacity*: If you landed no new clients starting today, for how many months could you meet your financial obligations without taking on debt? This is your runway. Your runway is the primary contributor to high risk capacity, though other things like flexibility

to quickly reduce expenses also contribute. Let's stick with the 3-point scale here:

   1. 0 to 3 months runway = 1
   2. 3 to 6 months runway = 2
   3. 6 to 12 months runway = 3

2. *Risk tolerance*: For our purposes here, let's call spending money with the intent of making more money—directly or indirectly—an investment. Remember as objectively as you can how you have invested money in the past. A willingness to invest is correlated with a tolerance to risk. Rate yourself on a 3-point scale:
   1. Have avoided investing = 1
   2. Have invested what, by your standards, are modest sums of money = 2
   3. Have invested what, by your standards, are large sums of money = 3
3. *Uncertainty*: To what extent does uncertainty create stress for you? This is a subjective measure, so do your best to assess your stress response to uncertainty. Use that 3-point scale again. None = 3, some = 2, a lot = 1.

At this point, you can sum these 3 scores or average them. I'd recommend averaging them to keep with the 3-point scoring scale.

At this point in the process, you have a risk profile score on a simple 3-point scale, though your score may be a fractional number. Let's call 1 on this scale a low risk profile, 2 medium, and 3 high. Brilliant, right? ;)

This is not fancy stuff, but it is a useful way to narrow down the list of options on your inventory to what, in some cases, will be a much more manageable number of options.

Naturally, you're going to eliminate the options that fall *above* the threshold created by your risk profile. If your risk profile is 2.3, then you'll want to eliminate the options for specialization that have a higher risk score than that.

Validate

Next, I'd suggest you undertake some very simple market sizing. You want to be sure any potential area of focus has a "goldilocks" size: not too small to support your business and not so big it makes you a small fish in a large pond.

David Baker has great data and conclusions in this area, so I'll just reference his published guidelines. You want a market with 10–200 competitors and 2,000–10,000 prospects (from https://www.recourses.com/how-man-competitors-and-prospects-should-you-have). Competitors would be specialized firms, and prospects would be companies big enough to afford you and small enough to take you seriously.

If you're pursuing a vertical focus, market sizing is relatively straightforward. If you're pursuing a horizontal focus, you have to ask: how do you know who your prospects are? Long term they're probably going to find you via really strong content marketing you'll produce, but you don't have that yet. For example, if you're Corey Quinn, who I mentioned above, and your market is companies that need their "horrifying AWS bill" fixed, how do you know how many companies that is? Not to be an ass, but you could hire me. Aside from that, you could use a few market sizing proxies, the easiest of which is the "conference hack".

Any horizontal market that has a national-level conference or association is likely to be a right-sized market. The conference/association won't necessarily map exactly to how you define your specialization, but if there's a rough fit you're probably OK.

For example, consider Corey Quinn again. How could he have known there are enough companies that need help with their AWS bill? Well, there's the AWS re:Invent conference. It's not a conference about AWS bills, but it's a conference about the platform, and it's a big one; somewhere over 40k attendees. This is a rough proxy for the size of the AWS market, especially when you compare the attendance of AWS Invent with the attendance of conferences for Google and Microsoft's PaaS products. With this data point, our hypothetical Corey knows he's not considering specializing in a platform that very few companies use.

You can exclude from your inventory anything that would point you to a market that is definitely too small, and you can refactor any focus that is too large by adding specificity. For example, the market vertical of manufacturing is almost sure to be too large, but pharmaceuticals is probably not.

If you have items on your inventory that represent an entrepreneurial theory, you may be in terrain that's a bit too complex for this article to help you navigate. Hopefully the main ideas help some anyway.

Decide

At this point, I've offered every reasonable way to narrow down your inventory. Since I can't know what you started with, I don't know how lengthy your "short list" is now. But if you're like most of the clients I've helped through this process, it's shorter in a way that clarifies your thinking and increases your confidence about your decision.

You haven't delegated your decision to chance, and you haven't made it using an algorithm. Instead, you've used a pragmatic process to generate and then narrow down a list of good options. In other words, you've taken what is usually an emotional and subjective process and added a layer of objectivity.

You may want to perform additional market research or validation before you actually change anything in your business. Depending on your personality, this will either be a form of procrastination or a good investment in the quality of the decision. If you've historically used research to procrastinate, that's probably what you'd be doing here. If that's you, consider the following next step instead of a lengthy research/procrastination process.

Reach out directly to some prospects. Ask them this question, and see where their response takes you:

"Who do you have helping you with [thing you might specialize in]?"

Application

If you have any questions about applying this guidance to your business, please feel free to contact me at philip@philipmorganconsulting.com.

If you'd like my direct help navigating the transition from generalist to specialist, my services may be a fit: https://philipmorganconsulting.com/services/