

How to simplify complex decisions by cleaving the facts

by Jason Cohen on November 25, 2016

Simplify complex decisions by separating upsides from downsides, investing in upsides, vetoing with downsides, and using an appropriate decision framework.



"It's not that I don't like you, Ted, you just don't fit my target demographic."

We face complex business decisions for which there is no one, correct answer, and yet a strong and occasionally permanent decision is required of us.

Sometimes the puzzle is created by fundamental uncertainty inherent in the decision. Examples: how a market will evolve, what competitors are investing in, whether new marketing campaigns will be successful, or how an important new hire will perform. There are specific strategies for dealing with inherent uncertainty.

But I've seen little startups, mid-sized companies like WP Engine, and large companies struggle with complexity in decisions even in conditions of relative certainty. If this sounds painfully familiar, the following technique might help.

Separate upside from downside, and clarify the outcome

The first thing is to separate the question of "how could things go right" (the upside) from the question of "how could things go wrong" (the downside). The symptom of

not following this advice is arguing in logical circles that inhibit decision-making. For example, take the decision "should we build feature X:"

1. If we add feature X, we will be unique in the market.
2. But we're behind in feature Y, which three competitors already have.
3. But if we build feature Y, those same competitors will then have built feature Z, and we'll still be behind.
4. But we're losing business today because we don't have feature Y.
5. But if we were unique in a different dimension, that would turn the conversation into feature X versus Y.
6. But if a competitor who already has Y ends up copying us on X, then we're back to being behind.
7. ...

All of those statements are true! That's why it's hard to reach a decision.

Instead, separate the upside and downside of "building feature X," and clarify the outcome of those two aspects:


Upside

Unique position in the market, leading to deals won through differentiation.

Downside

Won't build feature Y, leading to deals lost to a specific competitor.

When you look at it this way, the decision is clearer. If we do feature X, we earn more deals (through differentiation). Also we gain some (less measurable) brand benefit from that differentiation. If we do feature Y, we also earn more deals, but only in a subset of the sales calls, and there's no additional brand benefits. So we should probably do feature X.



"So, which 'whole' do we want to go with?"

A counter-argument could be: We're losing \$10m/yr not having feature Y; we anticipate gaining \$1m/yr if we have feature X. In this case, *quantification* of the upside/downside would also lead us to a clear winner of doing feature Y. Nevertheless, this example demonstrates that, even in the absence of data (which is frequently the case, especially with early-stage startups), this technique helps us arrive at a clear and sensible decision.

Decide with upside, veto with weakness

Let's overlay a deeper perspective on our upside/downside process, while also switching to a new example, to demonstrate the universal applicability of this framework.

Suppose we're hiring a high-level, experienced position, like "VP Marketing". The "perfect candidate" is mythical—a person who is world-class at vision, strategy, operations, people, org-structure, communication within the company, communication outside the company, gets things done quickly, mentorship, etc.. So, again it's easy to go in circles when deciding on a given candidate, encouraged by her merits but worried about her weaknesses, similar to the "feature" discussion.

With leaders, however, it's widely accepted that you should hire based on their exceptional strengths, and then fill in weaknesses using the composition of the rest of the team. This is because you first and foremost need a 10x player. This is especially true of Product Managers,

and even true of startup founders themselves, where it's your job to fill the company with people much better than you at every position.

So first you separate the upside ("in what aspects is she truly exceptional") from the downside ("in what aspects is she lacking?"). But now we apply a more sophisticated idea: That we must **primarily decide based on the upside** and then ask ourselves if we can **mitigate the downside**.

In this example, to be specific:

1. For the most critical business problems we have today, that this leader will be expected to solve, is she world-class in solving those particular problems? (*Because, if not, you'll still have your critical business problems, so that's a no-hire regardless of how much you like the upsides or don't care about the downsides.*)
2. For the weaknesses of this leader, do I know what those are (because if not, you can't plan for them), and are they either un-impactful, or do I understand exactly how I or they will mitigate them? (*Since no one is perfect, there will always be items in this category. If any are impossible for you to ameliorate, it's a no-hire. Otherwise, it's a plan!*)

By basing the primary decision on the upside, and using the downside only as a "veto" when unworkable, you have further clarified how to make the decision.


Invest in upside, ignore or dampen downside

In fact, we don't just want to "base the primary decision" on the upside—we want to intentionally over-invest in it.

To explain this, and to further emphasize the broad generality in which these rules apply, let's switch examples yet again. Now the decision is: "What aspects of our product should we invest in, over the next 12 months?" (As an exercise, afterwards try re-running this example with "my own strengths/weaknesses as a person," and see if you get some insights!)

Suppose we plot our strengths in six key areas, using some relative measure:


Current Strengths



People naturally focus on weaknesses. This appears in statements like “We’re getting killed in the market for not having F” or “I’m sick of customers complaining—rightly! —that we’re bad at E” or “20% of customers leave us because we’re so bad at D.”

As a result, we often invest time and money into mitigating weakness. Sure, you can’t turn a catastrophic weakness into a super-strength, but you could at least bring it to a passable neutral:

Typical: Neutralize Weakness



The trouble with this line of thinking is that it ignores a deeper truth, which is that developing a 10x strength is far more valuable than shifting a weakness to neutral. This is true personally, professionally, and in Product.

For example, in the v1 launch of the iPhone, it was far more important for it to be 10x on its strengths (e.g. form-factor and browser/email experience) than for it to

shore up weaknesses (e.g. lack of copy/paste, poor call experience).

Or for example, as an engineer you will be far more effective in producing quality code quickly if you’re extremely deep in one language/framework/problem-domain, than if you spread out your time becoming passable at ten different languages.

Or for example, Heroku won the hearts and minds of Rails developers because it was 10x better at the deploy/stage/production system, and therefore developers (begrudgingly!) put up with (what were *then* considered weaknesses) a read-only filesystem, having to use Bundler, having to use PostgreSQL, over-paying for CPU, and tying your fate to a platform.

Or as a counter-example, when we had poor web design at WP Engine (weakness), we invested and got zero return, because it turned out that our words and product/market fit were 100x more important than design.

Therefore, this is how you should be investing:

Better: Trading Weakness for 10x Strength




This drives home the idea of separating upsides from downsides, strengths from weaknesses, then deciding using and further investing in your greatest upsides and strengths, while using downside as constraint to design around, or possibly a veto in extreme cases.

Decide using ≤ 3 key dimensions

I pulled a trick earlier. It was necessary to make the decision clear. You might not have noticed, but you'll want to do it intentionally.

In the first example of whether to implement feature X or Y, I boiled it down to "both increase sales, but only one also increases differentiation." Sure, when things are simple, the choice is clear. But the choice wasn't so clear from the text that preceded the summary. And surely there are other considerations such as: how long would it take to implement, do we have the right team assembled, and what is the likelihood of abject failure. Why was it valid to boil it down to such a simple decision?

In one sense, this *can't be* valid. Any simplification that ignores a dozen important dimensions can't be an accurate model of the problem! However, if you don't oversimplify, you will never reach a firm and clear decision.



"Do it, don't do it. We don't care anymore."

To see why, consider the decision faced by American voters in this 2016 presidential election. Here you have two candidates which empirically are the least-liked in histo-

ry. If you attempt to vote on the issues, you eventually realize there are too many to consider: climate, energy, health, taxes, economy, trade, war, education, technology, corporations, Wall Street, abortion, drug-legalization, civil rights, and many more. It's almost impossible to agree with any one candidate on *all* the issues. And it's impossible to predict which handful of issues will actually get attention and change over the next four years; mostly governments produce gridlock, not change. This is not unlike the feature X / feature Y dilemma—too many considerations, and unknown which of those considerations might actually matter in the end.

What a lot of voters do, perhaps unconsciously, is select a few issues for which they possess a special affinity, and vote on only those. So for example, someone might select "elections are corrupted by money" as a key issue, and therefore support Bernie Sanders even if they agree quietly that his economic plan doesn't add up. Or someone might select "social justice" as a key issue, and therefore support Hillary Clinton (and the justices she will appoint), even if they agree it can't be appropriate to delete emails that are under subpoena. Or someone might want to "throw a bomb into the institution of government" and elect Donald Trump, even if they agree he has and will continue to say and do atrocious things.

This sort of simplification is logical, and necessary. We've talked about this before in the context of SaaS metrics. We sometimes have to be so simple as to be reductive, so we can make a clear decision.


Here are two tools that use this insight to arrive at good decisions:

- Binstack: How to make multi-dimensional choices.
- Fermi ROI: The better way to run rubrics or "ROI" calculations.

If after the examples in those articles you're still uncomfortable about ignoring a dozen important things, consider this: If you treat all those things as "important," you'll end up with a set of "good" choices. *Any* one of them is, objectively, good. Therefore, if you further refine your

process to narrow down to just one of those, *almost any process whatsoever* is acceptable, because you're picking from a set of already-good choices!

Inevitably, however, with whatever options you excluded, you're likely to field complaints, disappointment, and argument from others who have some vested interest in the losing options, even if their interest is simply that it was their idea, or they had gotten attached to it. For them, it will be easy to say "but how is your choice *really* better?" And they are right, but only because you had multiple, equally-good choices. Whatever you do, someone is going to say that. So, you should not include this "social pressure" in your decision.



"Reaction to the news was mixed, largely because we asked more than one person."

Decide how to decide, and get on with it

Making decisions quickly is valuable. You can usually make a different one if the first one proved wrong. (If you can't, that's a reason to take more time evaluating the decision.)

So pick a process, use it, and move on.

Separate the upside from the downside. Clarify the choice using one, two, or at most three dimensions. Base the primary decision on the magnitude and likelihood of the upside, and use the downside to veto untenable options. Use Binstack or Fermi ROI as frameworks that guide a group to a good, joint, explainable decision. Remember that defining the conditions of the decision is 80% of the decision, so invest your time in that.

And don't be too hard on yourself after the fact, even if it goes badly. Life is an experiment with little predictive power, no control group, with $N=1$, and which cannot be re-run.

They can't all be zingers.

Printed from: *A Smart Bear*

<https://longform.asmartbear.com/complex-decisions/>

© 2007-2024 Jason Cohen  @asmartbear