

Overtake on the Turn, Overwhelm on the Straight

Sometimes, doing the right thing is just way too hard. So you have use the best approximate substitute available. When you can't fly like a bird, you can aspire to be a frog that can jump really high, or a flying squirrel.

Decision-making is like that. There is, in my opinion, a "right way" to do decision-making in complex, dynamic environments (VUCA conditions — Volatility, Uncertainty, Complexity, Ambiguity), but most of the time, the right way is way too hard. So like most people, I use approximations tailored to current conditions (I am partial to the geeky joke that life isn't just hard, it's NP-hard).

To explain the right way and the approximate way, it helps to think in terms of high-speed maneuvering as a metaphor. Think of the dog-fighting in-an-asteroid-field scene in Star Wars. There are unpredictable moving obstacles and adversaries in the environment, and potential/kinetic energy considerations arising from the physics and energy levels of your own vehicle.

The "right" way to engage such a domain is with high situation awareness and calm mindfulness. Such a mental state allows you to maneuver smoothly and efficiently, with surgically precise moves that minimize entropy generation while achieving your objectives. This is the peak-flow-state, with your OODA-loop humming away at Enlightenment Level 42.

Unfortunately, if you're like me, you're in that state perhaps 1% of the time. What do you do at other times?

OODA for Zombies

The rest of the time, there is a good chance you are an over-caffeinated, under-exercised, manic-depressive, financially precarious, junk-food-eating zombie. And even if you're not a total train-wreck, you may have over-reached so far beyond your capabilities that you might as well be.

It's nice to have the meta-goal of increasing your peak-flow-state periods, and [we've talked about that a few times](#), but truth be told, it's nice to have a crutch to get you through zombie periods.

Once such crutch is replacing the satisfying full-blown dogfight-in-an-asteroid-field metaphor with a simpler and less satisfying one: an auto racing metaphor that is locally and approximately correct a lot of the time.

The racing metaphor is simpler because it involves no moving dumb obstacles like asteroids, adversaries who maneuver in simplified ways, a constrained competition model, and a playing field that's a closed circuit with a simple geometry.

The best-known use of the racing metaphor is the heuristic, *overtake on the turns*.

The reasoning behind this heuristic is that a turn, where potential and kinetic energy must be traded off in controlled ways to reorient a vehicle, is where skill and higher situation awareness can beat raw power. When power is roughly equal between competitors, skill differences at the turn are the only thing that can change the leader board.

A closed-circuit race *decouples* agility-first and energy-first epochs. If the course is simple but very long, you can even map planning/execution to turns and straights.

The Turn and the Straight

But I don't hear much about the obvious companion heuristic (I had to make up a mnemonic phrase): *overwhelm on the straight*.

If you have a true power advantage, you have to use the relatively straight portions of the course to pull away from the rest of the field.

And you have to *overwhelm*, which means using your power advantage to really put a lot of distance between you and competitors. It is not enough to get just a little ahead. You have to get as far ahead as you can.

Why? Because under zombie, non-peak-flow operating conditions, a trade-off holds between power on the straight and agility in the turn. You are unlikely to be good at both.

To take a simple example, imagine that you're racing with just one other adversary on a rounded-corner rectangular circuit. You gain 6 meters per straight due to greater power, and lose 5 meters per turn on every turn due to poorer turning skills. So your net advantage isn't 24 meters/lap. It is 4 meters per lap, and you have to go all out to get it.

You cannot be lazy or hold back. Depending on your *overall* power reserves for the race, however, you can and should cleverly time when you choose to go full throttle.

The key to the timing there is the *length* of the race.

Short versus Long Races

The shorter the race, the greater the advantage for the competitor with better turn performance.

What's more, when you're neck-to-neck, the more agile competitor has an even higher advantage, because the more powerful adversary is likely even worse on a neck-to-neck turn than a solo turn. So the agile adversary will likely gain more per turn when the gap is smaller.

On the other hand, long races favor the more powerful players (so long as they stay focused on the race and don't get distracted).

There are two reasons for this. The first is **decision fatigue**.

Navigating a turn smoothly drains the executive function because it is more demanding. By contrast, powering through a straight requires far less executive function control. It is more a function of raw energy levels.

This means, the advantage of an agile adversary is going to slowly decay through multiple turns, as he/she gets mentally fatigued.

The second reason is local learning.

Situational learning is easier and cheaper than generalized learning. An agile competitor has a starting advantage based on *generally* superior turning skills that are most potent when the course is equally unknown to all. But on a specific simple course, such as our simple rounded rectangle, a power-competitor can quickly and cheaply learn the local turns and neutralize the advantage of the generally agile competitor.

Fast-Following as Co-Opted Agility

When the turns are a repeating pattern of the same four turns, rather than a sequence of unexpected turns, generalized agility quickly becomes useless. It is equivalent to a starting position advantage on a completely straight and long course: a finite advantage that just delays the inevitable against a more energetic racer.

This means the longer the race, the more you can rely on greater power. In our simple example, assume that the more powerful competitor always gains 6 meters per lap. But if the turn performance advantage of the agile adversary erodes to nothing after (say) 100 turns (25 laps), the race turns into a pure straight race with a 24 meters/lap advantage instead of a 4 meters/lap advantage.

Once you've fatigued the executive function of an adversary, and done your local learning, you can pull away steadily. In fact with a limited overall energy budget, it might be useful to stay slightly behind initially, as a racing strategy, making the turn-advantage competitor work harder for their turn gains, fatiguing their executive function more quickly while learn-

ing faster yourself.

So the presence of a more *generally* agile competitor means you can learn a localized (and therefore cheaper) version of their applied generalized tricks via imitation.

Once you've exhausted them, you can start developing your straight advantage and pull away. This strategy works *even* better if the adversary also has a lower *total* energy store. You simply wait for physical or mental exhaustion (whichever happens first) before *you* get started, just keeping up until then.

They waste their limited energy learning, you imitate and save your greater energy reserves for scaling and winning.

Many fast followers (such as Microsoft) seem to operate this way, effectively co-opting more agile competitors as unwitting scouts. This is the reason [imitation is so much more successful than innovation](#), in terms of returns.

Exponential Breakaway

The racing metaphor actually understates the advantage of being a power player, because in many real-world situations, gains compound rather than accumulating linearly.

Imagine a weird sort of race where the more you are ahead, the more you can pull ahead. So instead of always gaining 6 meters/straight, you gain 6 meters/straight if you are 1 meter ahead, 12 meters/straight if you are 2 meters ahead, 24 meters/straight when you are 3 meters ahead.

You get the idea. You can *really* overwhelm on the straight in long races, and the longer the race, the more you become impossible to catch. It can be a crippling demotivator to watch an exponential breakaway from behind.

I've struggled to articulate this in previous posts, but I think I finally understand the phenomenon correctly. You can check out [my old post](#) on the subject if you want a more complete (if more confused) deep dive.

I am interested though, in whether there can be a similar advantage for the agile under sustained VUCA conditions.

Exponential Turnaway

Let's call the corresponding idea of a runaway agility advantage an *exponential turnaway*.

I don't think such an advantage can exist in a simple situation like racing on a known closed

course.

But when we return to the full-blown metaphor of dogfighting in an asteroid field, it may be possible.

Here, exponential turnaway may not depend on compounding gains for the more agile competitor, but compounding *losses* for the less agile one.

Military strategists often talk about the idea of a “target-rich environment.” An asteroid field is sort of the flip-side of that idea: a “projectile rich environment” that disproportionately penalizes bigger players (size is generally correlated with power and lower agility).

Instead of a known sequence of turns coming up with known periodicity, the non-agile player faces an unpredictable environment of unknown maneuvering challenges.

From the point of view of how quickly decision fatigue sets in, I suspect, the less agile competitor will quickly get overwhelmed, with compounding errors, bad decisions leading to worse situations, and eventually, an impossible-problem situation.

Of course, the more agile competitor is also experiencing decision fatigue, but he/she doesn't have to survive forever in the asteroid field. Just long enough for the competitor to crash and burn. At that point, he/she can exit the asteroid field, and recover in relative peace.

Exponential turnaway can particularly affect big companies that try to “fast follow” an entire *swarm* of little startups into a new market, instead of just a single agile scout. The leader emerges via a shakeout in the swarm rather than a powerful incumbent fast-following an initial leader. Neither Microsoft nor Google won the social network game. Facebook and LinkedIn did, via a shakeout in an initial field of many players (so it is inaccurate to say that Facebook was a “fast follower” to MySpace. It wasn't. The sector wasn't really viable at that point in the game, and Facebook did not really win by imitating MySpace the way Internet Explorer imitated Netscape).

In such a situation, even acquisition may be an impossible strategy. By the time a clear leader emerges via a shakeout, it may be too big for a wannabe “fast follower” to swallow. This is why, despite the generally greater returns for imitation over innovation, the innovation game is still worth playing, at least for the rare shakeout winner.

If you are a pioneering startup afraid of a big company fast-following you into markets you created, it might actually help you to help out other related startups a bit. They will tweak your formula and pursue it slightly differently, creating a shakeout game rather than a fast-follower game.

This is not about socialist cooperation, it is about creating VUCA for a more powerful adversary.

There is a Hunger Games joke in there somewhere (I just watched the first movie yesterday).

Startup Maneuvering

The Lean Startup types among you will have noticed some obvious connections to those ideas. In fact, this post partly grew out of my efforts to try and figure out what bothered me about the model.

It is by now widely recognized that the Lean Startup model can, in many ways, be regarded as an approximate version of the OODA loop. But the precise nature of the approximation involved has been eluding me.

The *symptom* of the simplification is that the Lean Startup replaces the core “get inside the tempo of your adversary” competition model of OODA with a “faster iteration” competition model.

Here “adversary” can stand for either a literal competitor, or a responsive market (i.e., an asteroid field that responds to your actions, such as customers with some control of their own decisions and learnable reactions to your moves).

I spent some time thinking about conditions under which the simplification is valid, and the best I’ve been able to come up with is this: “faster iteration” is a valid substitute for “inside the tempo” when for some reason, the learning per iteration has a fundamental rate limit.

What do I mean by this? Imagine a dumb startup and a smart startup trying the same first A/B test on identical minimum viable products (MVP). Both get the same results, but the latter *learns* a lot more, by extracting more intelligence, and is able to pick a second experiment (or pick a pivot direction) much better.

Under these conditions, iteration speed is not very relevant. By simply extracting more intelligence per iteration and pivoting smarter, the smart startup can gain on the dumb startup even if they are iterating slower. In optimization terms, the smart startup will be extracting the signal and climbing the right hill. The dumb startup will be thrashing around in response to the noise and failing to see the hill.

But *if* for some reason, the amount of net learning per iteration is rate-limited, an inefficient learner iterating faster can pull ahead.

Near-Random Domains

This corresponds to conditions where there is almost no scope for generalization either for future decisions in the same local domain, or other similar domains. Or to use a generalized form of our race metaphor, everybody is equally bad at the turns, so the race goes to the better straight performer by default, no matter how short the straights or how long the race.

This is a world where agility is meaningless. There is no such “skill” in the picture. You have a pure straight race where everybody gets knocked down in a completely unpredictable way by nature every few dozen steps. So it all depends on getting back up faster and running as hard as you can till you are knocked down again.

So a lean startup mindset is good if your learning resembles memorizing the digits of a random number, with one digit per iteration. Being smarter doesn’t really help. Being faster does. Every bit learned is expensively bought via a knockdown.

So it is not surprising that the lean startup is most popular in a domain that is effectively close to random user behavior around Web technologies.

Any truly generalizable insight spreads quickly via imitation, since UX design IP is hard to protect. What remains is the random-bits bleeding edge of highly localized user behavior learning (localized down to the single click in a single fixed context, with an additional advantage for enterprise software, which is more localized than the consumer world). The faster you can accumulate these results, the faster you can move (a good sign is that people aren’t particularly good at predicting the results of non-trivial A/B tests).

A clear illustration of this principle is this well-known factoid about memory: chess grandmasters are no better than random people off the street when it comes to memorizing random board positions. But when the board positions are realistic and legal, they do far better.

For a startup, near-random domains make the lean model useful, especially for enterprise software where there is a lot of corporate arbitrariness to be learned for every individual customer. For big consumer web companies, this translates to massive amounts of continuous, almost automated testing.

But when there is more general domain structure, it pays to give up some iteration speed for smarter learning: going for the “inside the tempo” approach.

So paradoxically, if you think you have an agility advantage, it pays to *not* be at the bleeding edge where all learning is almost random. It pays to withdraw to a less random domain where your superior generalization skills are an advantage.

I am still not completely satisfied with this model, but it’s getting there. I’ve been working on these ideas for almost two years now, and they are finally getting precise enough that I could

mathematically model them if I wanted to (starting with this metaphor of a race-to-asteroid-belt spectrum, and generalizing to exponential breakaway/turnaway type domains).

For now, what I've arrived at is a set of basic rules for playing the meta-game.

Recognizing, Picking and Creating Games

In business and life, the game you pick or define is far more important than how you compete within the game. If you are a power player, you need to pick a simpler game that is mostly straights. If you're competing with an agile player, try to draw him/her "out into the open." If you're an agile player, try and draw the power competitors into an asteroid field. Guerrillas retreat into the hills. Dictators try to flush them out into the open. It's an age-old game.

There are at least four levels of the meta-game here.

1. *Recognizing*: As a beginner, you must first learn to figure out what sort of game you are in: racing circuit, asteroid field, or true-random. Pursuing a racing strategy in an asteroid field, or vice-versa, is dumb. Trying to be meaningfully "agile" in learning a random number is beyond dumb. You must also figure out which of the games suits you best.
2. *Picking*: As an intermediate player, you must learn to pick your game. If you have an agility advantage, find a more complex game — an asteroid field. If you have a power advantage, find a simpler game (a race circuit), or a truly random game, which is simple in its own way (note to computational complexity geeks: there is a potential phase-transition diagram here with agility in the middle, I'll leave you to sketch it out yourself).
3. *Creating for yourself*: As an advanced player, you have to learn how to *create* a game that suits your strengths. If you're a power player, you must learn how to simplify the game through meta moves. If your huge starship is being drawn into an asteroid field by a maneuvering rebel fighter, use big-bang type weapons to pulverize the field. If you are an agile player, complicate and confuse the game for your powerful adversary.
4. *Creating for others*: As an enlightened player, you have to create a game not just for yourself, but your adversaries. One of the most extreme techniques is to create pseudorandomness in a domain that actually supports learning. This means taking a relatively clean signal, using it to learn yourself, and mixing in enough noise for your adversaries that they slip into "faster and faster" mode, mistaking an agility-friendly domain for a power-friendly one.

At the heart of this process is increasing self-awareness. You have to understand your own relative energy and agility capacities, and how they change as you gain experience through one game after another.

What makes competition fun is the loop between Step 4 and Step 1, which creates an arms race. It is often impossible to tell whether you are truly in a near-random domain, or whether a smarter player is creating pseudorandomness for you. When you recognize the latter case, you level up.

I am interested in armchair-thinking through these ideas in the context of real cases and examples, so if you're interested and have a suitable case, give me a call via [Clarity.fm](https://clarity.fm).

Data is Eating Clocks

It struck me recently that Marc Andreessen's now-famous observation, that software is eating everything, has a special case that is particularly interesting for students of the history of the industrial revolution.

Data is eating clocks.

Fifteen years ago, I used to wear a watch. One day, I lost it and never replaced it. The only time I look at a clock these days is when I have to catch a train or plane. I only think about the date when I have to sign a legal document. Most of the time, the day of the week matters more.

The clock was both a motif for the industrial revolution and a critical piece of technology driving it. Every small town in Europe gradually acquired a village clock tower. In the US, time zones emerged alongside transcontinental railroad clocks.

One reason precise time-keeping was so important in the industrial age is that when data is scarce, synchronization becomes critical to many activities. If you don't know where your friend is, you have to set a precise time and place to meet: "let's meet at Starbucks at 10:30. But if you can text, you can coordinate in much looser ways: "I'll text you when I am close to downtown and we can figure out where to meet."

Behavior becomes more responsive to real-time situational details, and more robust to delays. Synchronization, a fragile coordination technique, becomes less necessary.

Interestingly enough, Chet Richards, a close associate of John Boyd, told me that Boyd hated the idea of synchronization, which was antithetical to his conception of maneuver warfare. Synchronization, however, was central to the idea of network-centric warfare, which is often viewed as an opposed doctrine.

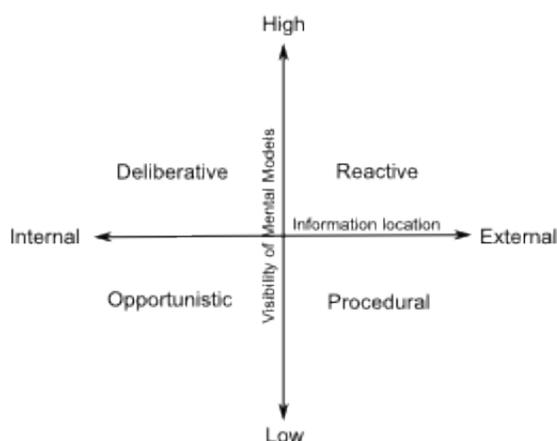
I think the *human* world is increasingly going to become liberated from clocks and calendars. This is the literal manifestation of atemporality. Clocks will remain extremely important to coordination between artificial technologies, however. Cellphones, satellites, data centers: all need very precise clocks to talk to each other properly.

The artificial world is going through its own industrial revolution apparently, going by the increasing importance of clocks to the inner workings of technology.

Annealing the Tactical Pattern Stack

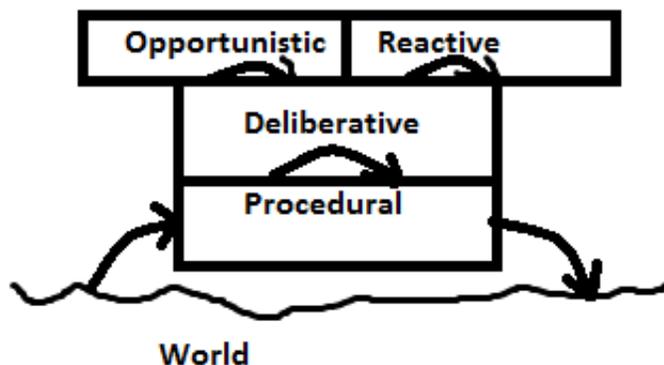
Human behaviors are complicated things. They are easy to describe, as fragments of narratives, but hard to unpack in useful and fundamental ways. In *Tempo*, I offered a model of behavior where **universal tactics** (universal in the sense of arising from universally shared **conceptual metaphors**, and being enacted in domain-specific ways) form a basic vocabulary, and are enacted through basic **decision patterns**, which are like basic sentence structures in language.

I suggested that there are four basic kinds of tactical pattern: *reactive*, *deliberative*, *procedural* and *opportunistic*, that could be conceptualized via this 2x2, where the *x*-axis represents the locus of the information driving the action (inside/outside your head) and the *y*-axis represents whether the information has high or low visibility (i.e. whether it is explicit and in awareness, or whether it is part of the frame/background, and below awareness).



The four basic decision patterns

While writing the book, I tried to figure out whether these behaviors also form a natural hierarchy of sorts. I was unable to make up my mind, so I did not include the idea in the book. Now I think I have a good model. The stack looks like this (the simplicity is deceptive):



Why? And how should you understand this diagram?

It took me a long time to get to this, so the logic is deceptively simple. Procedural patterns are the most completely learned behaviors inside your head *or* the most solidly codified and embedded patterns in the environment (via field-flow complexes for instance). Sometimes these patterns have sunk so low that they are not in the brain at all. Inside the body, that corresponds to muscle memory. Outside, it corresponds to things that are not in the form of signage or obvious parts of user experiences, but things that have become very deeply **normalized**, so that we *cannot* think of them in other ways except under abnormal circumstances (such as the idea that we call our hand-held computers “phones”).

This layer contains *most* of the data informing behavior, and is embodied in ways that represent optimization for efficient execution, not comprehension, appreciation or awareness. It is what AI people call “frame information.”

The next level up is deliberative. This is the layer that contains our pre-computed inferences and decisions that arise from explicit mental models. Inside our heads, this maps to “planning” but that’s actually a relatively uninteresting special case.

The most common and rich mapping of deliberative patterns is to work-in-progress rituals that haven’t stabilized. Incomplete learning states in other words. Things that are not unconscious, have a certain visible and comprehended logic to them, and have become robustly algorithmic, thanks to having been repeated countless times, but are not yet automatic.

Finally, the top level is driven mostly by our attractions and aversions. Opportunism is mostly driven by attractions, while reactions are mostly driven by aversions (there is a lot more to say about this, but I’ll leave it at that for now).

Daily Routines as an Example

A good example is your daily routine. It is a relatively predictable behavior, composed of a set of individual behaviors built out to fill out a defined **epoch** (from sleep to sleep; reboot to reboot).

How do you actually execute a “daily routine”? Is it actually as predictable as the phrase indicates?

Here’s how I do mine (by design, most of my days are empty of hard calendar constraints; I avoid all but the most unavoidable of business meetings and move as much as possible to asynchronous interactions):

- **Pick a template:** You start at the deliberative layer, not the procedural layer, and make up a rough scaffolding for your day, based on a few different templates. For me, for instance, one template is: *pack lunch and gym bag, go to cafe #1, eat lunch, go to cafe #2, go to gym*. Another is *go to cafe #1, come home, eat lunch, nap, go to gym*. This template is something I pick without even thinking much, unless it is a “special” day in some sense (I have visitors, evening plans or meetings).
- **Populate the template:** Then I make a tentative commitment to *how* I am going to fill up that template. This is still deliberative. For instance, I might decide that I’ll spend the morning session working on a blog post, and the afternoon session on putting together some material a client has asked for. While picking out my gym clothes, I generally decide whether I plan to run on the treadmill or do some weights. I usually don’t pick the cafe I want to work out of for the day until I am actually in the car. I’ve sort of trained myself to figure out my mood and pick an appropriate cafe the moment I get into the car.
- **Navigate Content Disruptions:** Then email and news hit, during the breakfast radar check. About a quarter of the time, something happens to derail the *content* of my plan, but not the template. An urgent request from a client or friend, or a great idea triggered by a news article that I need to capture and flesh out via a brainstorm. These types of disruptions generally happen early in the day for me, if they happen at all.
- **Navigate Template Disruptions:** Templates rarely break due to logical reasons (like an emergency). They usually break for emotional reasons. I might get depressed and lose the motivation to work out at the gym, or my morning session might leave me so frazzled, I have to switch patterns and make it a nap day instead of a non-nap day. Or I get hungrier or more upset than I expected to get and eat junk/eat too much, which derails the rest of the day. Sometimes the disruption is via dissipation, and I get sucked into some sort of online click-click-click pattern of stupidity so that my plans get derailed entirely because of lack of fuel. I still stick to the structure, but the content becomes worthless. In all cases, the key to structural disruption of the template I chose is due to *unmanaged emotions*. As I get more practice with a new template, it becomes more robust to emotional disruption. This is the heart of *ritualization*. You build in subtle cues and micro-behaviors that make it *emotionally* resilient (as opposed to resilient to functional/logical breaks).
- **Navigate bugs:** If your chosen template is relatively new, it is probably buggy. For instance, I recently moved to a home that was relatively far from my favorite cafes and gym, compared to my old home. So coming home for a lunch and a nap became expensive in terms of time/gas. After eating out for lunch a little too often over a week, I started packing lunch. After realizing that I didn’t like to do the longer drive home in sweaty clothes, and that I wanted to sometimes work out during the lunch break between cafe sessions, rather than the end of the day, I started packing a change of clothes and show-

ering at the gym. These “bugs” in a new template need to be worked out while it is settling into a ritual. Once it is ritualized, it becomes harder to change.

- **Navigate deep breaks:** And sometimes, the disruption is serious enough that I go meta and start thinking about my framing routines and procedures. If I don’t exercise discipline, this can regress into pointless introspection. If I end up spending too many days in a row failing in a predictable way (for instance, planning to do two 4-hour work sessions, but consistently lacking the energy to get through the second one, which happened recently), I end up analyzing what is going on and modifying my behaviors appropriately at a deeper level. This is the *only* time I dive into the procedural layer of the stack, which contains the implicit logic assumed by the various deliberative templates. For instance, I used to work out thrice a week, but now I try to get in four or five workouts if I can, even if it is only some quick calisthenics. This solved my “enough energy for two sessions a day” problem.

People have different constraints, and bring different levels of discipline to daily routine development. Some have packed calendars that drains them of energy but requires less thinking daily. Others try to explicitly write down their plans for the next day at bedtime. Others don’t find the “morning derail” as valuable as I do (reactive and opportunistic patterns are how I get a lot of value captured, and are low cost for me, but for others it is expensive and low-value), and discipline themselves not to look at email until late in the day.

The Logic of the Stack

When you actually analyze a routine this way, you realize why the patterns are stacked the way they are. The stack reflects logical dependencies and obvious, overt abstraction layers in generating behaviors, but also sustains an entire range of learning, emotion-management, stabilization, error-correction and robustification behaviors. Good rituals don’t just happen. In our complicated times, they cannot even be reliably copied or imitated. You have to develop the meta-cognitive skills to constantly improve your routines consciously, otherwise they will fall apart. This is the deliberate practice of life itself, rather than of a particular skill like piano playing.

The stack as I’ve drawn it is similar to the natural architecture of the brain (fast, more unconscious inner feedback loops enveloped by slower more conscious, and more open loops, with feedforward). It is also similar to what is known as the “subsumption” architecture, a common high-level way to organize robots. It is also the usual way to construct abstraction hierarchies in computers.

Why is this stack so common?

It is because you have to think about two aspects of information in a behavior. The *cost* of

handling it, and the *predictability* of the information. High cost, high predictability information gravitates (or sinks) to procedural layers. Low predictability behaviors, whether or not it is costly, necessarily has to live near the top of the stack. A good system has very few things being handled by the top and a lot being handled by the bottom. If things are flipped, you are living a very behaviorally expensive life.

There is really no low-cost way to handle truly unpredictable things. But you can evolve to anticipate more and more (so you can *make* things more predictable as you learn). Your meta-cognition can become more efficient as well, so that the process of setting up a stack and “turning it on” in a new environment takes less time, each time (this is *adaptability* basically).

The procedural layer is something most of us aren’t good at hacking. Some people become extreme creatures of habit, and eliminate or block out environmental information flows that might destabilize the procedural layer.

Others are good at hacking the layer, but get so addicted to it that they never stabilize a procedural pattern long enough to use it effectively to sustain activity at the higher layers (higher is “more distant from the environment” in the diagram).

In OODA loop terms, the idea of getting “inside the tempo” of an opponent generally relies on hacking the procedural layer. I will be talking about this idea at the upcoming Boyd and Beyond conference in a couple of weeks.

Ritualization as Behavior Integration

Let’s take a stab at understanding this in a more advanced way.

Learning to manage this stack is not easy at all, and it is best to start with simple and short templates at the deliberative layer. So using something like the Pomodoro technique might be a starting point for some. Others might want to start with bedtime or morning rituals. Or whole-day rituals that are artificially constrained in complexity through external aids (this is what any sort of “school” achieves).

From there you can build “out” in complexity and duration to days, weeks, months, years and a lifetime. You can also build “inwards” and make things far simpler and shorter. This is the direction of various meditative practices. I think our home zone is 20-25 minute periods with 2-3 moving parts in the deliberative templates. Both longer/more complicated and shorter/less complicated take effort.

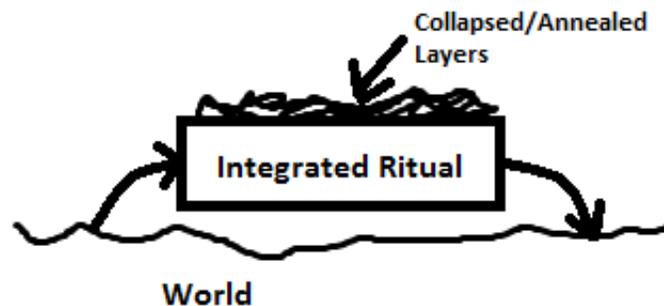
To improve your ability to program yourself, you have to constantly go both inwards and outwards. Build out too much, and you might end up with Jenga towers of instability because the bigger the structure, the more robust and fine-grained the building blocks have to

be. Go inwards too much, and you might lose the ability to handle larger ranges of volatility and complexity: like building perfect bricks, but never learning the broader principles of building houses.

It is sort of a fractally multi-scale approach to a mindful life that has a dynamic range from emptiness and tenths of a second on one extreme to extremely cluttered complexity and decade-long attacks on [wicked problems](#) at the other.

When you work on a stack repeatedly like this, it is sort of like annealing a piece of metal. The explicit divisions in the stack start to dissolve, and the whole thing slowly becomes one highly resilient, unified and integrated behavior that appropriately blends a potential for reaction, opportunism, deliberation and procedure into one unified whole. The *meta cognition* itself becomes proceduralized in the middle stages, at which point you hit increasing returns. Towards the end, you exhaust the variety in the local environment and the opportunistic and reactive layers start to atrophy due to lack of use.

So the whole stack slowly collapses into the procedural layer. Since this is the layer that gets “first shot” at reacting to an input, this is like all the logic of the stack gravitating to the surface, in a manner reminiscent of holography and the idea that the information in a volume can live on the surface. If the annealing has been decreasing the “grain” size of your templates and patterns all along, the whole behavior starts to get very dense. All your moves become [rich moves](#) in a sense, and achieve extreme narrative density. The metaphor of a neutron star is very compelling here.



The downside is that unless you happen to be the black hole at the center of the universe, this beautifully ritualized state of mindful living is a local optimum. If you are thrown into a sufficiently alien environment, you *will* struggle and go through the whole process of creating and collapsing a new behavioral stack of patterns.

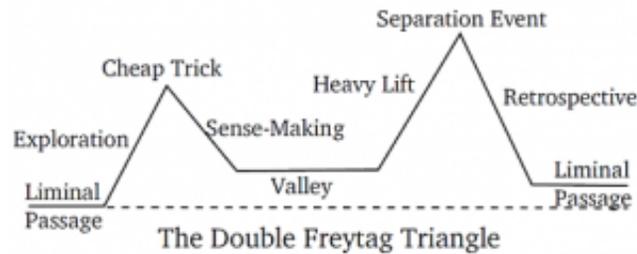
Accepting this process of discontinuous learning across many local optima is sort of the defi-

inition of mindful living. There is a limit to how far you can go though, since you have a finite brain with a design lifetime of about 70 years and high but limited plasticity.

Thanks to Greg Rader and Ho-Sheng Hsiao for useful discussions that helped me lay out the ideas in this post systematically.

Analysis-Paralysis and The Sensemaking Trap

Analysis-paralysis is when you get into a loop of continuous analysis that prevents you from breaking on through to the “other side” where action can begin. I am beginning to get a handle on the problem, but it is not going to make much sense to you unless you’ve read the [book](#). So this is in the advanced/extra-credit department. Perhaps after some more thought I’ll be able to capture this idea in a simpler way.



In the [Double Freytag](#) model of narrative decision-making, analysis-paralysis corresponds to getting stuck in the [sense-making](#) phase. Why does this happen?

The problem has to do with the nature of the insight (or [cheap trick](#)) that triggers a gear shift from [exploration](#) to sense-making. If it is an *appreciative* insight, it will help you make sense of what’s going on in the situation without necessarily creating the focal point for action. But on the other hand, it is not a purely *manipulative* insight either (i.e., it is more than a mindless hack). I covered appreciative versus manipulative knowledge in the last post.

My current definition of the cheap trick covers this ground, but clumsily. I am trying to craft a better definition. The Boydian/classical German notion of *Fingerspitzengefühl* gets at part of what I am getting at here, but it too is conceptually inadequate, since it is more of a continuous notion than the *coup d’oeil* from which “cheap trick” is derived. The related idea of *schwerpunkt* also gets at some of what I am getting at, but it describes the locus of action concentration where a breakthrough is most likely, rather than the organizing larger insight that helps you discover where the *schwerpunkt* might be.

So to slam three concepts together in sequence, a *coup d’oeil* suggests a *fingerspitzengefühl* pattern of sense-making action that triggers a *schwerpunkt* detection, followed by force concentration and breakthrough to the valley, where you can apply steady effort towards eventual accelerating returns. In the 2.0 notion of cheap trick I am trying to construct, the success of the second and third pieces hinges on the quality and DNA of the first piece (the German outflanking of the Maginot line in World War II is a classic illustration of this sequence).

So there is a lot of phenomenology wrapped into the single cheap trick moment. It is the moment when thinking must shift across the “air gap” between discovery to synthesis. For a suf-

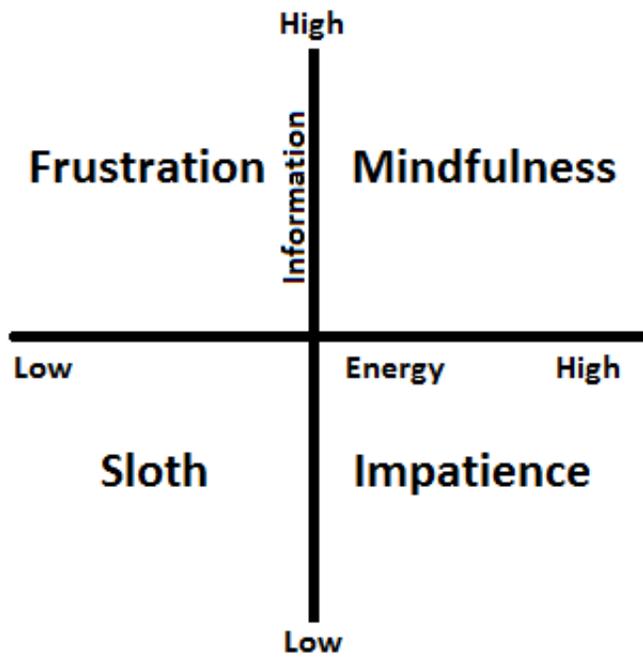
ficiently complex problem, there is always a jump. I've never experienced a smooth transition from exploration to (successful) sense-making.

The key to the problem is the idea that there is no such thing as pure analysis. There are only two possible destinations when you make the leap of faith from exploration and discovery: analysis-paralysis and analysis-synthesis. This is another way of saying that analysis and synthesis cannot be decoupled. If you try, you will find that you cannot break through to decisive action. The shift from sense-making to the stable-tempo action phase (the [valley](#) in the Double Freytag) will turn into an impassable barrier. But there is nothing important happening there. The cause of the barrier lies in the cheap trick at the beginning of the sense-making phase: it has some sort of fatal flaw at birth.

If the cheap trick is the right concoction of appreciative and manipulative insight, the shift to the valley will be natural and effortless. If you run into a wall there, you have no option but to go back to discovery and look for a fresh cheap trick. This may not be possible, since life doesn't offer as many do-over chances as people believe.

Steer, Ready, Fire

I like various permutations and adaptations of the phrase *ready, aim, fire* to think about decision-making between the extremes of pure contemplation and pure action. Playing around with this phrase led me to this 2x2 (I seem to be thinking a lot in 2x2 form these days). I'll connect the dots in a minute.



Aiming versus Feedback

The apparently logical sequence, *ready, aim, fire* describes a feedforward model. You get your mind in the right place, then you figure out how to be effective (*aim* can map to waterfall planning at any level), then you take action.

The phrase *ready, fire, aim*, preferred by the action-oriented in uncertain and dynamic environments, is a response to the analysis-paralysis that can happen if you try to get to ideal starting conditions and perfect information before starting to act.

The absurdity of aiming after firing can only be resolved via appeal to the logic of iteration and feedback. You converge on the successful course of action through feedback from failed actions. This works well as a motto for startup types and others who believe in the *release early and often, and fail fast* approach to projects.

Then there is the phrase, *ready, fire, steer*. I am not sure who came up with that one, but I've heard it attributed to Paul Saffo. This replacement of *aim* with *steer* suggests that real-time

feedback and control can be continuous. It is the logical limit of iterating faster and faster. Heat-seeking or radar-guided missiles are perfect examples.

The Role of “Fire”

The variant *ready, fire, steer* made me wonder about why *fire* is even necessary. Within your basic firearm metaphor, firing gives you all your momentum (kinetic energy) in one big dose. Of course, you also have whatever positional advantages (potential energy) you possess. It maps well to situations like getting investment in a startup, coming into a trust fund, or using a rocket to launch satellites.

But there are also cars and airplanes, with more continuous energy-generation models. There are also renewable energy models like sail ships, and models that create a net surplus of energy, like a solar car with more energy than it needs.

These don't need a *fire* step. You could do with just *ready, steer* thinking (or *ready, start, steer* if you insist). A lot of bootstrapped business models would qualify, as you use tiny or zero cash investments to get started, and nurture cash flows slowly to get where you want. You may be accumulating a surplus of cash or attention that you can conserve for later use.

It takes a lot more foresight to work without the boost of a *fire* stage, but in return you get more control and efficient use of resources, in cases where the *fire* represents borrowed energy, provided on terms that you don't like.

In fact, you can often dispense with *ready* as well. The idea that you need a *ready*, independent of information preparedness is more psychological fiction than reality. While you are contemplating doing anything, your readiness level changes over time, even before you adopt any sort of intention. As you process relevant information, your situation awareness may increase or degrade in quality, and you may become more or less oriented.

Ready really only matters in situations where there are decisive go/no-go thresholds defined by irreversible (or very expensive to reverse) actions, such as quitting your job or getting married, but *ready* as an internal state doesn't really capture that. You'll never be really ready. But as a continuously-changing state, your *readiness* may cross a minimum threshold associated with a given irreversible decision. That threshold is set by external conditions.

This means that you start steering the moment even a tiny amount of readiness bubbles up into your consciousness. After that, the feedback process that is *steer* automatically moves your readiness level along.

So *steer* is really at the heart of it all. Continuous feedback control of energy, using information.

Ready is useful to add in where there is an important, unavoidable and irreversible decision inside the decision process.

Creating an Opening

Fire can actually come at the end as well, and this is the case that interests me the most these days .

In cases where you maneuver for an opening starting from unfavorable conditions (*ready, steer*), you could be accumulating a surplus capacity for action while waiting for a good opportunity to use it.

This could be a purely passive wait, or you could be actively trying to engineer an opening through “set up” moves.

This accumulating surplus might be money, information, a slowly-grown marketing asset like a blog, or going to night school to get a degree. Or it might simply involve waiting and watching for environmental conditions, trending in a certain direction, to hit a threshold.

Within a large corporation, this could be a matter of making specific allies and accumulating a strong position around a currently unattractive business asset (such as a dog of a product that people think cannot do well in the future, or a sales region that nobody wants) and waiting for, or engineering, a way to work it.

For example, there was an optimal window of time for streaming video businesses to be launched, based on falling bandwidth costs. If you were in that business, you’d have been wise to adopt a *ready, steer* hold-and-accumulate strategy, waiting for your moment to *fire*.

Today, the emerging sector of 3D printing is in the *wait* zone for many people: once the technology becomes sufficiently cheap and some basic technology to exploit it has emerged (such as stable, cheap and easy to use software for generating designs), a lot of people are going to jump in.

Bootstrapping to Big

Since *ready* has to do with crossing externally-determined irreversibility thresholds more than being in some mystic state of perfect readiness, the *steer, ready, fire* sequence is great for maneuvering to create an opening, and then triggering an irreversible action that requires a burst of *informed* energy. This is what is typically referred to as a *go-big-or-go-home* moment.

One application of *steer-ready-fire* thinking is bootstrapped businesses that intend to grow big at the right time. These days, we’ve somehow bought into the illusion that bootstrapping is

for lifestyle businesses and that you need professional investors to go big.

This is obviously false. If you *steer to ready* with sufficient foresight, carefully build cash-flow assets, and wait for or create the right opening, you can bootstrap *and* go big. Many big businesses before the 1940s were grown in precisely this fashion. Before investment banking became a big business in its own right in the 1870s in America (and later, the sub-sector of venture capital in the post World War II era), big fortunes — including those of the two biggest Robber Barons, Vanderbilt and Rockefeller — were built through this sort of bootstrapped, leveraged model. There were times when Rockefeller in fact had more capacity to move the markets from the outside, than his famous finance contemporary, J. P. Morgan, had on the inside.

Stepping back a bit, what's common to all these approaches to thinking about decision processes is the interplay of energy and information in some abstract sense (where energy can be money or marketing potential for instance, in our running startup sector example). Acting with either too much or too little information, given your energy levels, is inefficient. Having neither information nor energy is of course a stable situation.

Mindfulness is when energy and information dance together well. Note that you don't necessarily have to keep them balanced at a specific moment. You can store both. So you might wait for energy to catch up with information, or vice-versa. Or you can accumulate both and unleash a ferocious burst of mindful action driven by a store of heavily-informed energy.

Sudden Actions, Entropy and OODA

That last part (accumulating both energy and information to enable sudden movements) took me a while to get to. For a long time, I was unable to reconcile sudden, high-power movements with the idea of mindfulness because I was fixated on the thought that mindful actions are necessarily *smooth* actions. They needn't be. Jerky movements have a role to play in our world.

But there *is* a deeper level at which "slow" and "smooth" matter. This is where an abstract notion of entropy is relevant. Slow, smooth actions cause low increases in entropy. Quick, jerky actions cause high increases in entropy. Unfortunately, you cannot always work with low-entropy behavior because there is a lot of messiness in the outside world — the world that you don't completely control. The smaller and more closed your world, the more you can approach the idea of working purely with slow, low-entropy actions.

This is why readiness is best thought of in relationship to irreversible-action thresholds determined by external conditions. In thermodynamics, isentropic processes (those that don't increase energy) are reversible. Entropic processes are not.

When you unleash a sudden action, entropy will increase. In decision-making terms, it means you'll trigger action that is so fast that you cannot process the information being generated by feedback, so it will effectively act as noise. But there are situations where you know enough to know that this chaos you are unleashing will mostly favor you. This is reflected in the attitude that "I think it will all work itself out." Eventually, when the dust settles, you will be able to get back to a more mindful engagement with the situation.

And of course, there will always be net entropy increases even after the dust settles. Being mindful about this realization is the same as accepting the inevitability of death.

Of course, this extended thermodynamic metaphor needs to be carefully applied in abstract situations, but I believe the correspondence is a very close one. This thermodynamic metaphor, and the interplay of *ready*, *fire*, *aim* and *steer* in various permutations and combinations, is one approach to understanding how Boyd's OODA model really works.

Smart Money and Dumb Money

You can extrapolate this sort of thinking to larger groups and organizations, and think about how energy (usually money in the human world) and information are distributed within a organization and the environment it operates in. You can talk about whether energy drives information or vice-versa.

In larger systems of people, *power* distributions often emerge out of the interplay of energy and information. Smart money represents information in control of energy. Dumb money represents the converse situation.

In the world of dumb money, entrepreneurs must chase investors. In the world of smart money, investors court entrepreneurs. Why?

In entrepreneurship, *smart money* is often used to refer to investment from people who can also provide information and advice. This is actually not particularly smart money. If an investor holds all the cards — money *and* information — what exactly does the entrepreneur bring to the table besides talent? That sort of relationship defines employment, not investment. Truly valuable information comes from unlikely places. Information from well-known sources, such as seasoned investors or former entrepreneurs, is unlikely to be particularly special or exclusive. It is in fact likely to be common knowledge — it will help you lower costs of doing business, but not provide a competitive advantage.

A collaboration between a party with too much energy, and one with too much information, is fraught with tension. It is very hard to merge the two in mindful ways. One party is impatient and the other party is frustrated. Meetings between parties with unbalanced and com-

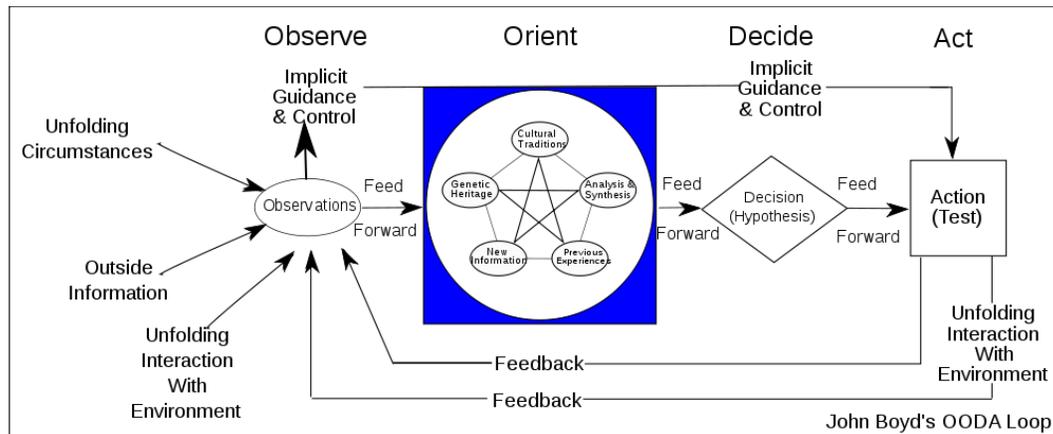
plementary assets, who are also mindful about what they have and what they need, are quite rare.

The result is that power dynamics are triggered while things are sorting themselves out. This is one reason I advocate a slightly evil philosophy. Engaging the world outside your personal control means dealing with all this. Trying to be purely good is like trying to work with just smooth, slow, isentropic actions. It is just not workable when there are transient openings and irreversibility thresholds in the environment.

So it isn't just individuals who have to gradually become more mindful decision-makers, gradually lowering the amount of sloth, impatience and frustration in their thinking. Organizations have to do it too. I can think of many frustrated, slothful or impatient organizations and groups, ranging in size from married couples to Fortune 500 companies and entire nations.

Tempo and OODA: The Backstory

John Boyd's [OODA loop](#) (observe, orient, decide, act) often comes up when I discuss *Tempo* with people from the more esoteric decision-making traditions. Very few people in the decision sciences are even aware of OODA, despite Boyd's significant technical contributions to fighter combat tactics and energy-maneuverability theory, which preceded his more conceptual, almost metaphysical OODA work. This is because, despite the very technical look of the classic OODA diagram, there is an element of mysticism surrounding OODA.



So I thought I'd tell the back-story of how OODA informed *Tempo*, and is continuing to inform the ongoing conversation that I hope will feed into a more ambitious second edition.

But first a heads-up: I'll be participating in a panel discussion about OODA and its relevance to the business/startup world next Wednesday, July 27th, 11:55 – 1:00. It's a free call-in webinar, but space is limited. If you're interested, [register here](#). The event is hosted by [Sean Murphy](#), one of my early supporters in getting the *Tempo* project off the ground, a few years ago.

Now for the backstory. There's two parts to it: the nature of the "Boydian community" itself, and how the ideas ended up informing *Tempo*.

Bottling Boydian Bushido

If you ignore that mystical element of OODA (as academia is wont to, and as I myself did, when I first encountered it sometime in grad school), you are likely to dismiss it entirely. Despite appearances, OODA is more Sun Tzu than feedback control block diagram. More Musashi than computer science flowchart. In fact, interpreted through those narrowly technical lenses, it can appear to be a clunky and impoverished model of human decision-making.

To get to the essence of OODA you have to take the apparently mystical element seriously.

You have to try to make sense of the cryptic utterances in OODA culture, such as "you must

get *inside* the tempo of your adversary” or “fight the enemy, not the terrain.” There are layers of meaning to peel back there. As my martial artist friend Ho-Sheng Hsiao would say, the OODA loop is more a *kata* than a recipe.

You could also think of it as a sort of *mandala*, to be meditatively pondered rather than merely executed. As you learn from his biography, Boyd was almost obsessively concerned with every detail of how he presented an idea, so you can be sure that every word and visual element in the diagram was put in with extreme care and deliberation. If something seems like an arbitrary or throwaway piece, you just haven’t unpacked its meaning yet.

If such mystic characterizations bother you, simply think of the diagram as a highly compressed (in the sense of file compression) representation of a much bigger gestalt of ideas. It must be uncompressed to be clearly understood.

I am not saying that the ideas in the mystic cloud around OODA are *actually* mystical (though some will no doubt disagree). They are simply very hard to demystify. We lack a clean vocabulary, so those who try in well-intentioned but naive ways often end up oversimplifying OODA to the point of triteness, rather than demystifying it.

As a result of this difficulty, OODA has been highly influential in the *practice* of decision-making, but not in the theoretical *study* of decision making in academia. Clean concepts and definitions are a prerequisite for that kind of scrutiny to be productive.

The influence of OODA among practitioners has been largely subterranean, passed on via illegible networks of influence, conversations, circulation of canonical documents that were never formally published, and very noisy word-of-mouth transmission of core concepts (in which I have participated in recent times).

There is of course a cost to this illegibility. Besides the noise that accompanies illegible intellectual cultures, Boydians can come across as almost cultish, and some of them to be honest, *are* cultish in their approach to the core ideas. This is one reason I don’t consider myself a true Boyddian. Apart from my general aversion to fanboyism, my other influences sometimes conflict with Boyddian orthodoxy (ironically, there *is* such a thing, despite the association of Boyddianism with unorthodoxy).

So I am going to be touching upon Boyddian themes (which are really part of a much longer intellectual culture) a lot in future blog posts, so as a heads up, that’s going to be my program of exploration: carefully, critically and selectively demystifying what I agree with and am able to clarify.

I hope to roll that up into a cogent treatment of adversarial decision-making in a future edi-

tion (Boydian thought is one of the handful of directions from which I am approaching adversarial decision-making).

But let me tell you about how this illegible, subterranean universe of ideas seeped into *Tempo*. It is the tale of an appropriately illegible pattern of influences on my own thinking.

Soaking up a Gestalt

The backstory is interesting primarily as an illustration of the degree to which a book is really the product of partly subconscious collective processes rather than the imagination of a single individual. As a result, there can be more inside a book than an author consciously puts in.

But that can only happen if you allow your subconscious to channel influences from a broader tradition of ideas into your work, and view yourself as an instrument of a grander historical process through which an idea discovers itself. You have to choose to take part in a particular bucket brigade. It's a matter of trusting your intuitions during writing, and keeping it partially permeable to osmosis processes.

OODA is one of the ideas that made it into *Tempo* more through such zeitgeist-channeling than any explicit inclusion decisions on my part. Though I only have a couple of passages explicitly devoted to OODA, the entire text has a somewhat Boydian flavor, which I never consciously intended to put in.

Allowing the zeitgeist channeling to occur is an admission that it takes a village to write a book, and those who are unwilling to make this admission end up suffering from what Bloom called [the anxiety of influence](#): an attempt to first socialize into and then individuate oneself out of, a broader intellectual culture. When the anxiety prevails, the work of individuals tends to get impoverished.

As I mentioned earlier, I mostly dismissed OODA when I first encountered it. Sean Murphy's feedback on a few early chapters however, made me take a second, harder look.

Back then, I *thought* I understood OODA, based on a casual first encounter and my own background as a control theorist (the field is all about feedback loops, and OODA at first sight just looks like a feedback loop diagram). With my second look however, I realized that there was a lot more to the loop. So I bookmarked it as a "Hmm! This really isn't a standard control-theoretic loop" but didn't dig much deeper.

In making the final cut of ideas to include in the book, I decided to leave out adversarial decision-making, so I limited my explicit discussion of the OODA loop to a few passages, and decided to defer further research into it. But my intuitions about the parts of OODA that I *hadn't* delved into back then, seem in retrospect to have guided how I wrote the rest of the book.

When I was nearly done with the manuscript, I finally had time to dig deeper. [Dan Pritchett](#)'s timely recommendation of Robert Coram's [Boyd biography](#) provided the perfect starting point. From there, I moved on to famous Boyd documents like [Destruction and Creation](#) and [Patterns of Conflict](#).

I was surprised to discover just how deeply the ideas in *Tempo* resonated with OODA. Though I was using a different vocabulary and exploring decision-making phenomena in non-adversarial settings, I seemed to have converged on many of the same core foundational themes, such as entropy, mental models, narratives and of course, tempo itself.

The validation was reassuring. Further conversations with people more immersed in OODA-lore, such as [Mark Safranski](#), increased my sense that despite my differing priorities, I'd basically been attacking the same mysteries that OODA takes on, from a different approach vector.

My exploration of OODA has only gotten more interesting since then. One particularly rich conversation I have ongoing is with martial artist [Ho-Sheng Hsiao](#) (Hosh). Thanks to his deep knowledge of esoteric martial arts and the philosophies of Musashi and Sun Tzu, he has been able to point out deeper layers of meaning in my own ideas. Though I originally introduced him to Boyd through my email newsletter, he now knows way more about the subject than I do. That's how these village-productions work, I guess. In particular, through some concrete examples, he helped me understand why "inside the tempo of an adversary" and "faster tempo than the competition" are not the same thing. Previously, I'd only had a vague intuition about this idea, based on a poor analogy to one signal interfering with another.

In some ways, with respect to some of the ideas I put into *Tempo*, driven more by intuition than analysis, people like Hosh and Mark understand what I've written better than I do myself.

I really like the thought of that: that perhaps by allowing myself to be unconsciously influenced, I've made the book smarter than myself. I suspect this is what Skinner was getting at in his famous "[on having a poem](#)" lecture (which is a sort of alternative statement of Bloom's anxiety of influence idea). One of the items on my to-do list is to reread *Tempo* thoroughly in about 6 months, to see if I can consciously pull out what I unconsciously put in.

Of course, there is no mystery or Jungian collective unconscious at work here. Thanks to Boyd's pervasive influence within the US military, I am sure I must have absorbed a lot of the ideas indirectly during my years as a postdoc working on Air Force sponsored research, during which I talked to pilots and pored through many military doctrine documents.

Ideas that Boydians understand through Boyd are also not unique to that community. There

are other traditions that also stumbled upon the same ideas in other domains. For example, Lakoff's model of conceptual metaphor is very Boydian, and Boyd's thought is very Lakoffian. The same can be said of James Scott's ideas on [legibility](#) and my own ideas on organizational dysfunction in the *Gervais Principle* series. At the risk of sounding wooly-headed, it's all connected. A lot of what I do is make those sorts of connections explicit and transparent.

Most recently, I was pleasantly surprised to discover, during a conversation with a chemist friend of mine at UCLA, that *he'd* been sort of rediscovering certain ideas I'd classify as "Boydian" through his study of how certain virus RNA molecules fold onto themselves. That conversation helped me connect some crucial dots that I'll talk about another day. Hint about the connection: the word *unfolding* occurs a few times in the OODA diagram; there is a certain sense in which OODA "unfolding" is related to protein folding.

With that cryptic teaser, I'll leave you to go ponder the OODA Mandala on your own.