Finding the Commitment

https://gerolds.github.io/textbook/textbook/posts/finding-the-commitment/

Contents

Finding the Commitment

A commitment you keep is worth more than a better commitment you abandon.

Making the Thing *(https://gerolds.github.io/posts/making-the-thing/)* argues that coherence requires a governing commitment, the thing you refuse to betray. But it assumes you already have one. This piece is about the prior question: how do you find it, and why does sticking with it matter so much?

The instinct is to search for the *right* commitment. Research, brainstorm, validate, iterate until you've found the direction that'll make the best game. But that frame misses where commitment's power lives. Not in correctness. In *stickiness*: the way a held commitment compounds decisions, aligns a team without constant negotiation, and turns doubt into forward motion.

Most projects don't fail because they picked the wrong direction. They fail because they never really committed to one. They kept options open, pivoted at the first sign of difficulty, and spent so much energy debating that they never moved far in any direction at all.

---

1. Why commitment is underestimated

Teams treat commitment like a hypothesis. "We'll try this direction and see if it works." Reasonable on the surface, but it misunderstands what commitment does.

A commitment isn't a prediction about what will work. It's an *agreement to stop asking certain questions*. When you commit to "the player excavates," you stop debating whether excavation is the right verb and start asking how to make excavation feel right. The questions change. The work changes.

Every hour spent debating the commitment is an hour not spent building under it. Every meeting re-litigating the vision could have solved a production problem instead. The team that commits to a

B+ direction and moves will outship the team that spends six months hunting for A+.

When you're committed, doubt becomes a design problem: "I'm worried excavation gets repetitive, how do we add depth?" When you're uncommitted, doubt becomes existential: "Maybe excavation isn't the right verb at all." The first question has answers. The second has only more questions.

A committed team can work in parallel. The artist knows what mood to hit. The programmer knows which systems matter. The designer knows what to cut. An uncommitted team must constantly synchronize, because nobody knows which direction the others are walking.

And commitment compounds. Every decision made under it reinforces the next one. The mechanics, the art, the tone, the pacing: they start to rhyme. A game built under a single commitment for two years has a coherence that can't be faked. A game that changed direction three times has seams everywhere.

---

2. **The cost of changing direction**

Everyone knows pivots are expensive. Few people know how expensive.

The obvious cost is rework: art thrown out, code refactored, levels redesigned. Visible, painful, and often accounted for.

The hidden cost is *alignment decay*. Changing the commitment doesn't just change the work. It changes everyone's mental model of the project. Some people update cleanly. Others carry the old commitment in their bones. For months after a pivot, you find decisions that assumed the old direction: a UI that served the old loop, an enemy designed for the old pacing, a level that tested the old mastery curve.

The deepest cost is *trust erosion*. Every pivot teaches the team that commitment is provisional. People stop investing fully in the current direction because it might change. They hedge their work, hold back their best ideas, wait until they're sure this direction will

stick. The team becomes risk-averse exactly when you need bold execution.

This is why frequent pivoting underperforms stubborn commitment. The stubborn team makes mistakes, but they make them together, in the same direction, and they fix them. The pivoting team never quite makes mistakes. They just keep starting over.

---

3. **The value of stubbornness**

There’s a mode of working that looks like stubbornness but is wisdom: refusing to reconsider the commitment just because the commitment got hard.

Every commitment gets hard. The loop doesn’t quite work. The scope balloons. The market shifts. A competitor ships something similar. A key person leaves.

These are real problems that need real solutions, but they almost never need a new commitment. They need *recommitment*: doubling down, finding a way through, trusting that the direction is worth the difficulty.

The teams that make great games nearly always hit a wall and kept going. They found solutions that wouldn’t exist if they’d pivoted. They discovered depth they’d never have found if they’d abandoned the commitment at the first sign of trouble.

Stubbornness has a dark twin, though: continuing without learning. Both feel like persistence. Both can be called “grit.” The difference is whether your effort is buying information or leverage, or just buying time to avoid a decision.

---

4. **The rationality of stubbornness**

Stubbornness is rational when it buys compounding. Games are integrative artifacts: a working loop tested against a clear commitment produces mechanics that rhyme, content that reinforces those mechanics, and depth that can’t be shortcut.

Abandoning a commitment mid-compound throws away not just the work but the *coherence* the work was building toward. Pivoting doesn't just lose the days. It loses the integration.

Stubbornness is also rational when you're still learning. If each prototype teaches you something about the commitment (what serves it, what fights it, what you misunderstood) then persistence is buying information. The failures are curriculum.

And it's rational when external pull is increasing: playtesters responding more strongly, market interest growing, collaborators arriving unbidden. These signals mean the commitment is entering compounding territory. Don't pivot when the evidence says continue.

Stubbornness becomes irrational when it stops buying information or leverage. This happens more often than people admit, because "I'm trying" *feels* like a signal but isn't one. The actual signals are forms of *pull*, external validation that reduces your reliance on self-assessment. If prototype after prototype fails to move playtesters, if the commitment has been tested repeatedly and the feeling never arrives, persistence isn't buying compounding. It's buying time to avoid facing what the evidence already says.

The hardest case is stubbornness that preserves identity rather than serves the game. You told people you're making *this* game. You built your self-concept around *this* commitment. Abandoning it feels like failure, so you continue not because the evidence supports it but because you can't face the alternative. The tell is continuing at high burn without changing exposure or seeking new signals. Same thing, same cost, flat results. You're paying for identity, not probability.

---

5. **Stage-gating the commitment**

The right default isn't "persist" or "pivot." It's *stage-gate*: commit to a sequence of experiments whose outcomes determine the next level of investment.

Stage-gating protects you from both failure modes: quitting before compounding can start, and continuing long after the evidence

says it won’t. Each gate should be defined by leading indicators that are legible to the market, not just to your self-concept. In game development, legible indicators are forms of pull:

Indicator	What it measures	Why it’s legible
Playtester conversion	Do strangers who play the prototype want to play more?	External desire, not your assessment.
Repeat selection	Do collaborators, publishers, or platforms come back?	They chose you with their scarcest resource: attention.
Audience response	Is there growth not driven by one-off virality?	Sustained interest survives the novelty cliff.
Gatekeeper conversion	Do pitches and showcases result in advancement?	These people see dozens of projects; they’re calibrated.

These metrics matter because they change your access to future resources. A rising pull metric means the commitment is working. A flat one after multiple attempts means it isn’t producing compounding, whether the problem is the commitment itself or the execution.

“Flat” doesn’t mean one bad playtest. It means you’ve run multiple experiments, varied your approach, and the signal hasn’t moved outside noise. Operationally: after 3+ prototype iterations shown to 10+ fresh testers each, conversion is still below baseline; after 3+ outreach cycles, stakeholders aren’t returning or escalating; after 3+ submissions with incorporated feedback, advancement rate isn’t improving. The threshold is simple: *has effort bought information that changes the odds?* If yes, continue. If no, the signals are flat.

Before you start, define your gates. Write down what signal you need (be specific: “40% of playtesters ask when they can play again,” not “positive feedback”), by when, and what you’ll do if it doesn’t arrive. A plan that includes “if this doesn’t work by X, I will Y” is stronger than one that relies on vague resolve.

If the indicators are flat but you’re not ready to abandon the commitment entirely, the rational move is a **barbell strategy**: split remaining resources between low-cost maintenance of the

commitment (smaller prototypes, constrained scope, experiments that could produce pull without full production burn) and investment in adjacent paths (other projects, contract work, capabilities that remain useful if this commitment fails). The middle ground, continuing at full burn with flat signals, is the dominated strategy. You're paying maximum cost for a lottery ticket whose odds aren't improving.

Why adjacent paths and not full exit? Because you can re-enter when signals improve. Persistence at low cost isn't giving up. It's refusing to pay a dominated price. A commitment is "expired" only in the sense that continuing at current burn is dominated by alternatives. Make the criteria explicit: runway thresholds, timeline thresholds, pull metrics that must move by a date. Vague suffering is what people later mislabel as "grit." Explicit criteria let you persist intelligently.

---

6. **Moving together beats debating direction**

This happens constantly:

The team disagrees about the commitment. Half want direction A, half want direction B. Both are viable. Weeks of arguing, prototyping both, gathering data, holding meetings. Eventually they pick one, but the losing half isn't truly convinced. The next time the project hits difficulty, the debate reopens.

The better approach is simpler:

Pick one. It doesn't matter which, as long as it's viable. Commit together. Move. If it's wrong, the game will tell you through building, not through debate. And if it's right, you'll have a six-week head start on the team that was still arguing.

Everyone pulling in one direction is worth more than pulling in the right direction. A team aligned on a good commitment will outperform a team fragmented over the perfect one. Coherence beats correctness.

This sounds backwards. Shouldn't you spend more time on the most important decision? But the most important decision isn't *which* commitment to pick. It's whether to commit at all.

Committing fully, stubbornly, together is what makes the commitment work.

7. **What makes a commitment viable**

None of this means any commitment will do. It has to be viable, capable of supporting a game through production:

Criterion	Test	Failure mode
Specific	Does it tell you what to cut?	“An emotional adventure” constrains nothing.
Robust	Can you imagine a core loop that serves it?	Beautiful but unbuildable with your team, timeline, or budget.
True	Would you regret abandoning it?	You’ll defend this for years; if you wouldn’t regret losing it, you won’t fight for it.
Aligned	Can you explain it to the investor, publisher, or marketing lead in terms they accept?	A commitment only you believe in dies at the first gate review.

These aren’t tests to run once. They’re conditions to maintain. A commitment that stops being viable needs reinforcement, or rarely, replacement.

8. **Where commitments come from**

Most real projects aren’t blank canvases. They start with obligations: a pitch that got greenlit, an IP license with requirements, an investor’s thesis, a sequel slot requiring continuity, a platform deal with expectations. These aren’t commitments. They’re constraints. The work is to find a commitment that *satisfies* the constraints and *adds* the creative direction they lack.

Example: the pitch said “co-op survival horror.” That’s a genre tag. The commitment might be: “Two players must depend on each

other in ways that create genuine tension between self-preservation and trust.” This satisfies the pitch, constrains design, and gives you something worth defending.

Commitments can also emerge from a feeling you want to recreate, a mechanic that fascinates you, a convention you want to prove unnecessary, an underserved audience, or simply the intersection of constraints when the constraints themselves leave only one fertile space. Where it came from matters less than whether it’s viable. A commitment born from constraint can be as powerful as one born from vision.

---

9. **Validating a candidate commitment**

Before committing resources, stress-test the candidate:

- **The cutting test.** Imagine cutting half the game. Does the commitment tell you what to cut? If not, it’s too vague.
- **The contradiction test.** Propose a feature that violates the commitment. Does the violation feel like a betrayal? If not, the commitment may not be load-bearing.
- **The team test.** Explain the commitment to someone else and ask them to make a design decision based on it. Do they reach the same conclusion you would?
- **The loop test.** Can you imagine a core loop that serves this commitment? If not, it may be aesthetic wrapper without mechanical substance.
- **The stamina test.** Imagine defending this commitment for two years, after setbacks, in difficult meetings, to skeptical stakeholders. Still worth protecting?

If the candidate fails a test, strengthen it or find another. But don’t spend forever searching. A good commitment, committed to, is worth more than a perfect commitment debated.

---

10. **When the commitment should change**

Sometimes the game tells you the commitment is wrong. Not hard, *wrong*. The mechanics fight it. The loop resists it. The feeling never arrives no matter how you tune. The pull metrics stay flat despite varied attempts.

That’s discovery. But distinguish it from two other things:

Discovery feels like revelation. “We were building toward X, but it was always about Y.” The new commitment explains the past better than the old one. You feel clarity.

Drift feels like relief. “X was hard, but now we don’t have to do the hard thing.” The new commitment is vaguer, easier, less constraining. You feel lighter, not clearer. This is retreat, not learning.

Stage-gate exit feels like data. “We ran the experiments. The signals didn’t move.” Neither revelation nor retreat, just the plan working as designed.

If it’s discovery, update the commitment and realign. The game taught you something.

If it’s drift, resist. The hard thing was probably the right thing.

If it’s a stage-gate exit, execute the plan. Move to the barbell. Shrink the burn. Preserve option value. This isn’t failure. It’s refusing to pay a dominated price when the evidence says your odds haven’t improved.

11. **Quick reference**

Concept	One-line definition
Stickiness	The way commitment compounds decisions, aligns teams, and converts doubt into motion.
Alignment decay	The hidden cost of pivoting: stale mental models, fractured trust, phantom decisions from the old direction.
Rational stubbornness	Persisting because effort is buying compounding, learning, or increasing external pull.
Irrational stubbornness	Persisting because you can't face the alternative, not because evidence supports continuing.
Stage-gate	Experiments whose outcomes determine the next level of investment. Define signal, deadline, and contingency before starting.
Flat signals	Multiple experiments, varied approaches, metrics that haven't moved outside noise.
The barbell	When signals are flat: low-cost maintenance of the commitment + investment in adjacent paths.
Discovery vs. drift	Discovery: the game taught you something (clarity). Drift: you're retreating from difficulty (relief).

---

12. **Related essays**

- *Making the Thing (https://gerolds.github.io/posts/making-the-thing/)* — assumes you have a commitment; this piece is about why keeping one matters.
- *Prototyping the Loop (https://gerolds.github.io/posts/prototyping-the-loop/)* — once you have a commitment, how to find a loop that serves it.

Drafting assistance: Claude. All claims mine; errors my responsibility.