Nuts and Bolts of New Ventures

LazyEarn track, MIT 15.393 OpenCourseWare entrepreneurship course

Original lectures by Joseph Hadzima

Transcript-derived course notes organized by LazyingArt LLC with Video2Book.

Contents

Chapter 1

Introduction: Most Startups Fail; How to Improve Your Odds

Overview. Joseph Hadzima begins this MIT OpenCourseWare course the right way: not with a doctrine, not with a slogan, but with a room full of motives and questions. He first asks why we are here, then what we do not know, then why we would attempt something so painful given such poor odds. Only after that does he permit himself equations, frameworks, and diagrams. These notes follow that sequence closely. They are curated by LazyingArt LLC from the source lecture, and they are meant to read as companion notes for a practical course in venture formation rather than as startup mythology.

1.1 Why We Are Here

The lecture opens with a guide's promise. We are beginning a six-session journey through the nuts and bolts of new ventures, and the first task is to identify the motives that have brought the room together. Hadzima does not sort those motives into pure and impure kinds. He simply lays them out and lets us hear ourselves in them.

Some are here because they have never taken a formal class in entrepreneurship and want to see what the subject even looks like. Some are here because they have encountered some miserable product or process and want to make it less bad. Some have an invention and want to bring it to life. Some want to change the world. Some are drawn by the glamour and money that entrepreneurship still advertises.

The lecture's early tone matters. We are not yet being asked to judge a venture. We are being asked to notice why the possibility of a venture exerts a pull at all.

The Dropbox anecdote gives this opening its first concrete shape. A student, stuck without access to his files while traveling, does not merely complain. He converts irritation into a design imperative: make a storage system that does not "suck." The lecture uses this story carefully. It is not yet a theory of ventures. It is an example of how dissatisfaction, if pursued, can become the beginning of disciplined entrepreneurial thought.

1.2 The Entrepreneur's First Questions

From motive, the lecture widens into uncertainty. Hadzima says that he will give us guideposts and asks us to carry his questions from session to session. The effect is cumulative. We do not begin with a neat framework because the subject does not first arrive to us neatly.

1. How do we start?
2. What, exactly, do we do first?
3. What problem are we solving?
4. Does the proposed solution really answer that problem?
5. Who cares about the idea, specifically, and who is the customer?
6. How will the venture make money?
7. How long will it take to get to market, what will it cost, and what resources will be required?
8. What legal entity, if any, should be formed?
9. How can the idea be protected, if protection is the right objective?
10. Will cofounders be needed, and how will the relationship be structured?
11. How will negotiation enter, with employees, consultants, customers, advisors, and investors?
12. How do we discover what we do not know?

The lecture then adds two important qualifications. First, entrepreneurs are always negotiating because they are assembling people and resources they do not yet fully control in order to pursue a vision they do not yet fully own. Second, mistakes are not the scandal. Repeating the same mistake is the scandal. The people teaching the course, he says, have scars, and part of their use is that they may help us recognize recurring potholes before we drive directly into them.

By this point the chapter has already adopted the lecture's rhythm. We begin from personal motives, but we are not left in autobiography. We are pushed outward into operational questions. The subject becomes practical before it becomes formal.

1.3 Why Do This If Most Startups Fail?

At exactly the point where the lecture might have become merely energizing, Hadzima turns and creates resistance. Why do this at all? Why enter a line of work defined by difficulty, embarrassment, and attrition?

The success image is easy to summon: fame, fortune, public recognition. The lecture then opposes to that image the much harsher testimony of startup life. The Nvidia anecdote serves this purpose. The founder, looking back across decades, says in effect that if he had fully known what the road required, he might not have taken it. The point is not anti-entrepreneurial. It is corrective. We are meant to feel the cost before the course offers any technique for improving our odds.

The lecture gives the base rates in blunt numerical form:

$$\Pr(\text{startup success}) \approx 0.10 \approx 10\%, \tag{1.1}$$

$$\Pr(\text{startup failure}) \approx 0.90 \approx 90\%, \tag{1.2}$$

$$\Pr(\text{top-tier VC success per investment}) \approx \frac{3}{10} = 0.300. \tag{1.3}$$

The baseball analogy is part of the argument, not a decorative aside. A venture capitalist succeeding three times out of ten is likened to a .300 hitter: someone who fails often and is still a star. The analogy has a pedagogical use. It takes a discouraging ratio and translates it into a more familiar domain where high failure and excellence already coexist.

One slide, cited in the lecture, sharpens the point further by attributing most failure to poor planning and poor experience. We should not over-read the exact number attached to that slide, but the argumentative purpose is clear. Hadzima is preparing the answer to the question he has just raised. If failure were purely random, there would be little reason for a course like this. If planning and preparation matter, then a course may genuinely improve the odds.

That is why he then introduces the local MIT counterweight. The Kauffman-style numbers are not given as a proof that ventures succeed here routinely; they are given as evidence that when ventures launched in this environment do succeed, they may succeed at very large scale:

$$\frac{26{,}000}{120{,}000} \approx 0.217 \approx \frac{1}{5}. \tag{1.4}$$

On the lecture's telling, that means roughly one active MIT-founded company for every five living alumni in the period covered by the report. The same framing attributes annual revenues on the order of $2 trillion to those companies and says that, if aggregated, they would amount to roughly the eleventh largest economy in the world.

So the lecture does not deny the ugliness of the odds. It accepts them. Then it narrows the ambition. The course will not make venture formation safe. It will try to improve the probability of landing on the successful side of an intrinsically asymmetric game.

1.4 Why This Course, and Who Is In The Room?

After the odds have been stated baldly, the lecture repairs confidence by explaining what kind of course this is. It is not about theory. It is about doing. The speakers are not academics who study entrepreneurship from a distance; they are people who either are doing the work or have already done it.

The course's origin story reinforces the same point. Students effectively demanded a January course on how to start a company. Hadzima, short on time, assembled practitioners and asked them for the most useful possible format: tell the room the three or four or five things you wish somebody had told you when you started. That origin story is itself a small entrepreneurial parable. Demand appeared first; the product was assembled in response.

Then the lecture narrows from the course as a whole to the room immediately in front of it. This matters because the first equation in the lecture does not appear on a pure mathematics slide. It appears on a slide about audience heterogeneity: students and non-student participants, different backgrounds, different interest groups, different expectations.

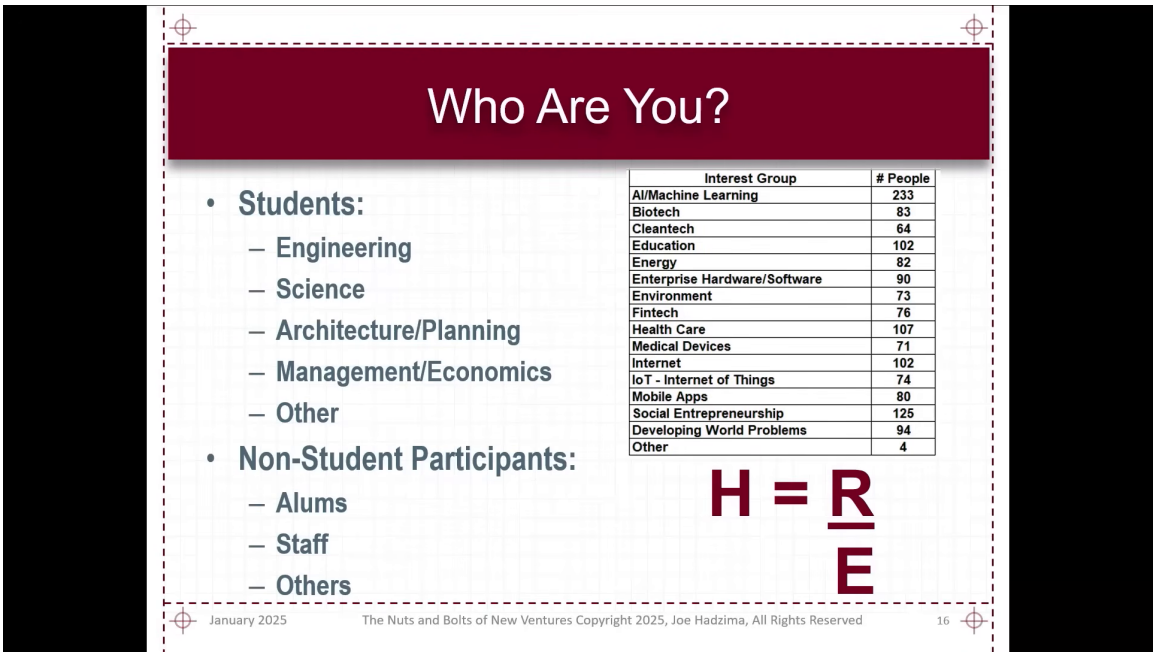

Figure 1.1: Audience mix and the course happiness equation.

The slide is really a room map. The audience breakdown at left and the small interest-group table at right explain why the equation belongs here. The lecture needs a compact rule for thinking about how one course can satisfy people who arrive with different prior knowledge and different hopes.

The slide writes the quotient in a stylized stacked form. We standardize it in ordinary mathematical notation:

$$H = \frac{R}{E}. \tag{1.5}$$

Definition 1.1 (The Happiness Equation)**.** In the lecture's notation,

$$H = \frac{R}{E},$$

where H denotes happiness, R denotes reality, and E denotes expectations.

Hadzima then states the course's objective in deliberately mathematical shorthand. The goal, he says, is to maximize H. We should hear that as lecture rhetoric rather than as a literal optimization problem, but the rhetoric is useful because it forces us to ask what in the ratio can be responsibly changed.

1.4.1 Question & Answer

Question. What are R and E in $H = \frac{R}{E}$?

Answer. R is reality and E is expectations. The lecture briefly records a revealing false start: some Sloan students suggest that R must be revenue and E must be expenses. Hadzima accepts that reading as a subset and then enlarges it. The equation is not an accounting identity. It is a more general rule for disappointment and satisfaction.

Once we write it cleanly, the lecture's verbal logic becomes immediate:

$$H = \frac{R}{E}, \tag{1.6}$$

$$H > 1 \iff R > E, \tag{1.7}$$

$$H = 1 \iff R = E, \tag{1.8}$$

$$H < 1 \iff R < E. \tag{1.9}$$

So if we underpromise and overdeliver, then $R > E$, and the ratio exceeds one. If we overpromise and underdeliver, then $R < E$, and the ratio falls below one. The lecture's little classroom equation is therefore already doing more than managing expectations for a single course. It is establishing a general rule about promises, delivery, and emotional aftermath.

Remark 1.2. The lecture jokes that if we could drive expectations to zero, we might produce infinite happiness. We should not literalize the joke into a usable model. The mathematical content we actually need is comparative: if expectations rise while reality does not, the quotient worsens.

This is also why the equation returns at the end of the lecture, after the discussion of pitch decks and Theranos. It is not a throwaway line. It is one of the lecture's organizing devices.

1.5 Frameworks for Judging a Venture

Now that the room has been calibrated, Hadzima begins compressing entrepreneurship into reusable filters. The first filter is almost embarrassingly simple, and that is part of its strength.

A venture must create value. If it creates no value, it has no subject matter. But creation alone is not enough. It must also capture enough of that value to continue. In the for-profit setting, that capture eventually takes the form of profit and cash flow. In the not-for-profit setting, it may take the form of enough attention, support, and incoming resources to keep operating. The structural distinction is important. Hadzima is careful not to reduce entrepreneurship to one legal or financial form.

The next compression is the "three whys," followed by a fourth question that appears only after the first three have begun to work:

1. **Why this?**
2. **Why now?**
3. **Why this team?**
4. **Why won't this work?**

Why this? The lecture's strongest illustration is Amy Smith's water-incubator pitch. Its power begins with scale: 1.9 billion people lacking access to clean water. Only after that scale has been felt does the mechanism appear. Biological contamination testing requires incubation. Existing incubators require electricity. The affected population frequently lacks electricity. Therefore a non-electric incubator attacks a real bottleneck. The structure is exact: large problem, narrow obstacle, specific bridge.

Why now? The FAA example gives this question its form. A technical solution may be genuinely good and still not be timely. But if institutional demand appears at the same moment that the solution becomes available, then timing and opportunity converge. The lecture's point is severe: five years early and five years late may both be failures even when the idea itself is sound.

Why this team? Prior experience can help, but it is not magic. Hadzima's Encore Computer example makes that clear. A team may have astonishing prior credentials and still fail. Prior experience can attract funding; it does not guarantee strategic coherence or executional success. For early founders without that prior experience, the question becomes sharper: what model, what judgment, what surrounding talent gives us reason to believe that this team is the right team?

Why won't this work? This question does not replace the first three. It arrives after them. Hadzima even frames it in sales language: by the time a prospect begins seriously probing limitations, interest already exists. That is why this fourth question belongs to de-risking. We are now asking not whether the venture sounds attractive, but whether its failure modes are understood well enough that we can plan around them.

By this point the lecture has moved from personal motive to evaluative machinery. But it has not yet become abstract. Each filter is anchored in a story, and each story exists to sharpen the filter rather than to replace it.

1.6 Ideas, Execution, Timing, and People

Hadzima says he dislikes frameworks, and then, sensibly, gives us one. The next compression is the lecture's central four-part model of venture success: ideas, execution, timing, and people. The lecture does something subtle here. It first presents the four-part structure in a general, almost static form, and only afterward begins to make one part of it dynamic.

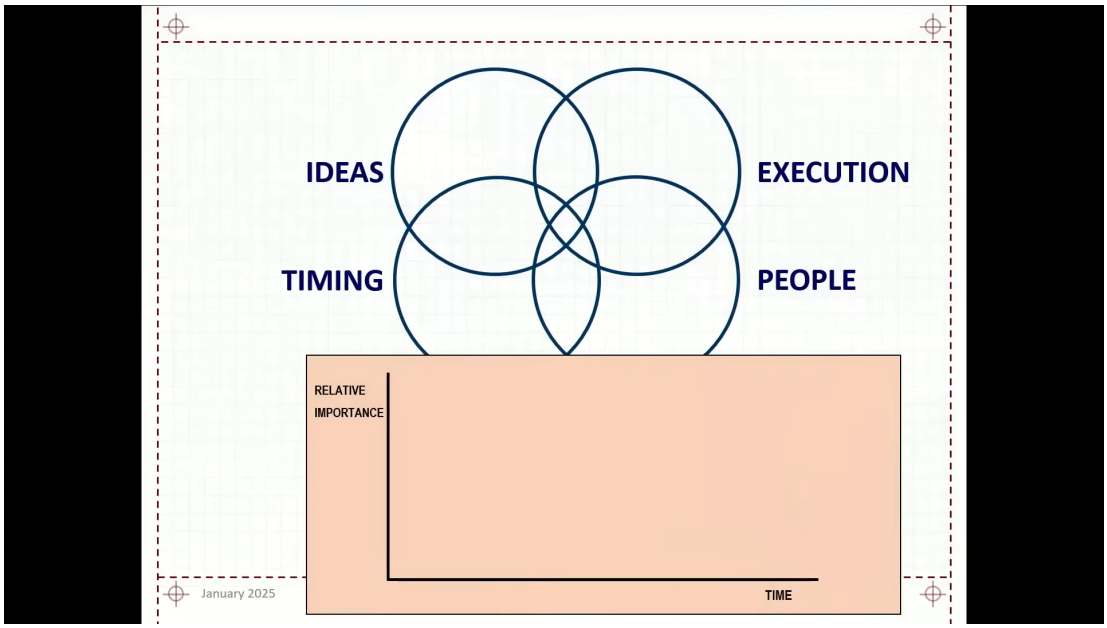

Figure 1.2: Four startup factors above an empty importance-time graph.

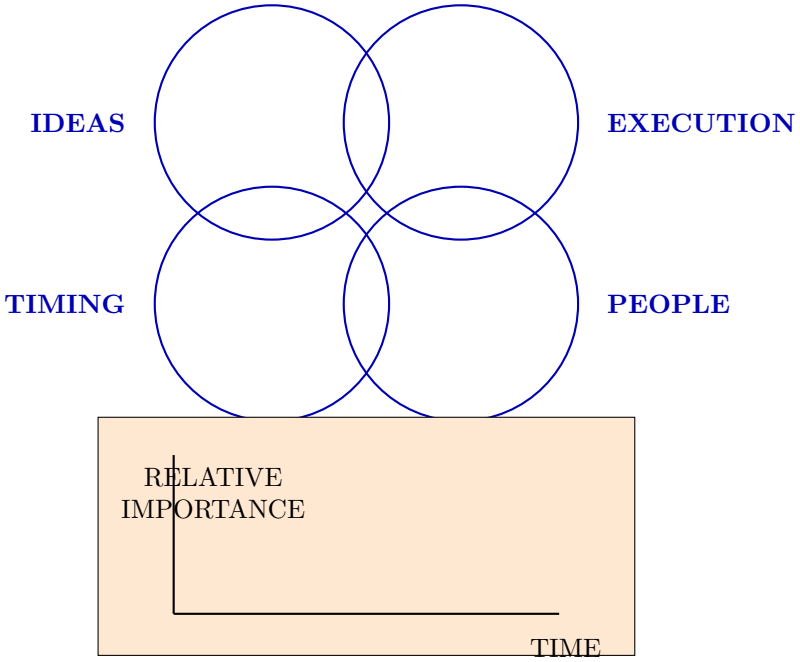

Figure 1.3: Pocket redraw of the four-factor framework before the lower graph is interpreted.

The upper panel says simultaneous alignment. The lower panel says, in effect, not yet. Although the lecture verbally introduces it as a "curve graph," the visible slide at this stage is only a scaffold: axes, labels, and empty space. That is exactly right narratively. The lower panel is being reserved for a specific problem that has not yet been solved.

1.6.1 Ideas

The lecture is impatient with the mere existence of ideas. Around MIT, ideas are everywhere. The real question is whether an idea is valuable, to whom it is valuable, how much value it creates, and whether that value can be defended long enough to support a venture. An idea that can be copied instantly may still be admirable; it may even be socially useful. But the lecture insists that admiration and venture viability are not the same thing.

1.6.2 Execution

Execution is where the lecture becomes particularly concrete. Hadzima uses examples of deals he passed on not to boast about discernment, but to show how execution uncertainty feels before an outcome is known.

Zipcar looked conceptually important. Renting a car by the hour is a genuine break in the consumption model of transportation. But the operational burden, especially around parking and city-by-city implementation, was for him obscure enough to justify passing. The precursor-to-eBay story works the same way. The market concept was visible early, but the trust and fraud machinery required for large-scale execution was not yet persuasive.

This is why the lecture pauses over *mens et manus*. Ideas and execution belong together. Yet even that pair is not sufficient.

1.6.3 Timing

Timing is treated with unusual seriousness. Hadzima says plainly that he has lost more money and time by being ahead of the curve than in many other ways. Timing, in this lecture, is not a matter of trendiness. It is a matter of whether the surrounding world has become compatible with the venture.

3D printing is the clean example. The technology was not "new" in the simple sense. What had changed was the supporting environment: materials, sensors, microelectronics, and the expiration of earlier constraints. Prodigy, by contrast, spent heavily to do online shopping before the user environment and browser infrastructure were ready. Fusion makes the point in an even harder form. It does not work in pieces. All of it must work together, and that makes timing inseparable from the completeness of the system.

So when the lecture asks "Why now?", it is not asking for a slogan about urgency. It is asking whether the world and the venture have arrived at the same time.

1.6.4 People

People, Hadzima says, are the single biggest source of failure in most ventures. He first frames this with the Cheshire Cat story: if we do not know where we are going, route choice is meaningless. He then makes the point vivid with the three-founder example. One founder wants the technology to become a standard, perhaps even open source. Another wants a fast-growing company that can go public. A third wants something important but more collegial and less intense. These are not three versions of the same ambition. They are three different journeys mistakenly housed in one venture.

The lecture then adds two more people problems. Sometimes teams are funded before they understand what they are doing, and then the pivot becomes internal rather than customer-facing. And sometimes teams are too certain too early. The E-Ink example belongs here. Not knowing what we do not know is itself a people problem because it determines whether a team is teachable.

1.6.5 Question & Answer

Question. How do technical and business contributions change over time inside a founding team?

Answer. The lecture answers this with a qualitative time-dependent graph. Early on, the technical side has very high relative importance because until something works technically, there is nothing to finance, market, or scale. The business founder, by contrast, may initially be working nights and weekends, perhaps still holding another job, trying to understand the market and the eventual commercial structure. Over time the balance changes. The business side rises in importance because distribution, team building, financing, negotiation, and organization become increasingly central.

The lecture's clean symbolic summary is qualitative rather than quantitative:

$$I_{\text{technical}}(t)\downarrow, \qquad I_{\text{business}}(t)\uparrow . \tag{1.10}$$

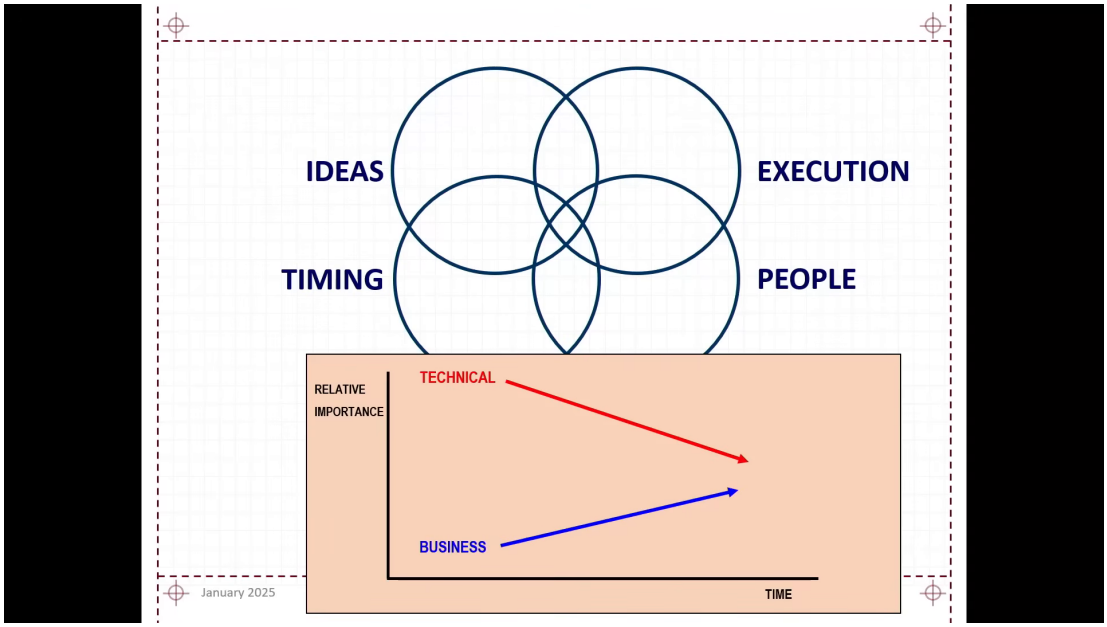

Figure 1.4: Technical importance falls as business importance rises.

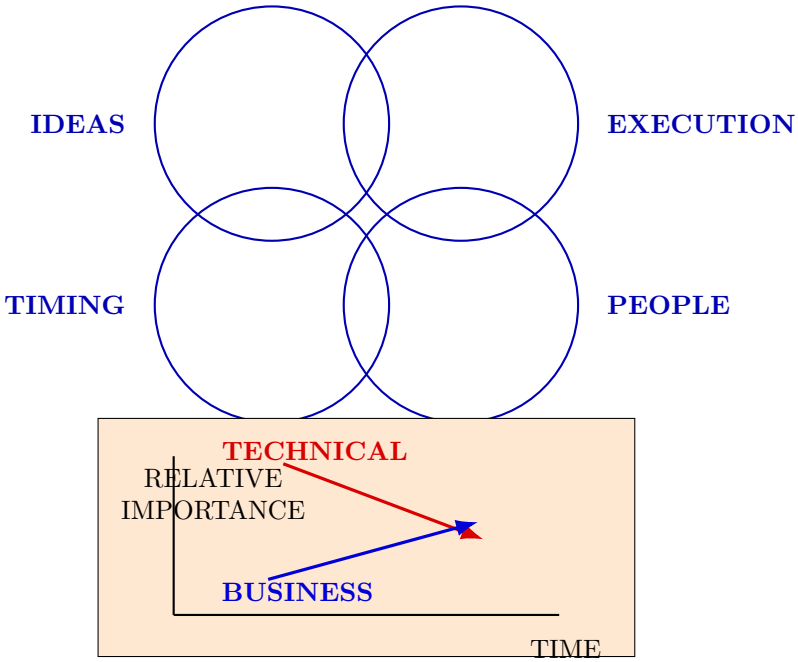

Figure 1.5: Pocket redraw of the lecture's qualitative team-balance graph.

Remark 1.3. The graph is schematic. The arrows are straight, unscaled, and not tied to a marked crossing time. The lecture explicitly says that whether the lines cross is not the point. The point is that the relative balance changes, and with it the conversation about legitimacy, structure, and equity.

Once the graph has opened the team problem, the lecture does not leave the point in diagrammatic form. It immediately turns to actual ventures and asks which of the four components aligned and which failed.

SpeechCo. This is the lecture's success case of alignment. The idea was strong. The timing was good because compute power had finally arrived. Execution improved by restricting the system to a domain of language use rather than trying to solve everything at once. And the people had spent enough time together before launch to form a functioning team.

VideoCo. This case initially looks like an execution failure in the making. The founders seem incapable even of choosing a company name. The payoff is that they were not merely naming the company; they were choosing a strategy. The DIVA/Avid connection reveals that what looked like hesitation had a directional logic.

HIV Co. Here the idea, timing, and scientific credibility were all strong. The weakness was people. Cofounder conflict drained the venture even in the presence of first-rate backers and advisers. This is Hadzima's cleanest case of a venture weakened primarily by internal human structure.

NanoCo. This is the lecture's timing case. NanoFuel had technical promise and environmental appeal, but the market arrived in the wrong price regime. The lecture even gives a back-of-the-

envelope threshold:

$$p_{\text{diesel}} \approx \$0.50/\text{gal}, \tag{1.11}$$

$$p_{\text{neutral}} \approx \$2.00/\text{gal}. \tag{1.12}$$

At launch,

$$p_{\text{diesel}} < p_{\text{neutral}}, \tag{1.13}$$

so the product fell below the lecturer's rough economic-neutral point. The lecture makes the reasoning even more vivid by observing that one could not buy a gallon of distilled water for \$0.50. The argument is intentionally simple, but it is enough. A technically interesting product may fail because the world supplies the wrong price at the wrong time.

1.7 From Vision to Pitch, Without Losing the Substance

Only after all this does the lecture turn to communication. The order is essential. We are not being taught to substitute rhetoric for venture substance. We are being taught how substance is compressed when it must be explained to other people.

Hadzima begins with the mission statement or vision, sometimes also treated as a value proposition. He then immediately insists that the vision must be supported. The lecture's pyramid is a hierarchy of compression.

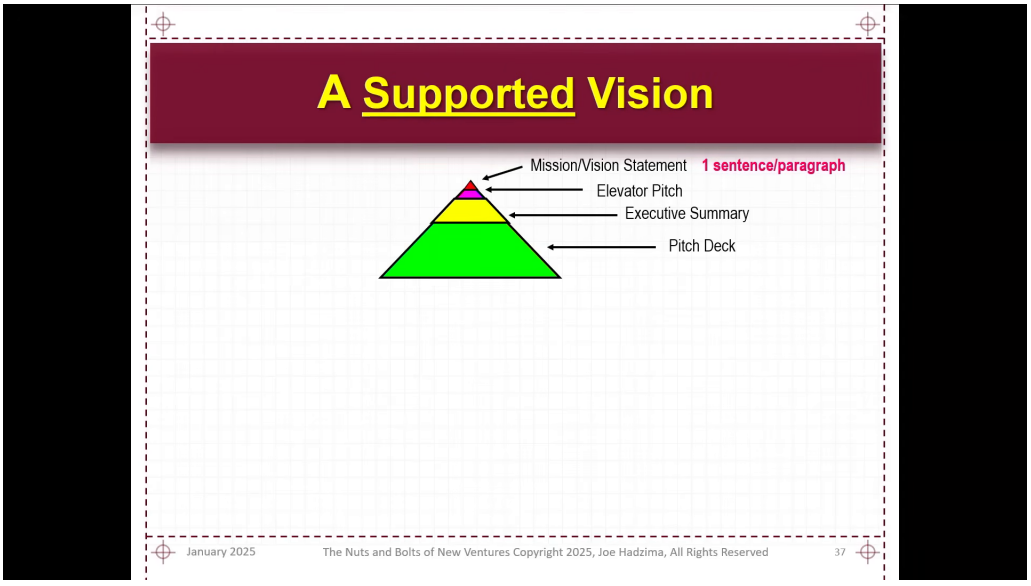

Figure 1.6: A supported vision from mission statement to pitch deck.

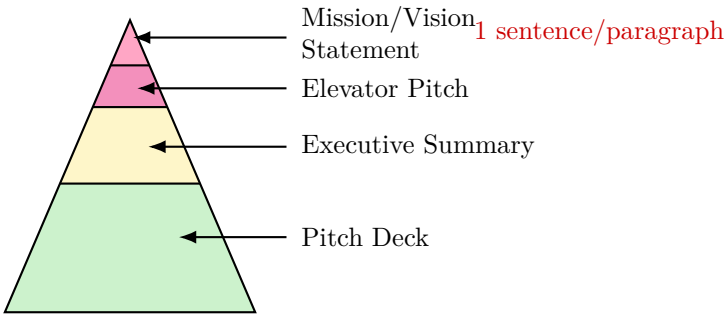

Figure 1.7: Pocket redraw of the supported-vision pyramid.

The original slide is worth keeping because the arrow targets fix the order of the tiers. The apex is the most compressed claim, and the lower tiers are progressively larger supports. One can read the entire device as repeated compression of the same venture at different bandwidths:

$$\text{mission statement} \sim 1 \text{ sentence or paragraph}, \tag{1.14}$$

$$\text{elevator pitch} \sim 30 \text{ seconds}, \tag{1.15}$$

$$\text{executive summary} \sim 1\text{–}4 \text{ pages}, \tag{1.16}$$

$$(\text{slides}, \text{minutes}, \text{font size}) = (10, 20, 30). \tag{1.17}$$

That last line is Guy Kawasaki's 10/20/30 rule: ten slides, twenty minutes, thirty-point font. Hadzima presents it as discipline. Each level asks the same question at a different scale: can the venture still make sense when compressed?

This is why the lecture then moves directly to the Theranos warning. Communication matters, but unsupported communication is not persuasion. It is deferred collapse. The lecture's earlier happiness equation becomes useful again. If reality is fixed while expectations are inflated, the ratio worsens:

$$R \text{ fixed}, \qquad E_2 > E_1 \qquad \Longrightarrow \qquad \frac{R}{E_2} < \frac{R}{E_1}. \tag{1.18}$$

So the lecture's final position is not anti-pitch. It is anti-separation. Sizzle without substance is empty, but substance that cannot be articulated will often fail to recruit the people, capital, and trust required for the venture to exist.

1.8 Summary

We can now see the lecture unfolding in the order Hadzima intended. It begins with motives, expands into operational ignorance, then creates the central tension by insisting on bad odds and painful realities. Only after that tension is fully present does it justify the course as a practical instrument for preparation. The room itself then becomes the site of the first equation, $H = \frac{R}{E}$, and that local rule about reality and expectations becomes a general rule about promises, delivery, and hype.

From there the lecture moves by successive compressions. First: create value and capture value. Then: why this, why now, why this team, and why won't this work? Then: ideas, execution, timing, and people. Finally: mission statement, elevator pitch, executive summary, and pitch deck. Each

compression is followed by examples, because the lecture never allows a framework to float free of lived ventures.

The ending is modest, and that modesty is part of the seriousness of the lecture. Luck remains. Preparation does not abolish it. It changes the odds of being ready when luck matters. That is why the lecture ends not with startup triumphalism but with a practical handoff: break the room into interest groups, meet collaborators, return for logistics, and then go find customers.

Chapter 2

Market Identification and Sales: Finding Your Customer

These notes follow Lecture 2 of MIT 15.393, in which Joseph Hadzima handles the course machinery and then turns the room over to Bob Jones for a sustained discussion of finding customers. The curation is by LazyingArt LLC. No validated mathematical frames survive from this lecture, so every displayed relation and every figure below is a cautious reconstruction from the transcript. The lecture's method is itself part of the lesson: begin with a case so simple that we are tempted to dismiss it, and then keep asking the next honest question until the real business either comes into focus or falls apart.

2.1 Logistics, Teams, and the Hand-Off to Customer-Finding

The lecture opens with logistics, but the logistics are already saying something entrepreneurial. Hadzima speaks about attendance, the written requirement, and the preference that the work be done in teams. One reason is comic and administrative: fewer separate submissions are easier to review. The serious reason is the one the lecture wants us to keep. Teams, he says, have a higher probability of success than sole entrepreneurs, which we may summarize as

$$\Pr(\text{success} \mid \text{team}) > \Pr(\text{success} \mid \text{solo}). \tag{2.1}$$

That inequality is a compact formalization of a spoken claim, not a theorem proved in the room; but it is the first practical relation of the evening, and it explains why the lecture begins with team formation rather than with product genius.

Only after the course structure is in place does Hadzima hand the room to Bob Jones. Jones immediately makes a contract with the audience. If he is going to ask for ninety minutes of attention, he says, he ought to return something worthwhile for that investment. He therefore asks two questions before he teaches anything substantive: who are you, and what do you want? The room turns out to be mixed. Some people are already on a second venture, some are in their first, some are seriously considering one after graduation, and some are simply exploring. The jokes matter here: medication is available for first-time founders; exploratory work can be as unacademic and as useful as springboard diving during IAP. The room is heterogeneous, and Jones wants the argument to be usable across that range.

This is why he makes his anti-academic turn. He mocks the sort of lecture that begins with a sentence about risk aversion, mean-variance utility maximization, and quadratic utility functions, gets dutifully written down, and leaves nobody any wiser. That joke is not the content of the lecture. It is the permission for the method. Jones tells us he will go in the opposite direction: start with examples so simple they are almost trivial, and let the complexity arrive only when the example itself demands it. That is the pedagogical pivot on which the whole chapter turns.

2.2 Start at the End: Money Needed, Customer Value, Customer Count

Jones now gives us the toy model. Suppose, he says, that he flees the overcharged intellectual world, moves to the Berkshires, becomes a grumpy old man in a tiny house, and decides to give guitar lessons. He can imagine the whole thing very easily: beginner chords, tuning, groove, rhythm, more advanced work, even how to play a mixolydian solo that still sounds like the blues. He would need a small studio, a few amplifiers, perhaps a drum machine. Then he stops himself: let us stop right there.

2.2.1 Question & Answer

What mistake did he make? The answer is the first conceptual obstacle of the lecture. We are often quick to ask whether we *could* do a venture. Jones insists that we must first ask whether we *should* do it. The bridge from the first question to the second is money. Most companies die, he says, because they run out of money or because they are on a path where they will never make any.

So he reverses direction and begins at the end of the story. We want to make enough money. That means we need revenues, and those revenues should exceed our costs:

$$R > C. \tag{2.2}$$

Revenues come from customers. So the real problem is not "I have a product idea" but "how many customers do I need, and is that number sane?"

The notation here is a cleaned reconstruction of Jones's spoken arithmetic, not notation visibly written in the room. Let M denote the amount of money we need over a chosen period, and let V_c denote the value of one customer over that same period. Then the required customer count is

$$N = \frac{M}{V_c}. \tag{2.3}$$

In this lecture V_c is not a grand lifetime-value formula. It is simply what one active customer is worth on the same monthly or yearly clock used for M.

Jones gives the numerical skeleton in one line:

$$\frac{1{,}000{,}000}{1{,}000} = 1{,}000. \tag{2.4}$$

If we need a million dollars and each customer is worth a thousand, then we need a thousand customers. Only after that division do we ask the practical question: how are we going to do that?

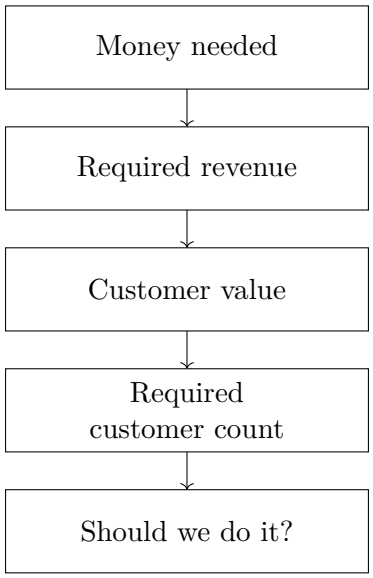

Figure 2.1: Transcript-derived reconstruction of Jones's backward logic. We do not start from enthusiasm for the offering; we start from the money requirement and work backward until the venture either looks feasible or does not.

Now the toy business is run through the arithmetic. Jones wants only a modest living. On a monthly time base,

$$M = \$5{,}000/\text{month}, \tag{2.5}$$

$$V_c = 25\ \$/\text{hour} \times 2\ \text{sessions/month} = 50\ \$/\text{customer/month}, \tag{2.6}$$

$$N = \frac{M}{V_c} = \frac{5{,}000}{50} = 100. \tag{2.7}$$

A hundred customers is not frightening for a large organization. It is a very different number for one grumpy guitarist starting alone. Jones does not hide that discouraging answer. The arithmetic is supposed to hurt a little. That pain is the point. It is cheaper to discover the problem in five lines of arithmetic than in two years of effort and most of a life savings.

2.3 Users, Payers, Champions, and What We Are Really Selling

Having forced the first go/no-go calculation, Jones now asks the question he had postponed: what exactly is broken here, and who pays us to fix it? The default answer is schoolchildren. They are the obvious users. But they are not the obvious buyers. Past a certain point, Jones says quite plainly, the users are not the customers. The parents are the ones who pay. Once we notice that, the economics remain the same but the meaning changes. The child may want fun, status, escape from boredom, or maybe music itself. The parent may be buying activity, discipline, social development, or merely an experiment.

The lecture then performs its first major reset. Jones says, in effect, erase that and start over. What if the relevant customer is not the child but the adult who used to play, used to sing, used to feel cool, and misses all of that? He gives a much richer version of the service: teach a recognizable tune on Tuesday, assemble an ensemble, and let the person play it with others on Thursday. The ticket

Customer definition	Primary payer	Approx. customer value per month	Customers needed for $5,000/month
Schoolchildren taking lessons	Parents	$50	100
Lapsed adults wanting to play again	Adult learner	≈ $400	≈ 12.5
Older adults seeking music and company	Adult learner	≈ $400	about a dozen

Table 2.1: Transcript-derived comparison of three customer definitions for roughly the same underlying teaching activity. The lesson is that the economics change when the payer and the perceived benefit change.

is no longer mere instruction. It is competence regained, identity recovered, and performance made possible. The arithmetic shifts with it:

$$V_c \approx 50 \text{ \$/hour} \times 2 \text{ sessions/week} \times 4 \text{ weeks/month} \tag{2.8}$$

$$\approx 400 \text{ \$/customer/month}, \tag{2.9}$$

$$N \approx \frac{5{,}000}{400} = 12.5. \tag{2.10}$$

Jones rounds the answer in words: about a dozen or so.

He pushes the point once more with the older adults he has seen through volunteer music work. Here the same visible format meets yet another need: social life, purpose, companionship, getting out of the house, perhaps even the possibility of romance. The lesson is not that one should randomly mutate the market. It is that changing the customer definition changes the economics because it changes the benefit being purchased.

2.3.1 Question & Answer

Who wins if I win? Jones slows the room down over this question because it widens the frame. The first answer is customers, and that answer is correct but incomplete. Customers may become a sales force by talking to one another. End users may become champions. Economic buyers may not be the same people as the users. Adjacent businesses may benefit from our success.

His music-store example is the cleanest illustration. A lapsed guitarist goes into a shop to buy strings, stares at the instruments on the wall, and explains that he has lost his chops and has no real setting in which to play. The helpful person in the store says, in effect, I know a guy who can fix that. If the customer starts taking lessons, the store may later sell strings, better instruments, and more equipment. We win; the store wins; the player wins. This is not generic referral marketing. It is a structural point about aligned incentives.

At this stage Jones asks the other crucial question: what are we really selling? One student says romance. Jones half laughs, half accepts, and then clarifies. We are selling comfort, socialization, purpose, collaboration, and the excitement that comes when ensemble performance suddenly clicks. The anecdote about Jones himself walking into a blues jam, playing with strangers, and feeling "magic" happen is not decoration. It is a phenomenology of the benefit. The guitar lesson is the visible product. The invisible product is the life the customer gets to step back into.

2.4 The Non-Negotiable Requirement: Customers

Having extracted a surprising amount from a tiny guitar business, Jones tightens the screws. He now lists the venture questions in a blunter form. We should ask:

1. What is broken?
2. Who specifically wants a solution?
3. Are there enough of those people?
4. Will they spend money to solve the problem?
5. How are they solving it now?
6. Why is our answer compellingly better than the alternatives?
7. Who actually pays?
8. How do we find them and let them know about us?

The questions are easy, he says. The answers are not. But if we do not at least face these questions, we will drive straight into the hole and only discover it when it is too late.

This is also the moment where the lecture pivots toward investor logic. Jones repeats Hadzima's hard framing: most investors lose most of their money. He jokes that perhaps only eight out of ten fail at MIT instead of nine out of ten elsewhere, but the odds are still dreadful. So the investor wants the answer to two questions: how will you make money, and how will I make money? The order matters. If the venture cannot answer the first question, there is no point even asking the second.

From that perspective, there is only one non-negotiable requirement. We must have customers. Everything else can be renegotiated. Jones says this with the force of a room trying to get away with something and being told that it cannot. We do not have to be rich. We do not have to be beautiful. We do not even have to look especially brilliant. If we have customers, the rest of the world begins to reinterpret us. That is the bridge into the first serious case study. Enough of the abstract, he says. Here is what it looks like when smart people get this wrong.

2.5 Customer Discovery by Failure: The Regain Case

The first medical case arrives only after the scaffold is in place, and the sequencing matters. Jones wants us to hear how reasonable a bad plan can sound. He was running the clinical nutrition division of a medical company in California, a division whose economics were excellent:

$$\text{cost per liter} \approx \$3, \qquad \text{sale price per liter} \approx \$70. \tag{2.11}$$

The broader company had been doing roughly two hundred million dollars a year for a long time, but managed care threatened the old profit pool. So Jones went looking for the next revenue source and found, as he says, an opportunity nobody else was pursuing. He immediately adds the correct diagnosis in hindsight: that should have been a red flag.

The field was kidney dialysis. The background numbers looked attractive:

$$\text{400,000 patients}, \qquad \text{2,300 dialysis centers}, \qquad \text{8–12 lb fluid gain between treatments}. \tag{2.12}$$

The contradiction also looked obvious. Dialysis patients often need nutritional supplementation, but the standard supplements were liquids, precisely the thing these patients should not be loading up on. So the company designed a bar-like product, high in what patients needed, low in what they should avoid, and without added fluid. The product was named *Regain*.

The company did the respectable things. It ran a clinical trial even though it was not required. It published the result. The result was good: blood chemistries improved, lives were extended. It then asked opinion leaders, clinicians, and dieticians whether they would recommend the product. The answer, extraordinarily, was yes. For how many patients? All of them. How often? Seven days a week. At what price? Roughly $3 per bar. The arithmetic looked almost too good. Jones says they began to think they had replaced the business they were afraid of losing.

Then the product was launched, and the floor dropped out:

$$\$400{,}000 \text{ forecast} \qquad \text{versus} \qquad \$35{,}000 \text{ actual.} \tag{2.13}$$

As a cleaned-up ratio,

$$\frac{35{,}000}{400{,}000} = 0.0875 < 0.10. \tag{2.14}$$

Month after month, the product came in at less than ten percent of forecast.

2.5.1 Question & Answer

What did we do wrong? Jones lets the question sit in the room because the answer comes in layers.

1. We interviewed clinicians, but the patients were the buyers.
2. We judged taste with the wrong palate.
3. We mistook clinical need for willingness to pay.
4. We misunderstood the motivational composition of the market.
5. We forgot that retailers were customers too.

The first two mistakes are concrete. Clinicians did not buy the product and did not eat it. Patients did both. When Jones's team finally spoke to patients, they learned that many did not want the product and could not afford it. The taste problem was equally sobering. In kidney failure the palate shifts, and a heavily protein-based product can taste like spoiled meat. The team had asked the wrong people how it tasted.

The third and fourth mistakes cut deeper. Jones says the real question should have been: who gets kidney failure? The transcript gives a rough decomposition,

$$33\% + 44\% = 77\% \approx \frac{3}{4}, \tag{2.15}$$

which he interprets as follows: about three quarters of the customer base had arrived there through long-term mismanagement of diabetes or high blood pressure. His conclusion is brutal but plain. A great many of these patients had not taken care of themselves before. Many were not going to start now. Clinical benefit was real, but motivation to spend was weak.

The final mistake is the one the lecture keeps widening toward. Patients still had to buy the product somewhere. Jones's company understood hospitals; it did not understand retail pharmacy.

That ignorance was expensive. A pharmacy manager has only seconds to hear a pitch. A shelf is a monetized asset. Remove a known product, insert an unknown one, and the retailer wants compensation. That is the logic of slotting fees, and Jones gives the number that made the lesson unforgettable:

$$\$1{,}000{,}000 \text{ per quarter for CVS}, \tag{2.16}$$

before guaranteed sales, refunds, and chain-wide stocking requirements were even fully counted.

The failure therefore was not one error but a stack of errors, each caused by a different bad definition of customer. Jones later found a way to reach the roughly twenty-five percent of the market that had not arrived through self-neglect, and the broader company recovered. But the lesson of the year is not rescue. It is diagnosis. The consequences of not thinking through the customer system got worse and worse.

2.6 Rebuilding the Business Around the Real Customer: Nighttime Diabetes

With characteristic dark humor, Jones says that entrepreneurship may be a form of mental illness, because after all this he went off and did more of it. The next medically grounded venture is the constructive answer to the Regain failure. This time the company enters diabetes, and the starting point is not a product gap but a lived problem.

The market is large:

$$10{,}000{,}000 \text{ diagnosed patients}, \qquad 4{,}000{,}000 \text{ insulin users}. \tag{2.17}$$

The practical aim of diabetes management, as Jones states it, is to keep blood glucose from fluctuating wildly. If we denote the overnight glucose level only schematically by $g(t)$, then we should read the next picture as a qualitative shape, not as a clinical graph with calibrated axes. During the day, disciplined patients can spread intake and insulin across small meals and small snacks. At night that discipline is much harder to sustain. Instead people tend to eat a large bedtime snack, take a large insulin bolus, and hope the two will titrate through the night.

The hope fails for a simple reason. The food turns to glucose quickly. The spike is higher, but the support does not last longer. Jones identifies the dangerous interval as

$$2 \text{ a.m.} \le t \le 6 \text{ a.m.}, \tag{2.18}$$

the period in which food support has run out but insulin may still be working.

At this point the lecture changes pace. Jones explains focus groups, then rebels against the standard focus-group script. After too much science and too much intellectual talk, he bursts out and asks what the customers are actually afraid of. The answer is the first true emotional discovery of the case: some people no longer sleep in a bed because they are afraid they will get too comfortable and die in their sleep. That is the sentence around which the business is rebuilt.

So the company creates a product whose components convert to glucose at different rates and last all night. Jones says that, being marketers at heart, they called it time-release glucose. Later in the lecture he refers to the resulting business by product name, *Night Bite*. This time, he adds pointedly, it also tasted good.

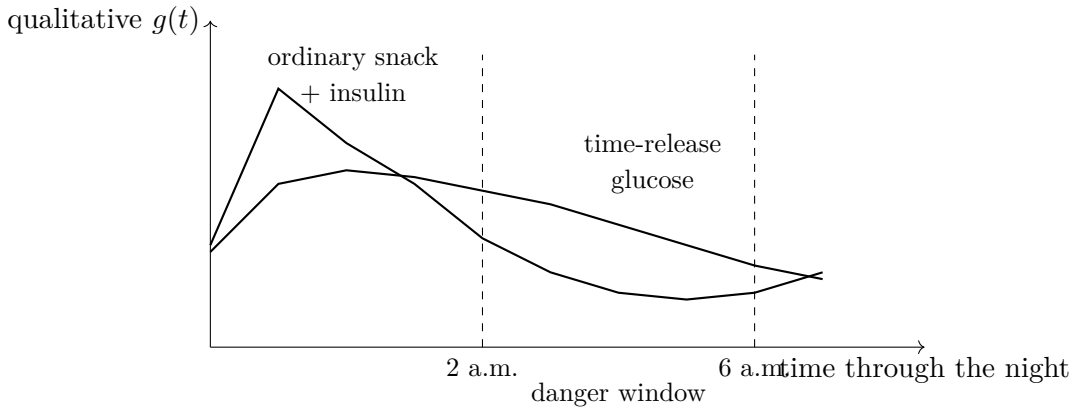

Figure 2.2: Transcript-derived schematic of the overnight problem. The ordinary bedtime snack gives a sharper early spike without widening the safe interval, whereas the redesigned product is meant to spread glucose availability further into the night.

2.6.1 Question & Answer

Why would putting "for diabetes" on the label kill the business? Jones raises this as a genuine puzzle because he originally fought to be able to say exactly that on the package. The customers corrected him. It is nobody's business, they say, that I have diabetes. I am not "a diabetic" in the public sense; I am a banker, a lawyer, an entrepreneur, an active professional. If I pull out something that looks like medicine during a meeting, people will think I am sick, frail, and professionally diminished. That is a career-limiting signal.

So the product has to look like something an elite athlete might use, not like medication. The chemistry can be medical; the public identity cannot. That is why the company gives it a non-medical name and why Jones later notes, with some satisfaction, that marathon runners ended up buying it too.

The redesign also becomes more precise because the company now talks not only to clinicians but to parents and patients. The unit is made deliberately small and divisible:

$$\text{100 calories per unit,} \tag{2.19}$$

with scored use that allows 50, 100, or 150 calories depending on the user. A better product has now emerged. But Jones is careful not to stop there. A better product is still only half a business. The next problem is distribution.

2.7 Channels, Free Sales Forces, and Selling Customers to Retailers

Jones now slows the story down and turns to what he calls the nuts and bolts. The first discovery is that channel changes perceived value. In a focus group the company puts paper in front of users and asks two questions: where would you expect to buy a product like this, and how much would you expect to pay? The answers sort into two piles:

$$\text{\$0.49 in a grocery store,} \qquad \text{\$1.25 in a pharmacy.} \tag{2.20}$$

Same product, different channel, different value. Jones says, dryly, call me silly, I like pharmacies.

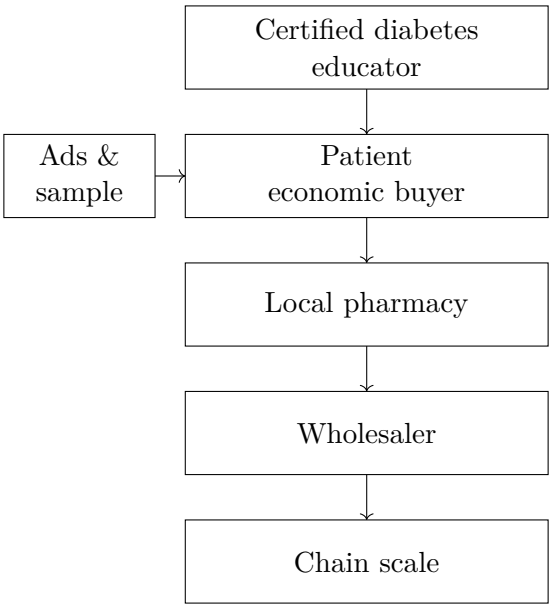

Figure 2.3: Transcript-derived reconstruction of the Night Bite customer system. The patient is the economic buyer, but the educator, the pharmacy, and the wholesaler each carry distinct incentives. The venture only scales when those roles are aligned rather than collapsed into one vague "customer."

The next discovery is organizational. A business like this does not sell to "the market" in one undifferentiated move. It sells through a customer system. The crucial intermediary turns out to be the certified diabetes educator. These professionals can literally be found by zip code through their association. More importantly, they have a professional frustration: patient compliance. They tell patients what to do, and the patients do not do it. If Jones can give them something patients will actually use and thank them for recommending, he has found people who win if he wins.

Several thousand educators eventually become a kind of unpaid sales force. Jones is careful about why. They are not unpaid because they are fools; they are unpaid because they are healers, and what they want is not commission but something their patients will comply with. The educators can then recommend the product to patients and even help pressure local pharmacies to stock it.

The company first sells directly, not because that is the intended endpoint but because Jones wants information. He wants reorder frequency so he can infer usage. Advertising generates inbound calls, and the numbers become a kind of live demand function:

$$100,\ 200,\ 300,\ 500 \text{ calls/day}. \tag{2.21}$$

At the same time, long-run usage turns out to be strikingly strong:

$$\text{reuse rate} > 300 \text{ occasions/year}. \tag{2.22}$$

That is not a market-research fantasy. It is behavioral evidence.

Jones also gives a small methodological lesson from inside sales. Closed-ended questions stop the conversation too quickly. "Did you like it?" can end with "no." An open-ended question such as "What did you think?" keeps the door open long enough to discover that taste may have been imperfect, efficacy was good, and the customer wants to reorder anyway.

The pharmacy-entry tactic is equally concrete. A store manager may have a purchase amount below which corporate approval is unnecessary. Jones guesses that threshold is about one hundred dollars and prices the starter kit just below it:

$$\$99.95 < \$100. \tag{2.23}$$

The risk is removed: if it does not sell, throw it away and get a refund. Then the company calls its existing customers and tells them where the product is now stocked. When the first batch sells out, the pharmacy itself becomes interested. The phrase Jones uses is exact and worth preserving: he bought a bunch of customers and then sold those customers to the retailers.

The same logic is pushed upstream. A large pharmacy may carry about eight thousand SKUs and does not want eight thousand invoices. So the company gets into wholesalers by seeding nearby pharmacies, giving the wholesaler its normal margin, and asking only to be entered into the system before reorder demand appears. Once that is done, scale becomes possible.

The planogram meeting at a major chain is the final demonstration. Jones and his colleague go into a fifteen-minute meeting with no money to pay slotting fees. They spend ten minutes saying who they are and why the product is good, and then open a briefcase full of printouts showing ten thousand existing diabetes customers in the states where the chain operates. Those customers all have one question: which pharmacy carries the product? How, they ask the buyer, would you like us to answer them? Jones then says nothing. He invokes the negotiating rule of thumb that the next person who speaks loses. After the silence, the buyer says yes. What won the meeting was not persuasion in the abstract. It was customers.

2.8 Segment by Motivation, Not Demographics

By the end of the Night Bite story the lecture has earned the right to become doctrinal. Jones turns back to the room and asks what he should have learned from all these experiences.

2.8.1 Question & Answer

The audience answers, and Jones sharpens the answers as he goes. The resulting principles are the real retrospective of the lecture:

1. Talk to your customers.
2. Let the customer definition evolve as the business evolves.
3. Reinvest early effort and money in better market knowledge.
4. Treat distribution as part of the venture, not as an afterthought.
5. Figure out what people are really buying.

That last point becomes a short theory of benefits. Nobody specifically wants the product. People want what the product does. Jones's example is vision: we may hire glasses, contact lenses, or a LASIK surgeon, but what we really want is to see better. The same principle reaches backward through the entire lecture. Nobody wanted "guitar lessons" in the narrow sense. They wanted social life, competence, romance, purpose, performance, collaboration. Nobody wanted "glucose" in the narrow sense. They wanted safety, sleep, and the ability to function without public stigma.

Jones then asks what market segmentation really means. Conventional marketers segment by geography, income, or gender. He argues that the entrepreneur, especially the undercapitalized entrepreneur, must segment by motivation. The transcript's cleanest example is the product that prevents a second heart attack. The person who has just survived the first one sits at the top of the motivation pile. Nearby family members may listen next. Social peers may pay some attention after that. The broad public is the low-motivation block at the bottom:

$$\text{most motivated} > \text{related observers} > \text{adjacent social group} > \text{general public}. \tag{2.24}$$

That ordering is qualitative, not measured. Its meaning is strategic. A startup cannot afford to spend scarce resources trying to convert the least motivated part of the world. Jones calls that big bottom block a graveyard.

The comparison between Regain and Night Bite makes the principle vivid. Regain improved clinical markers, but the patient felt no immediate difference. Every six months the doctor might say, nice going, your blood chemistries are better. Night Bite produced immediate morning feedback: I think it worked. In Jones's joking form, the user did not wake up dead. The line is facetious, but the point is exact. Motivation rises when the benefit is not only real but felt.

He then compresses his own view of marketing almost to absurdity. After acknowledging all the sophisticated mechanics, he says he can save us two years of graduate school with two bullet points, and the two bullet points are the same: find out what your customers want; then find out what your customers want. The repetition is not carelessness. It is emphasis.

The very last quick story before the closing homework lands the lesson one final time. Jones says a group of them once invented a sleep-support product and, in a foolish top-down move, decided the market was working moms. They assembled a huge list. They heard nothing. Then a few repeat purchasers emerged, and when Jones called them personally and asked why they were buying, every answer turned out to be some version of the same thing: I run marathons, I need better recovery, better sleep helps me recover, and therefore your product helps me perform. Jones thought he was selling sensitivity and care. He was selling to warriors. He says flatly that he would have died of old age before learning that from a sequence of Google ads. Talk to them.

That is why the lecture ends by narrowing, not widening. The next session will be about pitching. The homework is to pitch the business in three minutes, then in thirty seconds, then, if possible, in three sentences and about ten seconds. The compression is not a new topic. It is the compressed form of everything in this lecture.

2.9 Summary

The lecture begins with attendance codes and team formation, but its real structure is mathematical and strategic. We begin at the end of the story, ask how much money we need, estimate what one customer is worth, and derive how many customers are required. Then we ask the harder questions: who is the user, who is the payer, who is the champion, who else wins if we win, and what benefit is actually being purchased?

The Regain case shows how completely a business can fail even when the clinical result is real and the arithmetic looks encouraging. The Night Bite case shows how the same kind of medically grounded business can succeed once fear, identity, channel, and reorder behavior are taken seriously. The final principle is severe and practical: do not aim at everybody. Find the people who want the offering most, learn what they are really buying, and build the venture around them.

Chapter 3

How Are You Going to Make Money? The Business/Venture Model

This lecture begins with a clean handoff. Bob Jones had already shown, in the preceding session, that customer choice and pricing choice are not afterthoughts; they are already the beginning of a business model. Rich Kivel now takes the next step in Joseph Hadzima's MIT OpenCourseWare course and asks the harder question: once we have an invention, a service, or a technology, how does it become revenue, investor return, and finally a sustainable business? The lecture answers this in a deliberate order. We are first given an operating story, then a definition, then a set of recognizable business-model types, then two major platform examples, and only after that a return to the industry-specific pressure of med tech. The notes should therefore unfold as the lecture unfolds: step by step, with each example motivating the next.

3.1 From Customer Discovery to Business Model

Kivel opens by making the scope of the evening unmistakable. We are no longer only asking who the customer is. We are asking how a company makes money, how it brings a product to market, and how it eventually produces returns for investors or strategic partners. In the lecture this is presented as the "nitty-gritty" of company building, and that phrase matters. The business model is not a slogan placed after the invention. It is the operating structure that carries the invention into economic life.

A compact reconstruction of the lecture's practical target is

$$\text{technology} \rightarrow \text{distribution} \rightarrow \text{adoption} \rightarrow \text{revenue} \rightarrow \text{sustainable business.} \tag{3.1}$$

Kivel does not write this in symbols, but it is the underlying motion of the lecture. A technology by itself is inert from the standpoint of a company. Between invention and revenue there lie channels, buyers, behavior, and scale.

His autobiographical setup is also part of the argument. He presents himself as someone who spent the first part of his career on the operating side of companies: sales, marketing, partnerships, then CEO work, then investing. This is why the lecture does not begin with a framework diagram. It begins with operating experience. We are meant to see that a business model is discovered, stressed, and revised in contact with actual buyers.

3.2 MolecularWare and the First Operating Model

The first sustained example is the MIT spinout MolecularWare, a bioinformatics company from the era of the Human Genome Project. The technical scene is vividly drawn. A lab is full of autonomous devices: scanners, robots, micro-arrayers, barcode refrigerators, each with its own computer and its own software. The consequence is not elegance but spreadsheet chaos. Scientists and engineers become the unwilling clerks of disconnected systems.

Kivel's first important sequence comes here:

$$\text{capture} \rightarrow \text{visualize} \rightarrow \text{data mine} \rightarrow \text{create value.} \tag{3.2}$$

This deserves to be preserved because it is the lecture's first clean reduction of a business problem to a usable structure. One first captures information from the instruments. One then visualizes it. Once the information is visible, one can interpret overexpression, underexpression, patterns, or anomalies. Only then does the downstream problem of data mining become meaningful, and only there does value begin to emerge in the commercial sense.

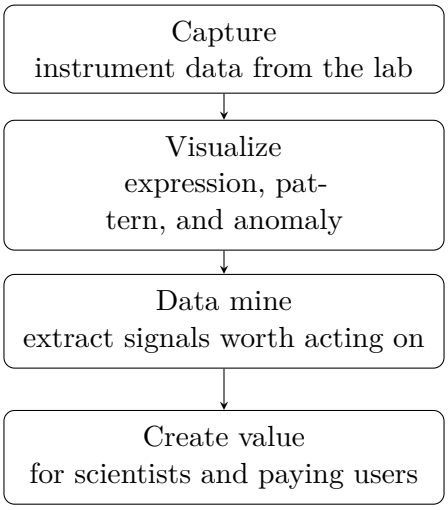

Figure 3.1: Transcript-derived reconstruction of the MolecularWare value chain.

Kivel then makes a point that is easy to lose if we over-academicize the story. He had no background in biotech or pharma. He jokes that he did not know how even to spell bioinformatics. But the point is serious: what mattered was not disciplinary identity; what mattered was whether the company was creating value out of the invention. That is why he broadens the example from gene-expression labs to drones, satellites, emissions sensing, and other domains. The lecture's refrain is that the value often comes from the data.

From there he moves, with almost no pause, from product structure to channel structure. When he arrived as CEO, MolecularWare had one business model: it sold to academics. That made sense because the founders sold to the people they knew. It did not scale well, and it was not very attractive to the investment community. So the business model had to change:

$$\text{direct to academics} \rightarrow \text{distributors} \rightarrow \text{OEM.} \tag{3.3}$$

The first extension was geographic. Building a sales force in Asia was impossible for a fifteen-person Kendall Square company, so distributors became the second part of the model. The second extension

was even more decisive. The company had strong software and no reason to build instruments. Other firms had strong instruments and weak software. OEM bundling then becomes the elegant move: put the software inside the instrument and let the instrument company sell the complete system.

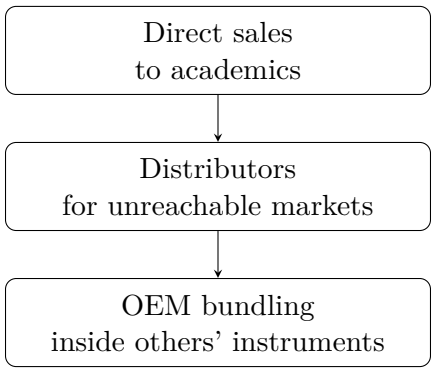

Figure 3.2: Transcript-derived reconstruction of the channel ladder in the MolecularWare case.

3.2.1 Question & Answer

Question. What is OEM, and why did it change the company?

Answer. OEM means original equipment manufacturer. In the lecture it is not merely a vocabulary check. It is the moment at which the company changes how it reaches the buyer. MolecularWare stops trying to sell every installation itself and instead places its software inside another firm's instrument. The manufacturer now bears much of the burden of selling, distribution, and customer reach. MolecularWare receives value when the bundled instrument is sold. The route to market changes, and with it the economics of growth.

Kivel then generalizes from this example to his larger operating experience. He has taken companies

$$8 \to 85 \to 250 \to 1{,}000 \text{ people}, \tag{3.4}$$

and he attributes that kind of growth not only to funding but to a business model flexible enough to support the growth. That claim prepares the ground for the formal definition.

3.3 Definition, Flexibility, and the Nine-Part Canvas

Only after the operating story does the lecture pause to define its subject. This order should be preserved. We are not given a framework and then asked to decorate it with examples. We are shown a company under pressure, and then we are given the abstraction that explains what we have just seen.

Definition 3.1. A business model describes the rationale by which an organization creates, delivers, and captures value.

A faithful symbolic compression of the lecture's definitional core is

$$\text{Business model} = \text{create value} + \text{deliver value} + \text{capture value}. \tag{3.5}$$

Remark 3.2. This displayed equation is a transcript-backed compression rather than a board-written formula. It is useful because it prevents us from treating creation, delivery, and capture as separable afterthoughts.

Kivel then turns to the *Business Model Generation* canvas and immediately refuses to let it become static. The nine components are not equal-sized slices of a permanent pie. Their weight changes with stage, financing, product maturity, and strategic pressure. They blend together. A manager may oversee key resources or key activities; a marketing team may be shaping the value proposition and the segmentation of customers; but the underlying claim is that these pieces must move together.

This is why the orchestra analogy appears. If the pieces of the model move harmoniously, the result is coherent. If one part lags or drifts, there are second- and third-order consequences. Cost structure is the lecturer's favored example. If the cost structure is not designed and revised properly, the firm starts to lose money; support weakens; repeat business fades; the model begins to damage itself.

There is a more general principle here, stated almost explicitly in the lecture: a successful business must be willing to modify its model quickly in response to the market. The model is not the frozen form of success. It is the mechanism by which a company remains able to move.

3.4 A Menu of Canonical Models

From the operating case and the formal definition, Kivel deliberately broadens the field. He reminds the room that none of this should feel exotic. We are all consumers. We live inside business models every day, whether or not we name them.

The lecture names a familiar cluster:

- subscription,
- freemium,
- marketplace,
- advertising,
- direct sales,
- reseller,
- OEM.

But the point is not taxonomy for its own sake. Kivel is training recognition. Years earlier, he says, business-model talk could feel almost arcane: are you going to sell online or retail, are you going to customize like Dell or stock inventory like everyone else? Then ubiquitous connectivity arrives, and suddenly models appear that would scarcely have made sense a short time before.

Subscription is introduced through ordinary life. We used to drive to Blockbuster, rent a movie, bring it home, and return it. There was no meaningful subscription in that world. The deeper shift, as Kivel phrases it, is from owning to experiencing. Freemium extends the same logic in another register: give access now, charge later, keep the user inside the system. Marketplace models let many sellers reach buyers through a shared platform. Advertising transforms attention into revenue rather than requiring equal direct payment from each user.

The older structures remain alive. Direct sales, reseller arrangements, and OEM still matter because the route to market is often the heart of the problem. In pharmaceutical biotech, direct sales forces survive because doctors are busy and physical education of the buyer still matters. The Intel example sharpens the point even further. Intel does not care very much about selling directly to end users. It sells to the manufacturer. In Kivel's improvised phrase, this is almost business-to-manufacturer.

The lecture therefore uses familiar models not to simplify the topic, but to show that the topic has been with us all along.

3.5 Netflix as a Moving Target

The lecture's first large platform example is Netflix. The narrative begins with Blockbuster, because Blockbuster makes the problem concrete. A retail chain owns stores, staffs them, secures them, manages inventory, handles returns, and charges late fees. Netflix appears first not as a studio or global streaming platform but as a company that mails DVDs.

That first move depends on the technological moment. Kivel stresses that at the beginning there was not enough bandwidth to make downloading a feature-length film practical. One could wait hours, consume memory, watch, delete, and repeat the process. The mailed-DVD model was not primitive. It was suited to the state of the world.

This first business model already contains several important features: convenience, no late fees, centralized inventory, and later a subscription layer that lets a household hold several DVDs out at once. The lecture is careful to keep the stages distinct. First there is mailing. Then there is the subscription tiering laid on top of mailing. Then the world changes.

Netflix does not survive by a theatrical leap of reinvention. It survives because it modifies its model before the old model becomes useless. Streaming enters gradually. Later, original content. Later, live events. Later still, small but powerful adjustments in pricing and account rules.

3.5.1 Question & Answer

Question. How did Netflix avoid dying with the DVD business?

Answer. It did not mistake its first successful form for its permanent identity. Mailing DVDs solved the immediate distribution problem in a low-bandwidth world. Streaming solved the next problem. Original programming solved the problem of dependence on others' content. Live events open yet another line of relevance. The company survives not by discarding the old model in panic but by modifying it before the old logic is exhausted.

The lecture then invites arithmetic. In a subscription business the natural first approximation is

$$R \approx p\,N, \tag{3.6}$$

with p the average subscription price and N the subscriber count.

If both price and subscriber count change, the first-order update is

$$R' \approx (p + \Delta p)(N + \Delta N) \tag{3.7}$$

$$= pN + N\Delta p + p\Delta N + \Delta p\,\Delta N. \tag{3.8}$$

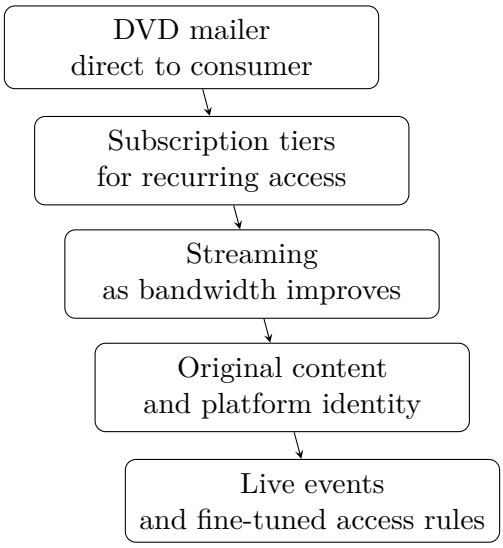

Figure 3.3: Transcript-derived reconstruction of the staged evolution of Netflix's business model.

This algebra is not written in the lecture, but it formalizes exactly the point Kivel is making verbally. In a system with a very large recurring base, tiny movements in price and nontrivial movements in the paying population are economically powerful. The lecture-era scale markers are

$$\Delta N \approx 18{,}000{,}000, \qquad N \approx 300{,}000{,}000. \tag{3.9}$$

So even a small Δp is amplified by a huge N, while even without a large price increase the term $p\Delta N$ can be material. This is why Kivel treats the password-sharing restriction not as a wholly new business model, but as a profitable fine-tuning of an existing one.

He then adds one more layer: personalization. Recommendation, search, and feedback make the platform sticky. The company does not merely stream a library; it uses data about our behavior to make the next interaction more likely.

Remark 3.3. The subscriber counts are clear enough in the transcript to use as scale markers. The later market-cap discussion around Netflix is garbled and should not be sharpened into a precise numerical claim.

3.6 Google, Data, and Predictive Value

If Netflix shows business-model revision across delivery format and content ownership, Google shows something different: the migration of value from visible service into data exhaust and prediction.

The opening move is search. Google begins, in Kivel's telling, as a better answer to a frustrating problem. Search engines existed, but search still felt bad. Humans should not have to master Boolean logic merely to find ordinary things. Better search, then, is already a value proposition.

But the lecture does not let Google remain "just search." Docs and Drive come next. Collaborative documents eliminate version chaos. Teams in different places can work in real time. Google Maps appears free and ordinary, but the analytics behind Maps are where the deeper value lives. The company knows geography, movement, and patterns at a scale that far exceeds the visible app.

A compact reconstruction is

$$\text{search/maps/docs} \rightarrow \text{data exhaust} \rightarrow \text{predictive signals} \rightarrow \text{monetizable value.} \tag{3.10}$$

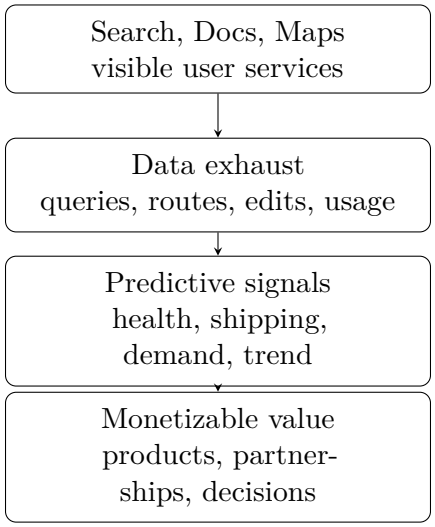

Figure 3.4: Transcript-derived reconstruction of how Google turns visible services into predictive value.

Kivel uses Google's size rhetorically rather than analytically:

$$\text{market cap} \approx \$2.4 \text{ trillion.} \tag{3.11}$$

The force of the number is not in its exact financial interpretation. It is in the contrast with the humble origin point: a company that began by fixing search became one of the most consequential data-and-infrastructure firms in the world.

The health example clarifies the mechanism. Regional search spikes for symptoms become biomarkers, then data points, then public-health signals. Those signals can inform health organizations, clinicians, and even pharmaceutical or over-the-counter product planning. The company that began by indexing web pages now participates in the practical anticipation of disease patterns.

3.6.1 Question & Answer

Question. Why is trailing data nearly useless, and what replaces it?

Answer. Because trailing data arrives after the strategic advantage has mostly evaporated. Quarterly reports tell us what happened. They do not tell us early enough what is about to happen. Kivel's substitute is the leading indicator: shipping activity, search behavior, satellite observation, lot inventory, and any other present signal that predicts future movement.

The lecture's asymmetry can be written as

$$\text{decision value(leading indicator)} > \text{decision value(trailing report).} \tag{3.12}$$

This is not a theorem but an operating principle. If official data tell us how many cars were sold last quarter, that information is already old. If current shipping patterns tell us that fewer cars are

leaving port now, then a prediction about next quarter becomes actionable before the official report exists.

Kivel folds this back into a more general enterprise example. A logistics product may promise

$$\Delta\text{productivity} \approx 15\%. \tag{3.13}$$

That number is not presented as a measured law. It is the sort of concrete business claim by which an offering becomes meaningful to a client. Yet even here the same warning returns: the first model may be simple, but the company must keep becoming more valuable, more relevant, and more sticky.

At this point the lecture pauses, briefly, for questions, and then turns. We have seen two major platform examples. Now Kivel narrows the field again and says, in effect: let us talk about industries.

3.7 Industry-Specific Models: MedTech, Diagnostics, and Pharma

The lecture's pivot back to industry specificity is crucial. Kivel explicitly says that pharmaceutical, diagnostic, and med-tech businesses are not interchangeable. Even inside life sciences, each sector carries a different business model.

Med tech becomes the sharpest example because it compresses technical success and commercial failure into the same space. A lab may achieve proof of concept. A device may work. A device may be safe. Yet the market question remains: who will adopt it, who will prescribe it, who will pay for it, and what service structure surrounds it?

The central tension can be written cleanly as

$$(\text{efficacious} \wedge \text{safe}) \not\Rightarrow \text{adopted}. \tag{3.14}$$

This is the lecture's most important formal reconstruction. It captures exactly the obstacle Kivel is building toward. In med tech, technical success is necessary, but it is not sufficient.

The lecture's validated visual evidence belongs here.

The slide itself is not a workflow. It is a radial framework of simultaneous pressures:

$$\{\text{efficacy},\ \text{data-driven},\ \text{device validation},\ \text{remote service}\}. \tag{3.15}$$

3.7.1 Question & Answer

Question. Why is "safe and effective" still not enough?

Answer. Because the commercial system is not the same as the technical system. One must ask which clinicians adopt first, which specialists are influential, how fast the specialty changes its behavior, whether reimbursement exists, whether the product is direct-to-consumer or prescription-mediated, and whether the service structure supports remote or repeated use. Kivel's examples are pointed. Surgeons may be eager for new instruments; psychiatrists may be slower to adopt new technologies. A patient may see an advertisement, but if the doctor does not prescribe, the product does not sell.

That is the decisive lesson of the MedTech beat. In a simple consumer product the buyer and the user may be the same person. In med tech, the chain from innovation to payment is often multi-layered and institutionally constrained.

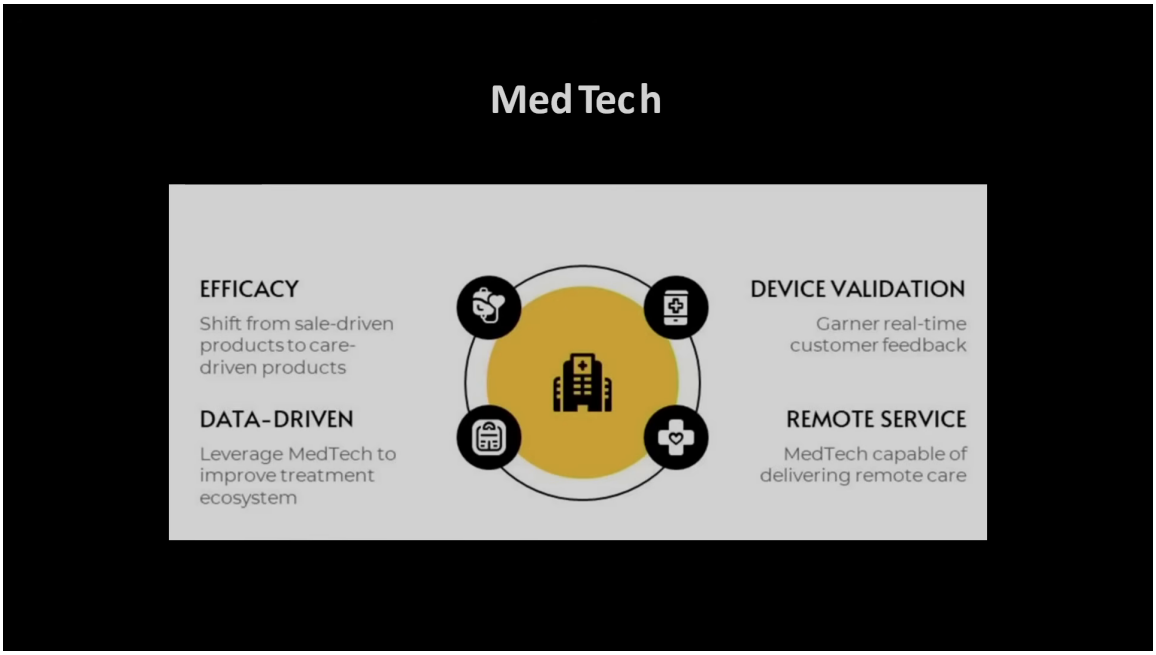

Figure 3.5: MedTech framework slide with efficacy, validation, data, and remote-service dimensions.

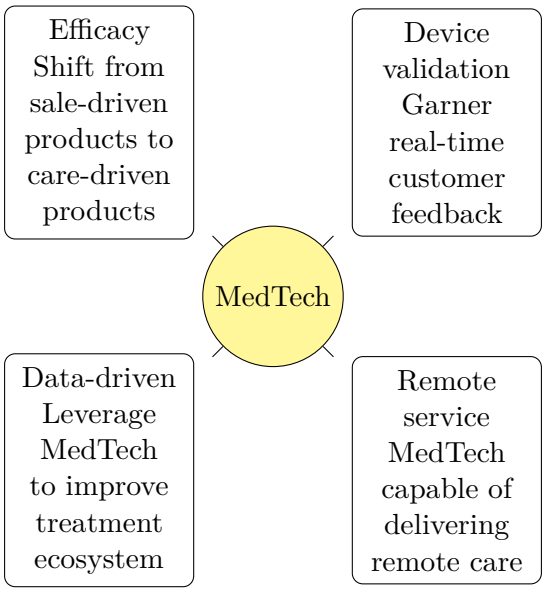

Figure 3.6: Pocket-safe redraw of the MedTech framework visible in the lecture slide.

3.7.2 Switching Costs, Segmentation, and the Late Examples

Once this sector-specific pressure has been made vivid, the lecture widens again. The remaining examples may appear miscellaneous if we flatten them, but they are not miscellaneous. They are a final sweep of heuristics.

First comes switching cost:

$$\text{low switching cost for buyer} \Rightarrow \text{high churn risk for seller.} \tag{3.16}$$

The lecture illustrates this with mobile carriers and gym memberships. Buyers like easy exit. Sellers dislike it because it creates turnover. Yet competitive pressure may force the seller to permit exactly that low-friction movement.

This is why Mint Mobile matters in the lecture. Mint uses T-Mobile's network, but reaches a different segment: bring your own device, no long contract, cheaper monthly access, less bandwidth privilege. The core idea is not merely telecom. It is segmentation. An incumbent model may work well and still miss an entire part of the market.

The same principle reappears in brand ladders. Toyota and Lexus, Nissan and Infiniti, Ralph Lauren and its internal hierarchy of labels: these examples all serve the same lesson. A firm can create adjacent brands or adjacent offers to reach people whom the original identity could not reach. The lecture's live exchange on exact automotive genealogy wanders somewhat, but the conceptual point remains sharp.

Digital commerce then supplies a new cluster of examples. Uber coordinates rides without owning cars. Airbnb coordinates lodging without owning rooms. Rent the Runway is not merely "sharing" but a personalized access model. The replenishment economy removes the need to remember to reorder household goods. The do-it-for-me service economy expands conveniences that once seemed reserved for the wealthy.

Travel becomes the last broad example before the close. What used to require a travel agent can now be built through search or generated through AI prompts. Here again Kivel's point is not that technology changes everything in the abstract. It is that changed technology permits new models of reaching the user.

The AI coda makes the same point one final time from the infrastructure side. A data center is a business. A data center in Norway, using a particular climate advantage, a particular energy profile, a particular funding structure, and a particular compute contract, begins to look like a business model. On the application side, tools such as Otter turn transcription into summary and summary into action items. The visible app is the surface; the model lies in where it sits in the value chain and what workflow it displaces.

The closing advice is therefore earned rather than generic. Do not overcomplicate the model. Understand the market dynamics. Be flexible. The companies that survive are the ones that see the future early enough to alter themselves.

3.8 Summary

This lecture begins with a handoff from customer discovery and ends with a warning about adaptability. Between those two points it builds a coherent sequence. First, a business model is the practical structure by which value is created, delivered, and captured. Second, channel design is not

secondary to product design; the MolecularWare case makes that unmistakable. Third, successful companies survive by modifying their model before the old one collapses. Netflix does this across format, content, and access rules. Google does it by turning visible services into predictive signals. Med tech shows the sharpest tension of all: technical success does not guarantee adoption. The final lesson is therefore sober and operational. We should ask of any company not merely what it makes, but how it reaches the buyer, how it gets paid, and how it will change when its first successful arrangement is no longer enough.

Chapter 4

Legal Issues

These notes follow Joseph Hadzima's MIT OpenCourseWare lecture on legal issues in new ventures, with curation in this volume by LazyingArt LLC. The lecture is practical from the first minute. We are not being invited to become amateur lawyers. We are being taught to notice the places where a venture can lose ownership, lose leverage, or create a tax problem simply by doing the right thing too late.

Hadzima therefore organizes the hour as a sequence of timing problems. He begins with the life cycle of the venture, moves through intellectual property and ownership, tightens around patents as exclusion rights and strategic assets, and then closes with entity choice, tax elections, and the Section 83 trap. The unifying question is always the same: what has to be done early, before success itself makes the legal position more expensive?

4.1 Framing Legal Issues Along the Venture Life Cycle

Hadzima lowers the temperature before he raises it. This is background information, not legal advice. The reason is twofold and worth preserving exactly in spirit: legal outcomes are fact-dependent, and the law itself changes. The lecture is thus not a catalogue of eternal rules. It is a hazard map. If we see a particular pattern, we should know enough to stop, ask questions, and get help before stumbling into a pothole.

He also tells us, almost apologetically, that he is going to overload the room. That warning matters. It establishes the pedagogical contract of the lecture. He has one shot to put the structure in our heads; the slides and video can be revisited later. So we should expect density, but not randomness.

The first structural device is the venture life cycle. Across the top sit stages: idea, pre-outside-financing, first financing, second financing, and beyond. Down the left run recurring topic rows: intellectual property, legal entity, people, and financing. Even without the slide itself, the logic is clear. We are being asked to read legal issues in temporal order. What matters at the idea stage is not identical to what matters after a priced round.

That is why he begins where founders usually postpone thinking: "I have an idea," and "I am getting ready for outside financing." At that stage we are already forced to ask who owns the idea, what kind of protection exists, what legal body should hold the rights, and whether the company has been formed soon enough to avoid avoidable trouble later. The investor will ask those questions in diligence. Hadzima wants the founders to ask them while the answers are still cheap.

4.2 Why Intellectual Property Matters, and What Kind of Protection We Actually Have

The lecture's first quantitative move is motivational. Why worry about intellectual property at all? Because the value of firms has moved away from plant and equipment and toward intangibles. Hadzima uses a 2015 slide in which the intangible share of S&P 500 market capitalization is about 84%. The mathematical skeleton of his hierarchy is

$$\text{Intellectual property} \subset \text{Intellectual capital} \subset \text{Intangible assets}, \qquad \text{Intangible share of S\&P 500 market cap} \approx 84\% \tag{4.1}$$

This nested relation matters. Intangible assets include people, processes, ideas, and organizational capability. Intellectual capital is the narrower set of ideas and processes. Intellectual property is narrower still: it is the subset that the law lets us protect, license, assign, or lose.

Definition 4.1. Intellectual property is the legally protectable subset of a venture's intellectual capital. In Hadzima's practical register, it is the part of intangible value that can be fenced, transferred, and tested in financing.

He then walks through the ladder of protection in a very deliberate order. The lecture starts with "none," which is important. We are meant to remember that legal protection is not automatic. We may simply have ideas that we choose not to protect. From there he moves to trade secret, trademark or service mark, copyright, and patent. Each step is motivated by what it protects, what it does not protect, how long it lasts, and what it tends to cost.

A concise summary of the lecture's timing structure is

$$\text{Term}_{\text{trade secret}} = \infty \quad \text{while secrecy is maintained}, \tag{4.2}$$

$$\text{Term}_{\text{patent, U.S.}} = 20 \text{ years from filing}, \tag{4.3}$$

$$\text{Intent-to-use window} \leq 3 \text{ years}, \tag{4.4}$$

$$\text{Copyright duration} \approx 70+ \text{ years}, \tag{4.5}$$

$$\text{Copyright registration cost} \approx \$100 - \$500. \tag{4.6}$$

Trade secret is the simplest to state and one of the hardest to manage well. It lasts as long as we keep the information secret. That is why Hadzima immediately connects it to non-disclosure agreements. The trade secret is not magic; it is a secret that gives market advantage, and the law only helps if we behave as though secrecy matters. The lecture's practical rule is not to explain the secret sauce unless there is a specific reason, a specific recipient, and usually an agreement around that disclosure.

Trademark and service mark are different. They identify source, not technical function. A trademark is for goods, a service mark for services. Rights arise through use, and federal registration strengthens them. Hadzima keeps the founder on practical ground here. Before registration, we mark the claimed brand with TM or SM; only after federal registration does the registered-mark symbol become available. The mark can be a word, symbol, phrase, sound, or even color. The lecture's red-sole-shoe example is meant to make that breadth memorable.

Two practical consequences follow. First, a mark should be fanciful rather than merely descriptive. "Apple" for computers is protectable in a way that "micro dose" for low-dose aspirin is not. Second, one should search before investing heavily in branding. Rights arise through use, but Hadzima emphasizes the *intent-to-use* filing precisely because founders may spend a great deal on design and launch preparation only to learn too late that the mark is unavailable.

He also preserves a subtlety that matters to founders who think federal registration erases everything that came before. It does not. The senior unregistered user can survive in the geographic area where use began earlier. His McDonald's example makes the point vividly: a small earlier user may survive locally, even after the national chain registers federally.

Copyright is introduced with equal care. It arises upon creation, not registration. If Hadzima writes a letter, he already has a copyright in the expression. But the right protects expression, not idea or function. That is why Shakespeare's treatment of a love triangle is protectable while the abstract idea of a love triangle is not. The distinction is not philosophical decoration. It explains why copyright is a natural fit for literature and music, and a more awkward fit for software.

The lecture then becomes founder-facing again. Registration is cheap relative to patents, but if one wants to sue, registration matters. The work must be deposited, and in software that raises obvious anxiety: must we deposit source code? Hadzima's answer is procedural rather than doctrinal. There are ways to deposit redacted portions so that authorship can be proved without giving away everything.

His Aero Maps stories serve the same didactic function. Phantom roads are inserted so copying becomes detectable. The Boston directory story is even sharper: the city stopped buying updated directories because it had copied the old one. That anecdote does not merely illustrate infringement. It demonstrates why a venture should put notices on what it wants to protect, and why one should not confuse "I made it" with "I can prove I own and can enforce it."

The most important copyright lesson for founders is ownership. The author of a work owns it unless the work falls within employment or a valid work-for-hire arrangement. Hadzima's import-export software story is painful for exactly that reason: a company can spend roughly half a million dollars building software and still not own it if the contract never assigned ownership. The companion lesson is symmetric. When a startup consults for a large company, the services agreement will often attempt to transfer everything created. Founders must negotiate the difference between the deliverable and the preexisting tools, modules, and know-how used to produce it.

The lecture adds two modern edges to this section. One is open-source compliance: if we modify open-source code in ways that trigger license obligations, that becomes part of diligence. The other is artificial intelligence. Hadzima's lecture-era caution is simple enough to keep: copyright requires human authorship, AI alone is not the author, and the boundary of "sufficient human authorship" remains unsettled.

4.2.1 Question & Answer

Why worry about intellectual property if we only have an idea?

Because at the earliest stage the venture is mostly intangible. The people, the process, the know-how, the technical insight, and the first sketches of product-market fit are the company. If value has shifted toward intangibles, then ownership and protection are not side issues. They are the first structural questions. Before we ask how fast the venture can grow, we must ask what exactly it

owns, what it can protect, and what later investors will be able to underwrite.

The lecture's rhythm also matters here. Hadzima does not begin with patent doctrine. He first makes the room care about intangible value. Only then does the taxonomy of protections make economic sense.

4.3 Patents as Exclusion Rights, Timing Constraints, and Strategic Assets

Once patent appears in the taxonomy, the lecture changes pace. The earlier categories are compared and illustrated; patent is analyzed. Hadzima calls it a limited-time monopoly, but immediately sharpens that phrase into its legal meaning. A patent is a federally granted exclusion right. It lets us prevent others from making, using, selling, or importing what falls within the claims. It does not, by itself, guarantee that we are free to practice every commercially useful embodiment of our invention.

Proposition 4.2. *A patent is a right to exclude, not a complete affirmative permission to practice an invention.*

Proof. Take Hadzima's cup example. Inventor A patents a vessel to hold a liquid. Inventor B later patents a handle. A commercial cup with a handle practices both inventions. Inventor A cannot commercialize the handled cup without infringing B's handle patent. Inventor B cannot commercialize the handled cup without infringing A's vessel patent. Each can exclude; neither has unilateral freedom to practice the combined product. A cross-license is required. □

The real-estate analogy now becomes exact enough for the lecture's purposes. The claims are the fence. They define what is inside the property right. But a fence is not a road. It marks exclusion, not universal freedom of use.

We should also preserve the deeper social logic Hadzima mentions. Patent law is a trade: society grants a limited monopoly in exchange for disclosure. That is why patent "cuts both ways." It gives a powerful exclusion right, but only after we explain the invention well enough that others can learn from it. The system is justified not because monopolies are generally loved, but because the disclosure permits later cumulative invention.

The lecture's checklist for patentability is standard but should remain explicit: the invention must be new, non-obvious, useful, and within patentable subject matter. It must also survive the timing traps created by disclosure and sale. Hadzima's compact timing rule is

$$\text{U.S. patent filing window after an enabling disclosure or an offer for sale} = 1 \text{ year}. \tag{4.7}$$

The word "enabling" does real work. "I have anti-gravity boots" is not enabling; it tells us nothing about how the device works. A detailed public explanation may be enabling. Once the disclosure is enabling, or once the invention has been offered for sale, the one-year U.S. clock begins. In many other jurisdictions the rule is harsher: pre-filing disclosure may destroy foreign patent rights immediately. Hadzima's story of the Egyptian team that published its result in a journal before considering patent consequences is included precisely to show how a scientific success can become a commercialization loss.

Non-obviousness receives a similar practical treatment. Hadzima does not turn it into doctrinal abstraction. He points instead to arguments of the form: the literature said this could not be done, and the market is now buying it, so it could not have been obvious to a person of ordinary skill in the art. The point is not that commercial success automatically proves non-obviousness. The point is that the analysis sits inside a real evidentiary world.

The lecture also marks a historical shift: first-inventor-to-file. The timing pressure is no longer merely "invent first." It is "file first, provided you are an inventor." That is why collaborations, joint ventures, and ambiguous middle-ground inventions become dangerous. Each side may think the jointly developed improvement is still in discussion while the other side is already racing to the office.

Before leaving the doctrinal core, Hadzima inserts another essential founder reminder: just as the author owns the copyright, the inventor owns the patent unless rights have been assigned. Companies therefore insist on invention disclosure and assignment agreements. The practical rule is simple: if the company needs to own the invention, the assignment must not be left to implication.

This leads directly to freedom to operate. A founder can spend time and money building a product only to discover that someone else's patent blocks launch. The freedom-to-operate opinion is the tool Hadzima mentions here. It is not magic protection and not a judicial decree of safety. It is a law firm's reasoned view that the proposed activity does not obviously infringe. That opinion also matters because willful infringement can produce treble damages, and so documented caution changes litigation posture.

4.3.1 Question & Answer

If we have a patent, are we free to practice the invention?

No. We are free only with respect to the claim boundary we own. Other patents may still cover indispensable components, adjacent improvements, manufacturing methods, or commercial embodiments. That is why Hadzima insists on keeping two questions distinct. "Can we get rights?" is a patentability question. "Can we launch without being blocked?" is a freedom-to-operate question.

This is one of the lecture's most important separations of mechanism from recommendation. The mechanism is exclusion by claim. The recommendation is to ask about infringement before market entry, not after.

4.4 Patent Process, Prior Art Mapping, and Cost Discipline for Startups

From here the lecture pivots from doctrine to operating strategy. If patents are exclusion rights in a crowded field, then the founder's real problem is no longer just "what can be patented?" but "what should be patented, when, at what cost, and in what competitive terrain?"

Patent mapping is introduced through an arithmetic failure. Hadzima describes sending out for an outside prior-art search. It took about six months. During those six months roughly 5000 new U.S. patents were issuing per week. The rough consequence is

$$5000 \text{ new U.S. patents/week} \times 26 \text{ weeks} \approx 1.3 \times 10^5 \text{ new patents.} \tag{4.8}$$

The precise number is only a back-of-the-envelope estimate, but the strategic conclusion is solid: a static search can be obsolete by the time it returns.

That motivates the patent map. Hadzima describes it carefully. A box is a patent or patent application. The left edge of the box marks issue or publication date. A tail extending left marks filing date. The lines connecting boxes are citations. These citations matter in a patent-specific way. They are not merely academic references. Prior art must be disclosed to the patent office, and willful failure to disclose relevant prior art can threaten the validity of an otherwise strong patent because it may be treated as a fraud on the office.

Tesla is the lecture's main illustration. One can look at a single Tesla patent and see both the prior art behind it and the later patents that cite it. One can also look at the portfolio over time and infer where the company is extending a core invention. A cluster of later patents citing a core patent may tell us where the company is building a continuation or extension strategy.

The deeper lesson is business-facing. Hadzima gives two opposite stories from SBIR-style commercialization work. In one case a cluster of patents suggested a thicket to avoid; a sparser region suggested a better entry point. In another case the same kind of cluster looked attractive, because the founders believed they had solved the very problem that those clustered filings were circling. In both cases the map functioned not as a legal ornament but as a tool for choosing a beachhead market and even for deciding whether acquisition-by-portfolio might be the right strategy.

That is the point at which the lecture asks what to patent. Large firms may reward sheer filing activity. Startups cannot afford that luxury. The founder's question is not "what claims could possibly be obtained?" but "which filings add value, create leverage, support licensing, or protect a market worth entering?"

When should we file? Hadzima's answer is layered. We file before losing rights through public disclosure or sale. We file in time to win the first-inventor-to-file race. We file in time to support a licensing or acquisition conversation with actual issued or pending rights. And if full filing is too expensive too early, we use the provisional system as a disciplined placeholder.

$$\text{Provisional priority period} = 1 \text{ year}, \tag{4.9}$$

$$\text{Provisional filing fee}_{\text{micro}} \approx \$64. \tag{4.10}$$

The lecture is precise about what a provisional does and does not do. It establishes a filing date. It need not contain claims. But it must contain a meaningful description of the invention. Only what is described gets the earlier priority. This is why Hadzima's fruits-and-vegetables example matters. If protein becomes part of the invention later, and protein was not described in the provisional, the original filing date does not magically expand to cover the new matter.

That makes the provisional less like a complete patent and more like a disciplined temporal marker. Hadzima even mentions the startup version of this: a conference paper can be turned into a provisional by filing it with a cover sheet, provided it truly describes the invention well enough. For a cash-constrained startup, that can be the difference between preserving rights and losing them.

His own example from IP Vision makes the discipline concrete. Development ran in six-week cycles. Before deployment, the team asked what had been built, what had been thought through but not

yet implemented, and what was important enough to file. Provisional applications became a routine timing tool rather than a rare heroic event.

The structure of a full application also enters here. Hadzima describes it almost as a term paper: field of invention, background, summary, detailed description, and then claims. The crucial doctrinal point returns: the detailed description must disclose the invention and its best mode of practice. The claims then define the boundary. Because claim drafting is the most technical part of the process, this is where he wants founders to spend money on qualified patent counsel rather than improvising.

Costs now become explicit:

$$\text{Patent preparation cost} \approx \$5{,}000 - \$15{,}000, \tag{4.11}$$

$$\text{Outlier patent-preparation bill} \approx \$35{,}000, \tag{4.12}$$

$$\text{Foreign patent cost}_{\text{10 major countries}} \approx \$330{,}000 - \$500{,}000. \tag{4.13}$$

The $35,000 anecdote is not a statistical estimate; it is a governance lesson. The co-founder kept calling the patent lawyer directly, adding changes, and running the meter. Hadzima's response was managerial, not theoretical: no one talks to the patent lawyer without coming through a single gatekeeper. Patent cost is partly doctrinal, partly procedural, and partly behavioral.

Prosecution cost is open-ended in the same way. The examiner often rejects first claims, and Hadzima even says that can be a good sign, because immediate allowance may mean the applicant tried to claim too little. But every back-and-forth round costs money. Fixed-fee patent packages may therefore yield claims that are formally allowed and strategically weak.

Foreign filing raises the economic question again. If an exclusive licensee can be brought in early, part of the bargain may be that the licensee carries later patent costs. This turns timing of licensing into a financing strategy for IP itself.

4.4.1 Question & Answer

Should we patent this idea, or disclose it to block others?

Sometimes the right move is not to patent everything, but to patent the core and disclose the periphery. Hadzima's "picket fence" example explains why. Suppose we hold a strong core patent on the vessel to hold a liquid. Others might patent the top, the sleeve, the handle, or the insulating shell and thereby block commercial use of the core invention. If the core patent is strong, a cheap counterstrategy is to disclose those peripheral possibilities without claiming them. The disclosure becomes prior art and makes those later blocking patents harder or impossible to obtain.

This is not a universal rule. It depends on confidence in the strength of the core patent. But it is exactly the kind of startup arithmetic the lecture wants us to learn: spend carefully, disclose strategically, and think in terms of market control rather than filing count.

Tesla is the larger version of the same point. Hadzima uses Tesla not to preach "open patents," but to broaden the notion of patent strategy. If the goal is to sell electric vehicles, complementary infrastructure matters. Giving some patent access away may stimulate the ecosystem the company needs. The legal instrument is still patenting; the business objective is now ecosystem design.

4.5 Working With Lawyers and Moving University IP Into a Startup

At this point the lecture compresses the IP discussion into founder checklists. Employees should sign invention disclosure and assignment agreements. Consultants should be governed by proper work-for-hire or assignment language. Non-disclosure agreements should be used where appropriate, but with eyes open: venture investors often refuse them because they see many similar companies and do not want later claims that they misused one founder's disclosure when funding another.

The operational principle remains the same. Do not disclose the secret sauce until the last responsible moment. Make sure the company owns what it thinks it owns. Understand the patent landscape over time, not once. Preserve rights by filing on time.

Hadzima uses even the provisional-application discussion to make a cultural point. Large companies sometimes resist provisional discipline because it is hard to track. Startups cannot afford that luxury. Their problem is not internal paperwork elegance. Their problem is losing rights while moving too slowly.

Then comes the lecture's comic but serious interlude on lawyers. Legal expertise is bought by the hour. If we dump a lab notebook on the desk and say "write me a patent," we are buying expensive confusion. If we wait until the one-year deadline is nearly gone and demand a filing by Monday, we are buying emergency pricing. If instead we prepare invention disclosures, read patents in the area, draft the technical description, and let counsel focus on claims and legal framing, we turn legal spend into leverage.

The next founder question then arrives naturally: what happens when the IP comes from a university? Hadzima had been asked about this on a founders' panel, and he answers by inserting a quick institutional history. Before Bayh-Dole, federally funded university inventions belonged to the federal government, were managed from Washington, and were typically licensed non-exclusively. For a startup that wants to build a business around a specific technology, this was a poor commercialization regime.

$$\text{Pre-Bayh-Dole commercially licensed patents} \approx 5\%. \tag{4.14}$$

Bayh-Dole changes the geometry. Universities retain ownership, file patents, license them, prefer small businesses, give the government a royalty-free license, build commercialization programs, and share royalties with inventors. The modern technology licensing office is the institutional consequence of that change.

MIT appears in the lecture as a comparatively founder-friendly case. Hadzima describes the Technology Licensing Office as easier to work with than many counterparts because its first priority is getting inventions into the world, not maximizing immediate revenue. That difference in mission shapes terms and process.

The ownership rules matter. Significant use of MIT facilities or MIT-administered funds typically means MIT owns the invention. MIT does not assign the IP outright; it licenses. Sponsors may have first licensing rights. The practical moral is plain: if we are building around university technology, we must know whether the invention belongs to the company, the founder, the sponsor, or the university. Investors will eventually force the issue if the founders do not.

Hadzima's Ali example and his Kendall Square explanation show how this becomes operational. If

the work was not done with significant use of MIT facilities or funds, the founder may obtain a letter clarifying non-ownership by MIT. If the invention does belong to MIT, the startup typically develops it across the street, in commercial space, not in university labs. That physical move mirrors the legal move.

4.5.1 Question & Answer

How do we take university IP into a startup without quietly mis-owning it?

By bringing it through the front door. Hadzima's recommended device is often an option to obtain a license. The option gives the startup time to explore markets, talk to investors, and raise capital without pretending that the technology has already been transferred into the company. It is the institutional analogue of the provisional patent application: a way to preserve position while uncertainty is still being resolved.

The lecture preserves several useful lecture-era financial ranges:

$$\text{University license issue fee} \approx \$50\text{K} - \$150\text{K}, \tag{4.15}$$

$$\text{University running royalty} \approx 3\% - 5\%, \tag{4.16}$$

$$\text{MIT startup equity stake} \approx \text{single-digit } \%. \tag{4.17}$$

He also describes a specifically startup-oriented equity practice in which the university may take a small equity position and keep that percentage fixed through the first outside financing round. In his example, if MIT holds 5% and an investor later buys 20%, MIT still sits at 5% through that first round rather than being diluted immediately with the founders. That is not a universal law of licensing. It is a lecture-era practice point, and it changes how founders should think about early capitalization.

4.6 Entity Choice, Tax Elections, and the Section 83 Trap

The lecture now tightens again. After asking who owns the technology, Hadzima asks what legal body owns the company. If founders do nothing, they may already have a general partnership: one or more persons carrying on a business for profit. That requires no written agreement, which is precisely why it is dangerous. Liability is joint and several. Ownership, absent a contrary agreement, follows money rather than founder effort. Neither feature is hospitable to a venture-backed startup.

The practical recommendation is therefore immediate incorporation, usually as a corporation, and often in Delaware if repeated outside financing is likely. Delaware matters here not as brand mystique but as a jurisdiction with well-developed corporate law and comparatively predictable outcomes. Hadzima also insists on corporate formalities. A corporation only protects if it is actually treated as a separate legal person in records, signatures, and operations.

Tax then enters through that corporate personhood. A corporation taxed under Subchapter C bears tax at the corporate level, and distributions to shareholders are taxed again. A Subchapter S election changes the structure by making the income pass through to the owners.

$$\text{C-corp asset sale} : \text{corporate-level tax} + \text{shareholder-level tax}, \tag{4.18}$$

$$\text{S-corp asset sale} : \text{pass-through to the owners without corporate-level tax}. \tag{4.19}$$

The eligibility conditions are narrow enough that Hadzima states them explicitly:

$$\text{S-corp shareholder cap} < 100, \qquad \text{classes of stock} = 1, \qquad \text{S-election filing window} \approx 3\,\text{months after incorporatio} \tag{4.20}$$

A venture capital investor usually destroys S status because the shareholder rules are restrictive and only one class of stock is permitted. But Hadzima does not let the audience dismiss the election as irrelevant. Founders say: we will be loss-making for years, so why worry now? His answer is the exit. Small acquisitions are often structured as asset sales. In that case, the C corporation pays tax on asset gain, while the S corporation does not bear corporate-level tax in the same way.

The lecture gives two vivid numbers here: an illustrative exit-tax difference on the order of 13%, and the DIVA story in which an early S election reportedly saved about $2,000,000. These are example numbers, not timeless constants, but they preserve the lecture's logic: early entity and tax choices can matter only at the end, which is exactly why founders neglect them.

Then Hadzima marks the climax audibly: if we are falling asleep, wake up. This is the big trap. Section 83 says that if we receive property in connection with services, we recognize ordinary income equal to the fair market value of the property minus what we paid for it. In the lecture's notation bank, the core formula is

$$I_{\text{ordinary}} = \text{FMV} - P. \tag{4.21}$$

To make the tax consequence explicit, we may write

$$T = \tau(\text{FMV} - P), \tag{4.22}$$

where T is the tax owed and τ is the marginal tax rate used for the example. Hadzima gestures at "37% or whatever" and then computes with 40%, so we keep the symbolic formula and treat 40% as a rough example rate.

Now the lecture turns the formula onto the founders themselves. Suppose an investor pays $1,000,000 for 50% of the company. What is the founder's 50% worth? Hadzima first lets the room infer a simple answer: perhaps about $500,000. If the founder paid essentially nothing for those shares, then

$$I_{\text{ordinary}} \approx \$500{,}000, \tag{4.23}$$

$$T = 0.40 \times \$500{,}000 = \$200{,}000. \tag{4.24}$$

He then pushes the room toward an even harsher inference. The wording is damaged in the transcript, but the intended arithmetic is clear enough: if the founder's block is instead treated as worth about $1,000,000, then

$$I_{\text{ordinary}} \approx \$1{,}000{,}000, \tag{4.25}$$

$$T = 0.40 \times \$1{,}000{,}000 = \$400{,}000. \tag{4.26}$$

The economic violence of the example is the point. Founder stock can create a personal tax bill long before the founder has cash liquidity.

Then comes the second blow: vesting. If the stock vests over time, the income may be measured when the stock vests rather than when it is first issued. So the tax base grows with success. The lecture's timing structure is

$$T(t_v) = \tau\big(\text{FMV}(t_v) - P\big), \qquad \text{FMV}(t_v) \gg \text{FMV}(t_0), \tag{4.27}$$

where t_0 is the early issuance date and t_v the later vesting date. Hadzima's plain-English version is better than any abstraction: when the company begins, the idea is a dime a dozen; when the stock vests two years later and the company has done well, the shares may be worth \$100 each. That is exactly when the tax becomes lethal.

His California CFO story is included to show that this is not theoretical. A later executive asked for the 83(b) election paperwork, and the company suddenly realized it had issued a large amount of vesting stock without understanding the tax consequences. What looked like a successful company was also a ticking time bomb.

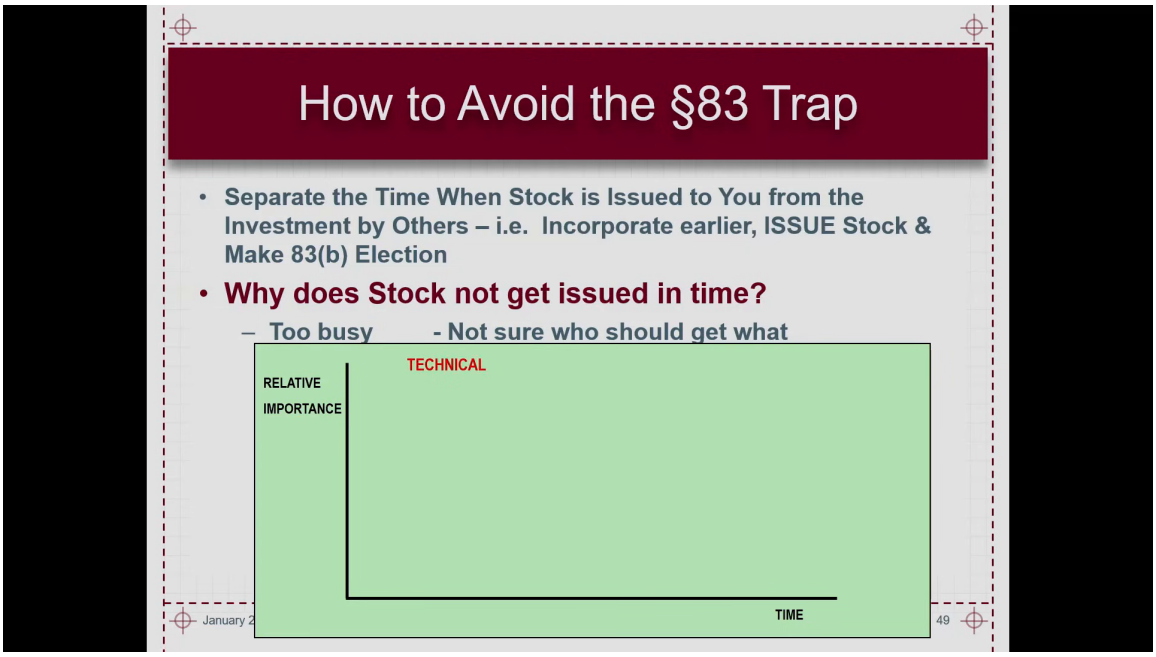

Figure 4.1: Section 83 trap slide with timing-versus-priority chart. The slide itself gives the operational recommendation: separate founder stock issuance from later outside investment and make the 83(b) election.

The screenshot matters because it preserves the lecturer's own wording. The top bullet says, in essence, separate in time the issuance of founder stock from the later investment by others:

incorporate earlier, issue stock, and make the 83(b) election. The second bullet asks why stock does not get issued in time. The visible answers are embarrassingly human: the founders are too busy, and they are not yet sure who should get what.

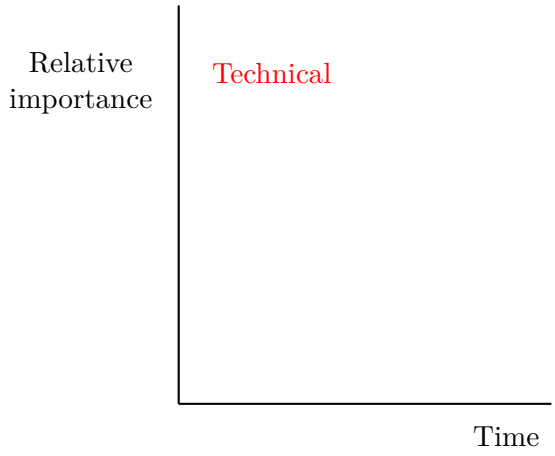

Figure 4.2: A pocket-scale reconstruction of the visible timing sketch. The frame supports only a qualitative claim: technical work is early and highly important, and that very urgency helps crowd out timely stock allocation.

4.6.1 Question & Answer

Why does stock not get issued in time, and how does an 83(b) election avoid the trap?

Because early ventures postpone the awkward split conversation while technical work, customer work, and financing consume attention. By the time the founders resolve who gets what, an outside investment or other valuation event may already have made the stock expensive for tax purposes. The 83(b) election changes the measuring time. It aims to tax the stock at the earlier, lower-value moment of transfer rather than the later, higher-value moment of vesting. But that only works if the stock has actually been issued, and if the issuance can be separated in time from the later event that makes valuation obvious.

That is why the lecture's recommendation is temporal, not merely formal: incorporate early, issue the stock early, let later cash investment occur later, and then file the 83(b) election promptly.

4.7 Founder Equity as Negotiated Structure, Not Myth

The lecture does not end by closing the tax folder and declaring victory. It spills, deliberately, into founder culture. How much equity should each founder have? Hadzima refuses to turn this into mythology. The amount of equity and who gets it is, in his words, about culture and negotiation.

His Keenan Systems example serves as a counter-myth. One founder kept all of the stock, paid employees very well, and later distributed bonuses when the company was acquired. That was a culture, not a universal fairness theorem. The lecture uses it to break the habit of thinking that there is a single morally canonical founder split.

The variables Hadzima wants on the table are concrete: value of past contribution, value of future

contribution over a specified time horizon, ownership of IP, sacrifice, commitment, reputation, opportunity cost, and the signaling effect of who is actually all-in. His professor-and-postdoc story preserves that method. Each person writes down what they contribute and what they give up. The professor brings name recognition, perhaps access to capital, perhaps one consulting day a week under MIT policy. The postdoc gives up a salary path and perhaps an academic career. The conversation becomes fairer not because someone solved it algebraically, but because the actual variables became visible.

The lecture then returns to numerical structure in the option-pool discussion. The equity compensation pool is not a decorative cap-table entry. It is a reserve sized for the next hiring horizon, typically two to three years and tied to the headcount plan. Hadzima's lecture-era range is

$$\text{Equity compensation pool after the first outside round} \approx 12\% - 18\%. \tag{4.28}$$

The lower end corresponds to a company that already has more of the needed senior team. The upper end corresponds to a company that still must recruit experience. Once again the lecture translates culture into timed structure. Who gets the pie and how large are the slices? We are supposed to have reasons, not slogans.

The ending remains intentionally unfinished. Hadzima notes that later lectures and slides will treat ownership examples, dilution, and financing instruments in more detail. That incompleteness is part of the chapter's rhythm. Founder equity is introduced here not as a closed calculus, but as the bridge from legal formation into financing.

4.8 Summary

The lecture moves from a high-level claim about intangible value to a founder-level tax trap, but the thread tying the whole argument together is timing. We ask when an intangible becomes property, when a disclosure becomes enabling, when a provisional must become a full application, when university ownership must be formalized by option or license, when a corporate election matters only because an exit may later occur, and when founder stock becomes taxable income.

If one lesson remains after the details blur, it is this: ventures are often injured not by exotic legal doctrine, but by ordinary legal rules encountered too late. Ownership, disclosure, assignment, filing, incorporation, election, vesting, and valuation all sit on a clock. Hadzima's practical achievement in this lecture is to make us see that clock before it starts billing us.

Chapter 5

Organizational and People Issues

This session from Joseph Hadzima's MIT OpenCourseWare course is not a doctrinal lecture but an interactive panel on how ventures are actually held together. That matters for the way we should write it. The argument does not proceed by a single theory of organizations. It proceeds by tension, question, example, correction, and return. The mathematics is correspondingly modest but still real. It appears here as threshold logic, scaling examples, decision rules, and compact tests for when an organization has crossed from structure into obstruction. These notes follow the transcript closely and are curated for book form by LazyingArt LLC.

5.1 People Issues As The Dominant Failure Mode

Hadzima begins by recalling the first day of the course. A new venture requires several elements to come together, but the present session narrows the field sharply. The factor most likely to sink the venture is not the technology and not the funding. It is the people side of the enterprise. We should read that as an ordering of practical causes, not as a numerical claim.

$$\text{people issues} \succ \text{technology}, \qquad \text{people issues} \succ \text{funding}. \tag{5.1}$$

That ordering explains the form of the evening. If the principal failure mode is human rather than purely technical, then a panel is more useful than a sermon. Vivian moderates. One panelist speaks from founder-scale people building and boards. Another speaks from large-company HR and acquisition. A third brings process and operations design. The room is then invited to interrupt, ask, challenge, and refine. The lecture's rhythm is therefore not accidental. It is part of the teaching method.

Remark 5.1. The displayed formulas in this chapter are transcript-backed heuristics rather than board equations from the room. They are included only to make the lecture's operating logic visible.

The opening claim also gives the chapter its tone. We are not dealing with "soft" material in the dismissive sense. The whole session argues the opposite. People are not one more function inside the company. They are the medium through which every other function is executed, delayed, clarified, or spoiled.

5.2 Culture Starts Before HR

The first substantive correction in the panel is decisive. Culture does not begin when a company finally hires an HR person. It begins at founding. Before any formal department exists, the founders have already started defining the company: what matters, how quality is judged, how they speak to one another, what behavior is rewarded, how mistakes are handled, and whether trust is real or merely assumed.

The lecture returns several times to a small set of early mechanisms.

1. founders must speak explicitly about what is important;
2. trust must be present in seed form and then built in practice;
3. hiring behavior already creates culture;
4. routine gives the company cadence and therefore security.

The point about routine is worth pausing over, because the panel gives it more weight than a reader might expect. Vivian compares management to parenting in one narrow but useful sense: regular routines do not merely organize time, they stabilize a human environment. In a tiny company, this matters. If everything is improvised, every error feels existential. If there is cadence, then the company can make mistakes without dissolving into panic.

5.2.1 Question & Answer

Question. When should we start thinking about the HR function, and how do we preserve the company's DNA as investors arrive?

Answer. The answer is earlier and later than one might think. Earlier, because the real content of culture begins immediately. Later, because the formal function often comes after founders, advisors, and a small set of routines have already carried the organization through its first stage.

Let n denote headcount, and let $\mathrm{support}(n)$ denote the amount of explicit people and organizational support the company requires at stage n. A fair compression of the panel's threshold logic is

$$\mathrm{support}(n) = \begin{cases} \text{founder-led / informal}, & n < 25, \\ \text{outside guidance useful}, & 25 \le n \le 50, \\ \text{explicit people support hard to avoid}, & n \gtrsim 50. \end{cases} \tag{5.2}$$

This is a heuristic, not a law. The exact threshold depends on the business. But the lecture is plain that somewhere around the 25 to 50 range outside help becomes materially useful, and by roughly 50 people the company is usually beyond pure improvisation.

Worked example. The value of the heuristic is not predictive precision. Its value is that it makes regime change visible.

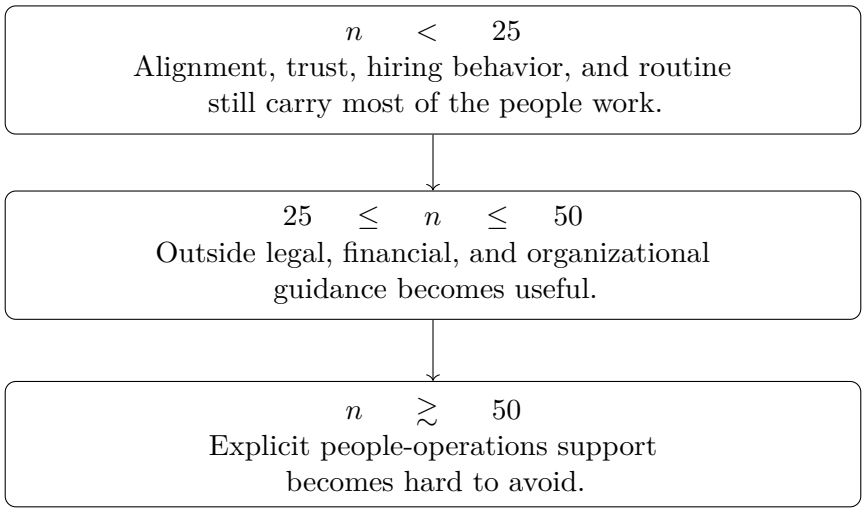

Figure 5.1: Transcript-derived stage heuristic for people and organizational support. The lecture presents these as rough thresholds, not exact boundaries.

$$n = 12 \Rightarrow \text{support}(n) = \text{founder-led / informal}, \tag{5.3}$$
$$n = 35 \Rightarrow \text{support}(n) = \text{outside guidance useful}, \tag{5.4}$$
$$n = 60 \Rightarrow \text{support}(n) = \text{explicit people support becomes urgent}. \tag{5.5}$$

At $n = 12$, founder attention, trust, and routine still carry most of the load. At $n = 35$, legal, financial, and organizational advice become useful. At $n = 60$, the firm can no longer assume that role design, fairness, and people process will take care of themselves.

The panel then adds two practical refinements. First, early companies often bring in legal or financial help before they bring in a human-resources specialist, because payroll, taxes, and basic financial control become urgent quickly. Second, once the people function does arrive, it should not simply disappear under finance. The lecture is unusually direct here: do not make HR a shadow of the finance office. The functions are connected, but not identical.

Stage and scale. At this point the panel turns from timing to meaning. The question is no longer only when support becomes necessary. It becomes what sort of company is being supported. Two scale markers sharpen the point:

$$n_{\text{RSA}} : 55 \to 2000, \tag{5.6}$$
$$n_{\text{Thermo}} \approx 130000. \tag{5.7}$$

The RSA example is used to show stage dependence. A company that moves from roughly 55 people to roughly 2,000 does not remain organizationally unchanged. At each stage of growth, different things are needed. But the values can remain stable across those changes. The Thermo Fisher example shows something complementary: even at very large scale, a simple mission can still bind the organization. The point is not sloganism. The point is that dispersed tasks must still be legible to the people performing them.

This is why the lecture immediately resists a false opposition. Vivian says that the organizations she most admired combined high achievement with human kindness. High standards were not softened. What mattered was that there was still a human way of doing difficult things. Even layoffs or structural changes could be carried out with respect.

Example. Michelle's "people first" company makes the same point from another angle. The founder's rule was simple: if you take care of your people, your people take care of your customers, and the rest has a chance to work. That is not sentimentality. It is a theory of organizational transmission. Customer treatment, revenue generation, and retention of talent are not separate lines. They are linked.

5.3 Roles, Friendship, And The Co-Founder Boundary

Once culture has been named, the next obstacle appears naturally. Small companies are personal companies. Friendship, prior collaboration, and mutual admiration often precede the formal structure. The panel does not deny that. What it insists upon is that friendship cannot substitute for role clarity.

That is why the lecture narrows here from culture in general to the structure of daily authority. If two people go into business together, they must speak early about who does what, how disagreements are handled, what norms of behavior govern the work, and where final authority lies. The point is preventive. We do not let these questions "work themselves out" under pressure. We make them explicit before pressure turns vagueness into resentment.

The panel's advice is vivid because it comes with lived examples. Some working friendships become deep precisely because people built something hard together. But that does not mean the friendship came first. Quite often the respect was formed in the work itself. This is one of the session's quieter but stronger insights: durable friendship is frequently the result of good structure, not a replacement for it.

5.3.1 Question & Answer

Question. What is a co-founder, and when should title, control, or equity wait until fit is clearer?

Answer. The panel treats "co-founder" as a real but somewhat ambiguous term. In the cases Vivian describes, co-founders usually came in together. It was not a title handed out later to someone asking for recognition before the fit was clear. The practical rule is therefore cautious: do not give away title, control, or equity too early when the relationship itself is still being tested.

This answer is not hostile to partnership. It is merely stage-aware. If a person brings indispensable expertise, and that expertise complements what the existing founder cannot do, then a deeper commitment may well be justified. But the panel advises founders to take time: work on a problem together, test the relationship, or start with a more limited structure before making the role permanent.

This is also where the lecture states one of its clearest decision rules. Co-CEO arrangements usually fail. One person generally has to be in charge. That does not mean the second person contributes

less value. It means the organization requires a clear center of responsibility. The lecture gives a concrete alternative: a founder may choose to be the evangelist for the product rather than the operator of the company. If that is decided early, then the business has a better chance of staying coherent.

5.4 Hiring As A Decision Process, Not A Credential Filter

Once titles and roles have been made more precise, the lecture asks the next question: how do we test important hires without binding ourselves too early? This is where the panel becomes especially practical. Hiring is not merely recruitment. It is culture formation under uncertainty. Early mistakes in hiring do not remain local. They alter the internal atmosphere of the company.

The core recommendation is simple: if the role is uncertain and the work permits it, try before you buy. Bring the person in as a consultant, or structure a trial relationship that reveals real contribution rather than just verbal self-description.

5.4.1 Question & Answer

Question. How do we test uncertain hires without locking ourselves into the wrong long-term relationship?

Answer. Introduce uncertainty into the contract rather than burying it inside an irreversible title. Let t_{fit} denote the time required to judge whether a working relationship is actually viable. The panel's rule of thumb is

$$t_{\text{fit}} \lesssim 3 \text{ months} \approx 90 \text{ days}. \tag{5.8}$$

The old 90-day probation period is mentioned as the institutional analogue, not as a hard prescription. In many ordinary cases one can judge within this horizon whether the person contributes, whether there is good give-and-take, and whether the relationship is likely to strengthen the company. Sometimes, the panel adds, one can know much sooner.

The compensation logic follows the same principle. If possible, pay cash for the trial rather than immediately reaching for stock or a deep long-term commitment. If cash is scarce, negotiate. But do not transform uncertainty into structure faster than the evidence warrants.

The lecture then turns the hiring process outward, from the résumé to the person. Vivian is explicit that she would often take executive candidates into an informal setting and watch how they behaved with others, including service staff. The point is not etiquette as a moral show. The point is diagnostic. A formal interview can conceal precisely the human subtleties that daily work will expose.

We can summarize the panel's hiring logic in four moves:

1. test the claimed capability;
2. observe real collaboration rather than self-presentation alone;
3. judge the whole person, not only the technical skill;

4. delay irreversible commitment until the fit becomes visible.

This is also where the lecture folds back to culture. Hiring is one of the principal ways culture is created. The company teaches people what it is by whom it admits, how it evaluates them, and how it treats them whether or not they are hired.

5.5 Growth, Acquisition, And The Structure–Bureaucracy Trade-Off

The panel now broadens the frame. Having asked how we build the inside of the firm, the room asks how the firm survives outside itself: how a startup competes with large incumbents, and how an acquisition can succeed without destroying what was purchased.

The contrast between small and large firms is stated in almost mechanical terms. Large firms sell size, scale, reach, and the ability to service a broad range of needs. Smaller firms win somewhere else: speed, flexibility, relationship, and often price. That is why the lecture insists that the small firm's advantage is not mystical. It lies in reduced rigidity.

The acquisition discussion is equally concrete. A large company does not buy only product. It often buys people, accumulated technical understanding, founder legitimacy, and the local loyalty of a team. The panel therefore stresses retention of founders and core talent as part of the value being acquired. The acquired unit may need to remain substantially intact for a while so that the buyer can understand what it has actually bought.

If t_{intact} denotes the period during which an acquired unit is left largely intact, and if n_{retrofit} denotes the headcount at which missing structure becomes painful to repair, then the panel offers the following two scale markers:

$$t_{\text{intact}} \approx 1 \text{ year}, \tag{5.9}$$

$$n_{\text{retrofit}} \approx 500. \tag{5.10}$$

The first is the integration example: one acquired company was allowed to run more or less intact for about a year before deeper integration. The second is the retrofit warning: Vivian describes joining a company at about 500 people where basic structure still lived on spreadsheets and no coherent job architecture existed. Later repair was therefore painful.

5.5.1 Question & Answer

Question. When do process and hierarchy help the company, and when do they start slowing the business down?

Answer. The lecture's answer is subtler than startup folklore usually allows. Michelle argues strongly for early SOPs and documentation. Put simple process in place early. Let people know what the work requires. Give them a place to return to. But Vivian immediately adds the necessary correction: process is not the same thing as bureaucracy.

Let P denote the active process load in the company, and let $B(P) = 1$ denote that process has crossed into bureaucracy. The lecture's test can then be written schematically as

$$B(P) = 1 \iff P \text{ impedes customer needs, hiring speed, or revenue without a compensating gain in clarity, fairn} \tag{5.11}$$

That is a functional definition rather than a moral one. Process is useful when it makes work clearer, fairer, and more reproducible. It becomes bureaucracy when it ceases to help the business run.

Proposition 5.2. *If a company postpones basic structure until it is already large, the eventual repair is more expensive than early simple process.*

Proof. At small scale, routines can be written and revised locally. At larger scale, payroll, job structure, reporting, benefits, fairness, and compliance are already coupled across the firm. Repair is then no longer local. It must be performed while the company continues to operate. The panel's 500-person spreadsheet example is precisely such a late retrofit. □

The lecture's positive model is therefore quite clear: simple process early, hierarchy only where it actually serves the work, decision-making as low in the organization as one can responsibly allow, and continual testing of whether the current process still makes business sense. A flat organization is not automatically better, but needless hierarchy is one of the panel's most reliable indicators that a company has started to confuse management with obstruction.

5.6 Founder Self-Awareness, Conflict, And Letting Go

Having moved through competition, acquisition, and structure, the lecture turns inward again. A founder asks how to lead while still being herself, especially if she is more relational than hard-edged. The panel's reply is one of the most important in the session: do not perform a borrowed executive character. Be yourself. Become more self-aware. Then hire complements around your weaknesses.

This part of the lecture matters because it changes the question. We are no longer only asking what sort of company should be built. We are asking what sort of person the founder must be inside that company. The panel rejects self-erasure. Growth does not require impersonating someone else. It requires accurate self-knowledge. If a founder is strong strategically but weak in execution, then the right move may be to place execution-oriented people nearby rather than to spend years trying to become a different human type.

The lecture also speaks directly to a related question: can people change? The panel's answer is yes, sometimes under profound stress or after serious life events. But the answer is immediately bounded. Founders in the early life of a company should not build their hiring logic around hoped-for future reform. Work with the person who is present, not the one who may someday emerge.

5.6.1 Question & Answer

Question. How do we surface conflict honestly, and how does a founder know when to step back rather than step in?

Answer. The first answer is about speech. Conflict becomes dangerous when the real issue remains covert. Vivian's language here is memorable: ask whether there is any "undelivered communication." The point is that an issue that remains hidden cannot be solved. Once it is made overt, it can be worked on.

$$\text{covert issue} \xrightarrow{\text{named directly}} \text{overt issue} \xrightarrow{\text{discussion}} \text{solvable issue.} \tag{5.12}$$

This is why the panel is willing to say that conflict can be good if it produces clarity.

$$\text{conflict} \to \text{clarity} \to \text{better execution.} \tag{5.13}$$

The second answer is about scale. Let F_{critical} be the set of functions that must be executed reliably if the company is to scale, and let S denote an informal indicator of organizational scalability. The lecture's late warning about founder control can be written as

$$\Big(\forall f \in F_{\text{critical}},\ \text{founder executes } f\Big) \Rightarrow S = 0. \tag{5.14}$$

Proposition 5.3. *If every critical function still routes personally through the founder, the operating model is not scalable.*

Proof. A single person becomes the bottleneck for judgment and action. As the volume of work grows, delay accumulates, frustration rises, and the founder begins to over-control. The organization then scales workload rather than capacity. This is exactly the logic the panel states in ordinary language: doing everything yourself all the time is not a scalable model. □

Hadzima's own case study makes the conflict logic concrete. A brilliant technical co-founder was causing missed schedules not because he was incompetent, but because the team felt belittled when they brought ideas to him. Information stopped moving because people stopped risking humiliation. The surface symptom was slippage. The underlying mechanism was suppressed communication.

Worked case. The repair in Hadzima's story has the structure of a short derivation.

1. Observe missed schedules.
2. Ask what is not being said directly.
3. Discover that the team feels made to look foolish and therefore stops surfacing problems.
4. Name the behavior and make the pattern visible.
5. Change the interaction so that issues can once again be raised without humiliation.

The panel is realistic about the limits of repair. Sometimes the problem is ego. Sometimes it is simply lack of awareness. Either way, radical candor is useful only if it is both honest and usable. The lecture never endorses cruelty. It endorses clarity.

5.7 Culture As A Selection Rule And A CEO Obligation

The closing movement loops back through several late questions: delegated hiring, insecurity inside the hiring team, whom the founder should trust, AI tools, self-awareness models, the brilliant-but-risky hire, the repeated question of when HR should appear, and finally the CEO's role in setting tone across a growing and even global firm. What ties these together is straightforward. The company must now scale judgment without losing its human coherence.

The late hiring questions sharpen the earlier discussion. Once managers are hiring beneath the founder, the founder must stay close enough to the process to see whether insecurity is distorting it. Group input is useful, but group indecision is expensive. The panel has seen overly inclusive hiring processes lose candidates by moving too slowly. It therefore reasserts a simple rule: gather input widely if needed, but keep clarity about who makes the final decision.

The lecture also asks whom the founder should trust. The answer is not abstract. Trust is built by observed behavior, reciprocity, discretion, and good faith. There will always be a more trusted colleague. The founder's task is to weight advice by demonstrated honesty rather than by flattery or fear.

The panel's remarks about AI are deliberately modest. Tools such as ChatGPT or LinkedIn-assisted systems can reduce the cost of drafting, search, and process setup. They can help build infrastructure cheaply and can support self-development. They do not replace judgment in hiring, leadership, or human diagnosis. The same is true of the Process Communication Model mentioned late in the session. Such tools may clarify stress patterns and communication styles; they do not eliminate the need for actual self-awareness.

5.7.1 Question & Answer

Question. Should we hire the brilliant but culturally risky person, and what kind of culture can actually scale?

Answer. The panel refuses a simple answer. Technical brilliance matters, but it is not absolute. The lecture insists on distinguishing between eccentricity and toxicity. Some highly technical people are difficult but manageable. Others carry core character flaws that will damage the team.

If T denotes technical contribution and R denotes cultural or behavioral risk, then the lecture's decision rule may be written as

$$D(T,R) = \begin{cases} \text{hire}, & T \text{ high},\ R \text{ low}, \\ \text{manage carefully}, & T \text{ high},\ R \text{ moderate and containable}, \\ \text{do not hire}, & R \text{ signals toxic character}. \end{cases} \tag{5.15}$$

The middle case matters. The panel gives two concrete reasons one might accept it. A technically exceptional person may advance the company's technology materially, and may also be worth keeping away from competitors. But the condition is strict. If the problem is mean-spiritedness, narcissism, or another deep character flaw, the hire is too dangerous for a small firm.

The late lecture then returns again to timing. Around the 25 to 50 range, outside organizational advice becomes useful; by roughly 50, a more explicit people function often becomes hard to avoid.

But when HR does appear, the earlier warning still stands: it should not simply be subordinated to finance. The people function is not a payroll echo. It is one of the instruments by which the company protects fairness, clarity, and culture.

Finally, the lecture returns to the CEO. Here the panel is emphatic. The founder or CEO has enormous power to create the organization. Culture scales when it is written, repeated, and embodied by leadership. The lecture offers several versions of this claim: the simple mission statement, the people-first founder, the leader who keeps relationships with employees rather than disappearing into hierarchy, and the global example of RSA, whose culture remained recognizable across roughly twenty countries even while each country kept local variation.

The lecture's final cultural test is therefore concrete rather than rhetorical. What scales is not abstraction alone. What scales is a mission people can remember, a set of values that can be written down, and a pattern of behavior that employees can actually observe. A company's culture is not what it says about itself once a year. It is what its people learn from how the company hires, corrects, trusts, rewards, and behaves under stress.

5.8 Summary

This session contains almost no formal mathematics in the blackboard sense, but it contains a precise operational logic. We move from an ordering of failure modes, to headcount thresholds, to trial horizons, to growth and acquisition timing, to a test for bureaucracy, and finally to a compact condition for whether the founder has actually built something scalable. Along the way the panel rejects two false oppositions: culture versus performance, and process versus agility.

The chapter's central claim is therefore the same as the lecture's opening claim, now made more explicit. A venture does not merely have people issues among its other issues. People are the mechanism through which every other issue is managed. That is why the session begins with failure, passes through culture and role clarity, turns outward to acquisition and structure, and ends with authenticity, trust, hiring judgment, and CEO-set tone. The company grows only insofar as its human system grows with it.

Chapter 6

Presenting Your Venture

Joseph Hadzima opens the evening as a two-part movement: first presenting the venture, then negotiating around it. Bob Jones takes the first part and, characteristically, does not begin with technique. He begins by justifying the subject. In a course built around practitioners, volunteers, and hard-won experience, presenting is not a cosmetic skill added after the business is understood. It is one of the places where we discover whether the business is understood at all. These notes follow that progression and keep the chapter anchored in the MIT OpenCourseWare lecture framed by Hadzima, delivered by Jones, and curated here by LazyingArt LLC.

6.1 Why Presenting a Venture Belongs in the Course

We begin where Jones begins: with a challenge. Why devote scarce course time to presenting a venture rather than to algorithms, market research, or some other more obviously technical subject? His answer is not abstract. A venture creates value only by persuading somebody to do something: buy, join, distribute, invest, stay, or eventually acquire. If we cannot describe the venture clearly enough to move one of those audiences, we have not merely failed at presentation; we have interrupted the mechanism by which the venture becomes real.

6.1.1 Question & Answer

Question. Why are we talking about presenting at all?

Answer. Because presentation sits inside value creation. Jones walks the room through the obvious and the less obvious cases. We may need to acquire customers, recruit a channel partner, hire an engineer who has another offer, reassure a current staff after bad quarters, raise capital, or one day sell the company. Even family members may become an audience if the founder is about to reject a secure job and step into a startup with only a few months of cash. The lecture is therefore not about public speaking in the abstract. It is about the repeated act of aligning another person's action with the venture's next need.

That answer is strengthened by a second distinction. Jones reminds us that the entrepreneurial

ladder has several rungs:

$$\text{Idea} \neq \text{Product}, \qquad \text{Product} \neq \text{Business}. \tag{6.1}$$

A good idea still has to survive packaging, pricing, and execution. A good product still has to become a business. And a business, if it is understood, ought to be describable. Jones's test is merciless and useful: if someone asks about our company and, half an hour later, we are still searching for words, that is already evidence against our mastery of the business.

So the chapter cannot begin with stylistic advice. It must begin with the lecturer's deeper claim: the ability to describe the venture clearly and persuasively is itself one sign that the venture has taken coherent shape.

6.2 Audiences, Objectives, and the Meaning of a Pitch

Only after that motivational work does Jones narrow the frame. A pitch, he says in effect, is not just talk. It is talk with a goal.

Definition 6.1. A pitch is a presentation addressed to a particular audience A, designed to produce some desired behavior Y, organized around what that audience wants W_A, and ending in the speaker's ask G.

We may compress his questions into a single formal line:

$$\text{Pitch} = (A,\ Y,\ W_A,\ G). \tag{6.2}$$

That notation is editorial, not original to the slide deck, but it is faithful to Jones's own sequence: whom are we talking to, why are we talking to them, what do they want, how do we give them what they want, and what do we want from them?

At this point the lecture pivots from general communication to audience-specific design. The surviving slide frame is valuable precisely because it records that pivot in mid-reveal rather than after the full list is already visible.

The visible image supports only the first two items:

$$A_{\text{slide}} = \{\text{Customers},\ \text{Potential employees}\}. \tag{6.3}$$

But Jones immediately extends the spoken list:

$$A_{\text{full}} \approx \{\text{customers},\ \text{potential employees},\ \text{current employees},\ \text{channel partners / retailers},\ \text{family members},\ \text{inves} \tag{6.4}$$

Remark 6.2. The image should not be over-read. It is evidence for the staged emergence of the audience taxonomy, not for the full taxonomy itself. The additional audiences come from the transcript, not from the frame.

This staged reveal matters. Jones is not merely listing stakeholders. He is preparing the next question: if these audiences are different, do they really share the same priorities?

6.2.1 Question & Answer

Question. Should every audience hear the same pitch?

Figure 6.1: Partial audience slide as revealed in class: customers and potential employees. The frame records an intermediate state of the argument rather than the full later list.

Answer. No. Jones asks this directly and answers it directly. The investor, the recruit, the customer, and the nervous parent are not asking the same question, so they should not receive the same framing. The underlying venture may be the same, but the order of emphasis, the benefits highlighted, the evidence offered, and the final ask must change with the audience.

6.3 The Audience's Hidden Question: WIFM

Now the lecture makes its first real compression. Jones introduces the sales acronym WIFM:

$$\text{WIFM}(A) = \text{“What's in it for me?”} \tag{6.5}$$

The point is not merely salesmanship. The point is that every audience carries a payoff function into the room. They may be too polite or too sophisticated to ask the question aloud, but they are asking it internally, and a good pitch answers it before they have to do that work themselves.

For the customer, the translation is straightforward. The customer wants a problem solved. Jones's examples are ordinary on purpose: will this app get me where I want to go quickly and efficiently? will this suit help me look right in the board meeting? The question is always local and practical. The audience wants to know whether the venture gets them what they want.

For the recruit, the question changes. A new engineer may think the technology is interesting, but that is not top of mind. Can I pay my bills? Will this company survive long enough for my effort to matter? Am I joining a place where I can make a real difference, or just admire the founder's technology from the cheap seats? The recruit is still asking WIFM, but it is a different WIFM.

The pitch therefore has an internal order. Jones circles around it repeatedly, and we may write it in a cautious reconstructed form as

$$\text{Pitch order}(A) = B_A \to R_A \to G, \tag{6.6}$$

where B_A is the benefit for audience A, R_A is the reason to believe, and G is the ask. If we broaden that one step farther back, we get the fuller narrative chain:

$$\text{Problem} \to B_A \to R_A \to G. \tag{6.7}$$

A compact diagram is useful here because the lecture itself is moving from audience taxonomy to a mechanism for tailoring the pitch:

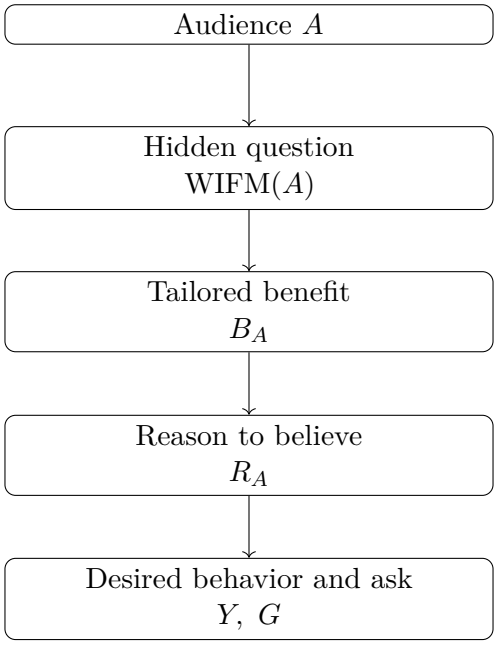

Figure 6.2: Transcript-derived pitch logic in narrow form. Once the audience is fixed, the pitch must be translated through that audience's own WIFM.

This is the chapter's main invariant. Once we lose sight of it, the audience is forced to translate for us, and the pitch begins to fail.

6.4 Baseline Pitches and What They Reveal

Jones changes tempo by making the room do the work. He announces an "equal opportunity embarrassment program" and invites volunteers to pitch. That move matters pedagogically. He does not want the rule to arrive as a clean abstraction from nowhere. He wants the class to watch the rule emerge from failure, repair, and comparison.

The first volunteer, Andres, is building a venture around generative tools for video production. Once pressed to name the audience, he chooses prospective engineers. But his three-minute pitch is about hard computer-vision problems, market growth, and technical challenge. It is not nonsense. It is simply pointed at the wrong target. The class notices at once that they have heard almost nothing

about the job itself. If we had to apply on the basis of the pitch alone, we would not know the role, the opportunity, the growth path, or even the practical reason to join.

Jones's repair is revealing. He keeps the unmet need in video editing, keeps the technology in the background as support, and pivots quickly to the recruit's benefit: come somewhere you can matter; come somewhere you are not one of 40,000 employees in a giant firm; come somewhere your skill changes the outcome. In other words, the repair works because the pitch is translated from technical description into audience-specific benefit.

The second volunteer, Subhash, brings an ultrasound venture and pitches to investors. This failure is different. Here the talk contains a real problem and a real technology, but the path from problem to business is muddy. The opening question about who likes hospitals wastes time. The user, the buyer, and the payer drift in and out of focus. The investor does not yet know who pays, how far along the venture is, whether traction exists, or why the solution is economically compelling. Again Jones intervenes, this time by giving a model pitch of his own. He starts from the burden placed on a specific high-risk patient group, moves through the inadequacy of ordinary telehealth, states the improved diagnostic capability, adds reimbursement as proof, and then arrives at the capital raise.

That is the diagnostic value of the exercise. A weak pitch is not merely "bad." It usually reveals exactly where the entrepreneur has failed to connect audience, benefit, proof, and ask.

6.4.1 Question & Answer

Question. What went wrong when the pitch did not match the audience?

Answer. In each case, the speaker made the listener perform the translation. The recruiting pitch never made the job legible to a recruit. The investor pitch made the investor wait too long for business logic. In both cases, the audience had to infer what the benefit was, who the relevant decision-maker was, and what action was being requested. Jones's critique is therefore less about eloquence than about structure.

6.5 Benefits First, Then Reason to Believe

Once the class has watched those failures, Jones becomes more explicit. The governing principle is simple and he states it more than once in different ways: benefits first, reason to believe second.

We can formalize the rule with the notation already in hand:

$$B_A = \text{the payoff the audience cares about,} \tag{6.8}$$

$$R_A = \text{the evidence that this payoff can actually be delivered.} \tag{6.9}$$

The novice entrepreneur often wants to start with R_A: the algorithm, the patent, the founders, the brilliance of the vision, the claim to be "visionary." Jones is not denying that such things matter. He is saying that they matter only if they give the listener a reason to believe the benefit.

This is why his reconstructed pitch order is

$$\text{Problem} \to B_A \to R_A \to G. \tag{6.10}$$

The listener first needs to recognize a problem, then care about the benefit, then believe that we can deliver it, and only then entertain the ask.

His examples sharpen the point from several angles. The cell-phone example is deliberately stripped down. If we begin by describing marvelous technology, the audience still has to decide whether to care. If we begin with the benefit—your phone can be instantly recharged just when the battery is dying—interest arrives before mechanism. The mechanism then returns as support.

The same logic explains the lecture's long paraphrase of Leonardo da Vinci's letter. Jones does not present Leonardo as saying, "Duke, observe how ingenious my engineering is." He presents him as saying, in effect, "Here is what I can do for your war." Bridges, siege devices, armored wagons, ways to destroy the enemy's position: benefit after benefit, all in the Duke's language. The Duke does not need to love the technology first. He needs to want the outcome.

The Nordstrom story lands the same rule in ordinary commerce. The salesman asks what work Jones does, infers what kind of suit he needs, and supplies the version of Jones he wants to become. Only afterward do cloth quality and buttonholes appear as reason to believe. The sale happens because the benefit has already been internalized.

Jones also strengthens the rule with a story principle. When he later imagines "Mabel" dreading the hospital trip, he is not ornamenting the pitch. He is linearizing it. Story pulls the audience through the causal chain in the right order. It is much easier to say

$$\text{Mabel's problem} \rightarrow \text{Mabel's burden} \rightarrow \text{better solution}$$

than to dump disconnected facts about telehealth, imaging, insurance, and device architecture and hope the audience assembles them correctly.

He adds one further qualification. Believability matters. Conviction and enthusiasm help, but not to the point of lunacy. The entrepreneur who refuses all feedback, like the salsa seller who insists the product is perfect when the customer says it is too hot, is not demonstrating confidence. He is destroying the sale. Support must support; it cannot bulldoze.

6.6 Delivery Under Constraints: Time, Memory, Voice, and Practice

The lecture now moves from content to constraint. Jones is careful here. The right structure can still fail if it arrives too fast, too long, or physically inaudible. Presentation is therefore a bounded system. It has hard limits.

The first limit is memory:

$$N_{\text{remembered}} \approx 3, \tag{6.11}$$

$$N_{\text{sentences}} = 3, \tag{6.12}$$

$$t_{\text{value prop}} \approx 10\,\text{s}. \tag{6.13}$$

Most people, Jones says, will not remember more than about three things, and usually the three things most relevant to them. That observation yields his closing challenge: express the value proposition in three sentences and ten seconds. He calls it brutal, and it is. But it gives the venture a north star. On dark nights, when the founder is not sure what exactly is being built, the ability to say it in that compressed form becomes a test of understanding.

The second limit is attention:

$$t_{\text{attention loss}} \approx 30\,\text{s}. \tag{6.14}$$

If the audience is bored or confused for much more than that, the phone comes out and we are done. Jones's advice here is not decorative. It is structural. If we have six minutes of material and only four minutes of time, we do not solve that by speaking faster. We solve it by editing.

His summary of brevity is comic and exact: be clear, be brief, and shut up.

The third limit is physical intelligibility. Jones tells a pitch-contest story in which he wins not because he had the most brilliant deck, but because he was the only speaker the judges could understand. That leads into the most practical technical instruction in the lecture:

$$d_{\text{mic}} \approx 2\text{–}4\,\text{in}. \tag{6.15}$$

These microphones are directional. We speak down the barrel, not over the top, not off to the side, and not from six or eight inches away. If we like to point, we should point with the other hand. Otherwise excitement itself becomes noise.

Finally comes practice. Here Jones sounds almost musical, which fits his own background. To look easy, one usually has to work hard. He describes giving talks to empty rooms with a timer, cutting material when the clock exposed bad editing, even taping pictures to chairs so that eye contact itself could be rehearsed. These are not theatrical extras. They are part of building new habits. The founder used to talking about substance in technical language now has to talk about benefits, story, and pace. That retraining takes repetitions.

6.7 Investor Logic and the Economics of Proof

Only after the broader pitch framework is in place does Jones specialize to venture capital. This order is important. Investors are one audience among several, but they are an audience whose logic is unusually easy to compress.

Before the meeting itself comes access. The cold call is nearly dead on arrival:

$$N_{\text{cold calls/week}} \approx 200. \tag{6.16}$$

Jones's point is not simply that venture capitalists are busy. It is that cold outreach has already failed the first filter. A warm introduction is better because it carries two hidden tests. First, is there fit? A software company may not belong in front of a healthcare-only investor. Second, will the introduction make the referrer look foolish? That is another form of WIFM. If the answer is yes, the introduction never happens.

6.7.1 Question & Answer

Question. What do investors actually want to know?

Answer. Jones reduces the situation to two large questions:

$$Q_{\text{inv}}^{(1)} = \text{How will you make money?}, \tag{6.17}$$

$$Q_{\text{inv}}^{(2)} = \text{How will I make money?} \tag{6.18}$$

Everything else is subordinate. Under $Q_{\text{inv}}^{(1)}$ sit the familiar business questions: what is broken, why is the solution better, are customers emerging, who is on the team, how will sales happen, what about manufacturing, what about IP? Under $Q_{\text{inv}}^{(2)}$ sit terms, ownership, exit, and return.

Jones is also careful about what an initial investor pitch is for. It is usually not for finishing diligence. It is for winning the next question, the next meeting, the next layer of attention.

The lecture immediately tests that logic on a live case. Rita pitches an investor-facing software-recommendation venture. Her initial move is to ask the room several show-of-hands questions about AI adoption and discovery. Jones's reaction is brutal and useful: this is a colossal waste of time. The first 45 seconds of an investor pitch are too expensive to spend on generic audience participation, especially when the speaker is already holding stronger proof. What had been buried near the end turns out to be the true opening: she knows the market, has roughly half a million followers in that space, and has already made about $80,000 in revenue. For investors, those facts are compressed forms of R_A. They say, "I know this market, I have access, and I am already seeing signal." Jones's rewritten opening therefore moves traction and credibility forward.

The lecture then returns to the ultrasound founder for a second try. This time the opening is linear and much stronger. He starts with headline market context, narrows to the high-risk segment, states the home-scan solution, and then names distinct benefits for doctors, hospitals, and insurers. Jones stops him early because the structure is already much better. The arithmetic in the classroom exchange is rough, but the order is right. The audience now sees the user, the problem, the solution, and the stakeholder-specific value in sequence.

Worked example. Jones's clearest piece of economic reasoning comes from his own catheter-lab story. The structure is worth preserving carefully because it separates user enthusiasm from payer logic.

The facts unfold in stages:

1. Doctors want the device because procedures go better.
2. Patients benefit because surgery is shorter and recovery is quicker.
3. But neither doctors nor patients pay for adoption.
4. The actual economic decision-maker is the hospital purchasing function.
5. The cath lab makes money in proportion to the number of procedures performed.
6. Turnover lag between procedures is therefore economically important.
7. The device reduces that lag enough to add one more procedure per day.

Jones summarizes the result numerically as

$$\Delta n_{\text{procedures/day}} = 1, \tag{6.19}$$

$$T_{\text{payback}} \approx 6\,\text{weeks}. \tag{6.20}$$

He does not present line-by-line arithmetic, and we should not manufacture it. But the logic is clean. "Doctors like it" is not proof of value. "The hospital earns one additional procedure per day and the device pays for itself in six weeks" is proof of value. That is exactly the kind of translation from benefit to payer-specific evidence that the lecture has been building toward from the beginning.

6.7.2 Question & Answer

Question. Who should give the pitch?

Answer. Jones answers with a distinction. Sometimes a team presentation is best. A chief technology officer may explain the technical side better than the CEO, and a chief marketing officer may explain go-to-market better than either. But he also says he has never seen an investor commit without confidence that the CEO can drive the bus. So the pitch may be distributed, but the founder is not exempt from it.

6.8 Summary

By the end of the lecture the method has been compressed almost to axioms. A pitch is not generic speech but a structured object:

$$\text{Pitch} = (A,\ Y,\ W_A,\ G).$$

The audience comes first. Every audience asks some version of WIFM. Therefore we do not begin with ourselves, our brilliance, or our technology. We begin with a problem, state the benefit in the listener's language, provide reason to believe, and then make the ask.

The later sections add the constraints that make the method operational. The audience remembers only a few things. Time cannot be cheated by speaking faster. Physical intelligibility matters. Practice is part of the construction of the pitch, not an afterthought.

Jones finally compresses all of this into a closing rule we can carry forward: know the audience, know the three things we want them to remember, show that we understand their problem, tell the story of the better solution, and practice until the delivery is economical and clear. In a venture course, that is not mere rhetoric. It is part of the venture's mechanism for becoming a business.

Chapter 7

Negotiation Skills

Once one can present a venture, Joseph Hadzima remarks, one must then engage the people who want to know more. Mindy Garber takes over exactly there. In this MIT OpenCourseWare session, and in these companion notes curated by LazyingArt LLC, negotiation is treated not as executive polish or post hoc legal cleanup, but as part of the mechanics by which a young venture either holds together or comes apart. We will follow the lecture as it unfolds: first the stakes, then the split between adversarial and collaborative models, then the central move from positions to interests, and only after that the concrete startup cases in which the method is tested.

7.1 Why Negotiation Belongs in a Startup Course

Garber opens by fixing the stakes. Ventures often fail, she says, not because the technology collapses and not because the financing is absent, but because the people cannot keep working together. That fact gives the lecture its governing motivation. Negotiation is not ornamental business etiquette. It is part of the structure by which intellectual property, effort, and opportunity are either turned into a company or wasted.

Her self-positioning matters as well. She is both engineer and mediator. The engineer gathers the facts and seeks the optimal solution. The mediator helps people have a conversation they could not yet have, so that they can understand one another well enough to solve their own problem. The lecture never abandons the engineer's desire for structure, but it refuses the fantasy that a startup is only a technical system. A venture is a technical system inhabited by people whose motives, fears, and constraints are not self-revealing.

As editorial shorthand, and not as notation visible in the lecture, we may compress a negotiation into

$$N = (\mathcal{P}, \{I_p\}_{p \in \mathcal{P}}, O, R), \tag{7.1}$$

where $\mathcal{P}$ is the set of parties, I_p the interests of party p, O the space of candidate options, and R the relationship and reputation constraints that survive the present exchange.

Definition 7.1. A negotiation is a structured attempt by a set of parties $\mathcal{P}$ to move toward an outcome in O, under the condition that each party carries interests I_p and that the relationship term R is usually not negligible.

That last clause is already Garber's world. She does not begin with the legal instrument. She begins with the people who must continue speaking, deciding, revising, and working together.

7.1.1 Question & Answer

Question. Why is negotiation a venture skill rather than a legal afterthought?

Answer. Because the venture is already negotiating before any formal legal event occurs. Co-founders negotiate motives, timing, sacrifice, roles, and ownership. Early employees negotiate culture and accountability. The company negotiates with customers, vendors, collaborators, and later with much larger organizations. If those conversations fail, the company may fail while the technology still works.

7.2 Win-Lose, Collaboration, and Reputation

From that opening Garber moves to the lecture's first conceptual split. There are many schools of negotiation, but at the level of first principles the lecture compresses them into two modes:

$$\text{mode}(N) \in \{\text{win-lose},\ \text{collaborative problem solving}\}. \tag{7.2}$$

In the win-lose model, one assumes a zero-sum game. If I gain, you lose. The relationship is peripheral. In the collaborative model, the two sides are working on a problem together. One still negotiates seriously, but one does so under the recognition that the relationship itself is part of the problem data.

Garber does not present this as mere ethical preference. The lecture's reasoning is structural. There is, she says, no such thing as one-and-done in startup life. The person across the table today may be the buyer, colleague, collaborator, or recommender tomorrow. Even if that exact individual moves to another company, the history does not vanish. In our shorthand, startup negotiation almost never lives in the degenerate regime

$$R = \varnothing.$$

More realistically,

$$R \neq \varnothing. \tag{7.3}$$

This is why reputation appears so early in the lecture. "Your reputation will precede you" is not a moral slogan tacked onto bargaining theory. It is part of the constraint set. A startup that treats every meeting as a one-shot extraction game has misunderstood the environment in which it must survive.

7.2.1 Question & Answer

Question. Is negotiation fundamentally adversarial or collaborative?

Answer. Garber's answer is that the adversarial temptation is always available, but the startup environment is structurally relational. Because the relationship persists, the faithful default is collaborative problem solving under constraint. The issue is not whether disagreement exists. The issue is whether we are willing to model the future consequences of how we handle it.

7.3 Interests, Positions, and the Expansion of Options

Only after the lecture has fixed the regime does it introduce its main analytic move. Garber asks what is, for the rest of the hour, the recurring question: what are your interests? She glosses these as needs, goals, hopes, and fears. The point is immediate. Many people cannot solve a problem because they have not yet identified what problem they are actually trying to solve.

Definition 7.2. A *position* P_i is a stated demand, preferred term, or announced answer. An *interest* I_i is the underlying need, constraint, aim, or fear that makes the position meaningful.

The lecture's core chain may therefore be written as

$$P_i \longrightarrow I_i \longrightarrow O(I_i), \tag{7.4}$$

where

$$O(I) = \{o : o \text{ satisfies the underlying interest } I\}. \tag{7.5}$$

This is the chapter's central derivation. Once the interest beneath the position becomes visible, the option set widens.

Proposition 7.3. *If a stated position P_i is only one means of satisfying an underlying interest I_i, then exposing I_i enlarges the visible solution set:*

$$\{P_i\} \subseteq O(I_i). \tag{7.6}$$

Proof. Garber's fence example is enough. A party says, "I want \$500." That is the position. Asked why, the party answers that the fence was damaged and must be repaired. The interest is not intrinsically the money; it is the repair. Once that is visible, direct repair becomes a new admissible option. The single visible candidate has been replaced by a larger set. □

Garber's examples are deliberately concrete:

$$P_{\text{fence}} = \$500, \qquad I_{\text{fence}} = \text{repair the fence}, \tag{7.7}$$

$$o_1 = \text{cash payment}, \qquad o_2 = \text{direct repair by the other side}. \tag{7.8}$$

The same logic applies in technical argument. Engineers often arrive with what looks like the correct answer already fixed:

$$P_{\text{code}} = \text{use C++}, \tag{7.9}$$

$$I_{\text{code}} = \text{minimize time to an implementable system}, \tag{7.10}$$

$$t_{\text{new language}} \approx 6 \text{ months}, \tag{7.11}$$

$$t_{\text{C++}} \approx 3 \text{ weeks}. \tag{7.12}$$

The lecture's point is subtle. "Use C++" sounds technical and definitive, but even here it may only be a position. The interest is schedule. Once that becomes explicit, we can reason about time-to-implementation rather than treating the familiar language as sacred.

We may summarize Garber's procedure in the order she teaches it:

1. Start with the visible position P_i.
2. Ask why that position is being defended.
3. Replace the declared answer by the underlying interest I_i.
4. Search over $O(I_i)$, not merely over variants of the original demand.
5. Compare options by how well they satisfy the interest.

This is why Garber can later pause, during the Sandra case, and remind the class that even "I want to raise capital" may still only be a position. The interest lies one level deeper: what is the capital for, what does it protect, what does it enable, what risk does it answer?

7.3.1 Question & Answer

Question. What is the difference between a position and an interest?

Answer. A position is the thing we say we want. An interest is the reason we want it. Garber's method is to keep asking why until the answer stops sounding like a bargaining move and starts sounding like a need, a fear, a constraint, or a goal. Only then do we see the actual option space.

7.4 Founder Negotiation: First Questions and the Circle of Interests

Having built the abstract mechanism, Garber narrows to startup life and, characteristically, makes the room work. The course is interactive. Students in the room and on Zoom are instructed to find partners and treat one another, for the moment, as cofounders. This matters. The lecture does not simply tell us the first questions; it makes the class feel what it is like to answer them under time pressure and social uncertainty.

Garber divides founder negotiation into three parts: first questions, founders agreements, and team agreements. She begins with the first questions. In compact form:

$$Q_0 = (\text{motivation},\ \text{success},\ \text{sacrifice}). \tag{7.13}$$

In words: Why are you doing this? What does success look like? What are you giving up?

The third question has special force in the lecture. A founder may be giving up graduate school or a high-status job at Google. That sacrifice is not an ornament to the story. It tells us how the person is pricing the venture psychologically, which later shapes conflict around risk, speed, control, or exit.

Garber adds a second operation before any formal agreement enters. The partner must restate the answer. This is her active-listening move. The goal is not merely information transfer but

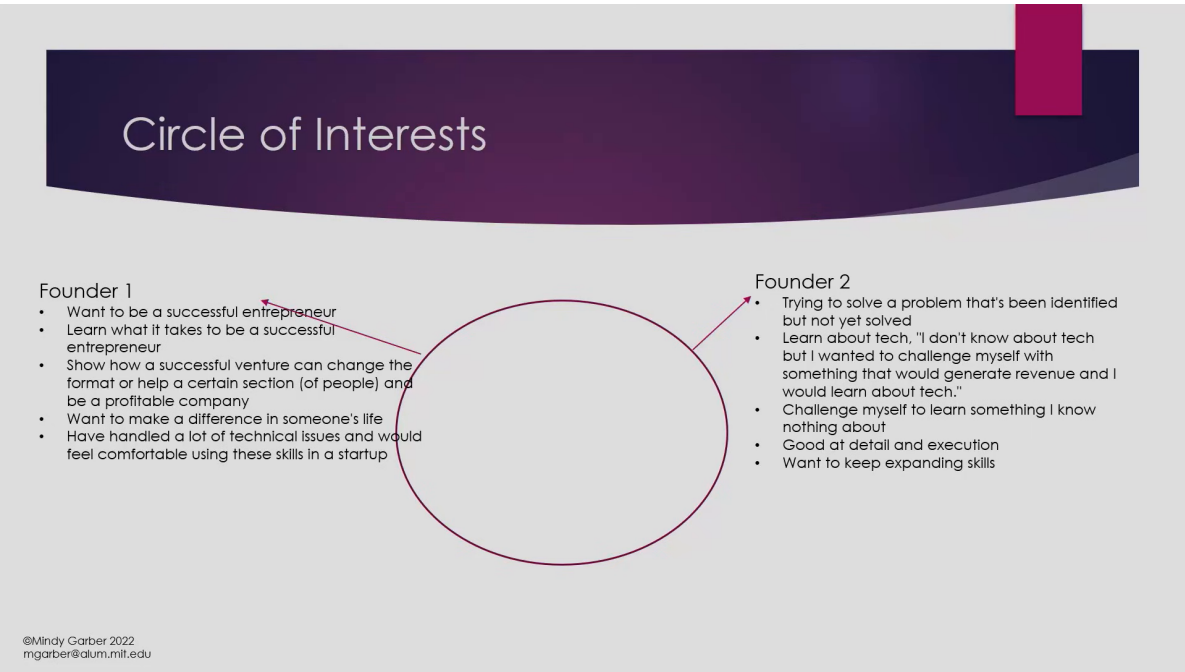

Figure 7.1: Circle-of-interests map for two founders.

evidence that both sides feel heard. Only after that does the lecture circle back to the central visual compression.

This slide is the lecture's clearest diagram, and its geometry matters. The circle in the center is blank. The content lies outside it, in the lists attached by Garber's "little sticks." One founder wants to be a successful entrepreneur, to learn what success requires, to make a difference, and to use technical skill in a startup. The other wants to solve an identified problem, learn about technology, challenge personal limits, and keep expanding detailed execution skills. The point is not to merge the lists into one slogan. The point is to place both sets of motives on one page at once.

As editorial shorthand we may therefore write

$$I_{\mathrm{F1}},\ I_{\mathrm{F2}} \longrightarrow C, \tag{7.14}$$

where C is the shared page or conversational space, not yet the company itself. The blankness of the center is part of the point. Before we decide terms, we make the motives visible.

Definition 7.4. A *circle of interests* is a one-page map that places several parties' interests around a shared discussion space so that hidden motives can be discussed before they harden into conflict.

Garber's example is revealing because the founders had a good idea and both wanted to pursue it, yet they had not actually discussed why each wanted to do so. The diagram does not decorate the negotiation. It prevents the parties from negotiating blind.

7.4.1 Question & Answer

Question. How do we make hidden founder motives visible before conflict arrives?

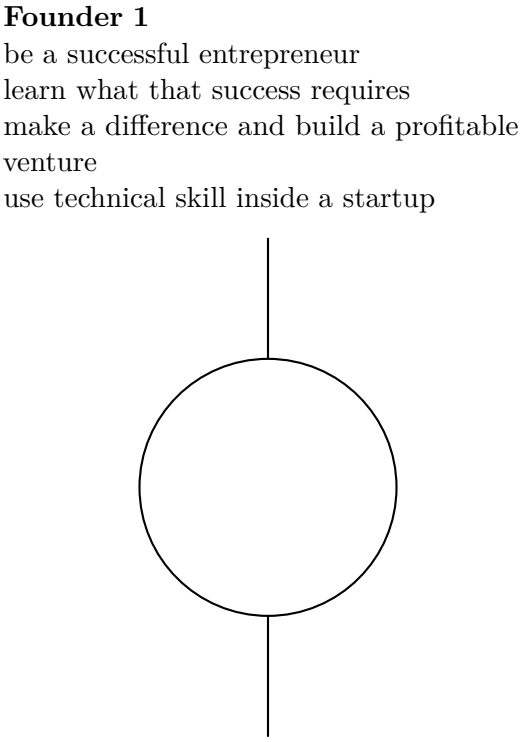

Figure 7.2: Editorial redraw of the circle-of-interests method in a taller, pocket-safe layout. The blank central circle is interpreted in the text as the shared discussion space C.

Answer. Garber's answer is procedural. Ask the first questions, restate the answers so that each side feels heard, and then write the resulting interests onto a common page. The circle does not solve the disagreement by itself; it creates the conditions under which the disagreement can be understood.

7.5 From Founder Fit to Formal Agreements

From the first questions, the lecture broadens to the thicker internal constitution of the company. This part of the hour is cumulative rather than theorem-like. Garber moves from honest assessment of cofounder fit, to founders agreements, to team agreements, and finally to the deadlock question that any equal founding pair must eventually face.

The honest-assessment portion is easily oversimplified if we read too fast. Garber is not saying that founders must begin with perfectly complementary technical backgrounds. Two mechanical engineers may found a company together. What matters is that they understand which skills are present, which are absent, and how the missing capabilities will later enter the firm.

The founders agreement then formalizes the obvious questions that people prefer to postpone: who owns what, who is contributing what, how equity will be discussed, and what happens if one founder leaves. Garber is especially insistent about this last point. One cannot safely defer the question of what occurs if the inventor departs, or if a founder who once did all the work later does less as the company grows. The lecture frames this not as moral complaint but as valuation: what are the pieces, who is contributing them, and how should those contributions be discussed honestly?

The team agreement moves from ownership to operating practice. Garber treats it as the startup's internal constitution. It covers, among other things,

- the culture the founders want to inhabit,
- norms of behavior and communication,
- meeting cadence and tools,
- roles and accountability,
- the procedure by which decisions are made,
- the language by which the team can later revisit what it once agreed.

That last item is central to Garber's view. An agreement is valuable not only because it fixes a rule, but because it gives the team neutral language for repair. We do not have to begin from accusation; we can begin from what we previously said we would do, and then ask whether it still works.

This is also where the lecture's cultural point belongs. If the founders do not define culture, the first employees will define it for them. Garber puts the matter almost constitutionally: the startup has a rare chance to decide how people will work together before habits harden.

Only after all of this does the lecture arrive at the formal deadlock question. Many cofounders begin in structural symmetry:

$$s_1 = s_2 = 0.5. \tag{7.15}$$

Garber does not reject the arrangement; she notes that it is extremely common. But she does insist that it creates a structural risk, so the team must have tools for disagreement. The classroom suggestions fall into several families: honest discussion, stepping back, consulting a neutral adviser, mediation, and distinguishing decisions that can be reversed from those that cannot. Editorially we may compress that distinction into

$$D \in \{D_{\text{one-way}},\ D_{\text{two-way}}\}. \tag{7.16}$$

A one-way decision is difficult to reverse and therefore requires slower resolution. A two-way decision can be tested and revisited.

Garber's lawyer anecdote then appears as a foil rather than a conclusion:

$$(s_1, s_2) = (0.51, 0.49). \tag{7.17}$$

Remark 7.5. The lecture does not endorse 0.51/0.49 as a theorem of startup design. It appears as a litigation-minded lawyer's response to deadlock. Garber's own aim is to give founders better tools than litigation.

Her actual advice is much closer to a decision rule:

$$D_{\text{two-way}} \Rightarrow \text{test, observe, revisit,} \qquad D_{\text{one-way}} \Rightarrow \text{slow down, deliberate, and if needed seek help.} \tag{7.18}$$

She even adds the temporal dimension: revisit these agreements periodically. Her own practice is to conduct cofounder check-ins every several months, not to ask only what the company is doing, but how the people are working together.

7.5.1 Question & Answer

Question. What should equal cofounders do when they disagree?

Answer. They should not reach first for a magical decider. Garber's sequence is more disciplined: surface the disagreement honestly, decide whether the decision is reversible, bring in neutral advice if needed, and make sure the company has already defined what its escalation path will be. Equality is workable, but only if the founders treat disagreement as a design problem rather than as a personal surprise.

7.6 The Sandra Case: A Negotiation Failure in Slow Motion

At this point the lecture slows down and becomes case-based. Garber introduces the Sandra example as a true MIT story, and she does so in a deliberately staged way: first the overview, then the separate interests, then possible agreements, and only then the reveal.

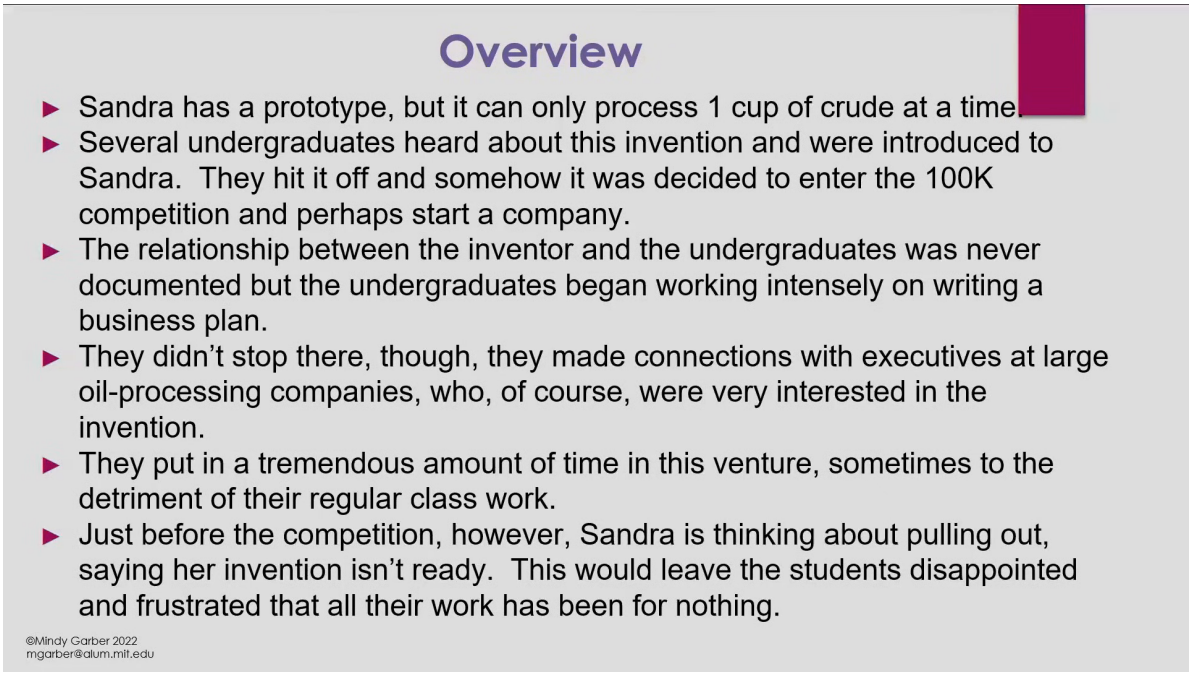

Figure 7.3: Overview of the Sandra-undergraduates negotiation case.

The slide gives the ordered structure. Sandra has a prototype, but it can process only one cup of crude oil at a time. Several undergraduates hear about the invention, meet Sandra, and move toward the 100K competition and perhaps a company. The relationship between the inventor and the undergraduates is never documented. The students then invest heavily in the business plan and

in connections with major oil-processing executives. Just before the competition, Sandra considers pulling out on the ground that the invention is not ready.

In compact form:

$$q_{\text{prototype}} = 1 \text{ cup at a time}, \tag{7.19}$$

$$A_{\text{written}} = \varnothing, \tag{7.20}$$

$$\text{student investment} = \text{already substantial}, \tag{7.21}$$

$$\text{withdrawal point} = \text{just before the competition}. \tag{7.22}$$

Garber then asks the class to do what the lecture has been building toward: identify the interests. The students may want experience, excitement, career direction, access to senior executives, resume value, or even a future job if the venture succeeds. Sandra wants to protect her intellectual property, keep control over what happens to her invention, avoid reputational harm in her field, and avoid presenting a technology as ready when scaling from one cup to industrial throughput is still unresolved.

This is where Garber carefully warns the room that inferred interests are not guaranteed to be correct. In mediation, she says, one can suggest an interest and be told immediately that it is wrong. Still, without trying to surface the interests, one has no serious way to search for agreement.

The class then generates agreements. The lecture does not collapse this into a single preferred answer, which is part of its power. Instead it shows the option space opening once the interests are named:

- Sandra could license the invention to a joint venture while remaining an advisor rather than a cofounder.
- Equity could be tied to the set of functions required for the company rather than to one undifferentiated notion of contribution.
- The IP could be protected while operational control is distributed.
- Instead of forcing an immediate startup, the parties could move first into a research-and-grants phase devoted to scaling and validation.

At exactly this point Garber pauses to refine the earlier mathematics. Even “I want to raise capital,” she says, may still be only a position. We must still ask what the capital is meant to do. The Sandra case therefore doubles as a second pass through the position-interest distinction.

Then comes the reveal: there was no agreement. More importantly, there was no real conversation about what the students wanted and what Sandra wanted. Garber had even offered to help as a neutral third party, but the students assumed the effort itself would somehow resolve the underlying structure. It did not. Sandra eventually took the invention and formed a company with a different group of people.

7.6.1 Question & Answer

Question. Why did effort, enthusiasm, and opportunity still fail to produce an agreement?

Answer. Because none of those things substitutes for negotiation. The students had labor and optimism; Sandra had control of the invention. But the parties never stabilized the relationship,

never documented it, and never tested whether their preferred future was even the same future. The case is a failure not of energy but of explicit conversation.

7.7 Negotiation Within and Beyond the Company

After the Sandra case, Garber widens the scope again. The same method survives, but the environment becomes more complicated. Once we move inside a company, negotiations may no longer be simple two-party exchanges; power becomes asymmetric, legal issues become live, and confidentiality becomes tighter.

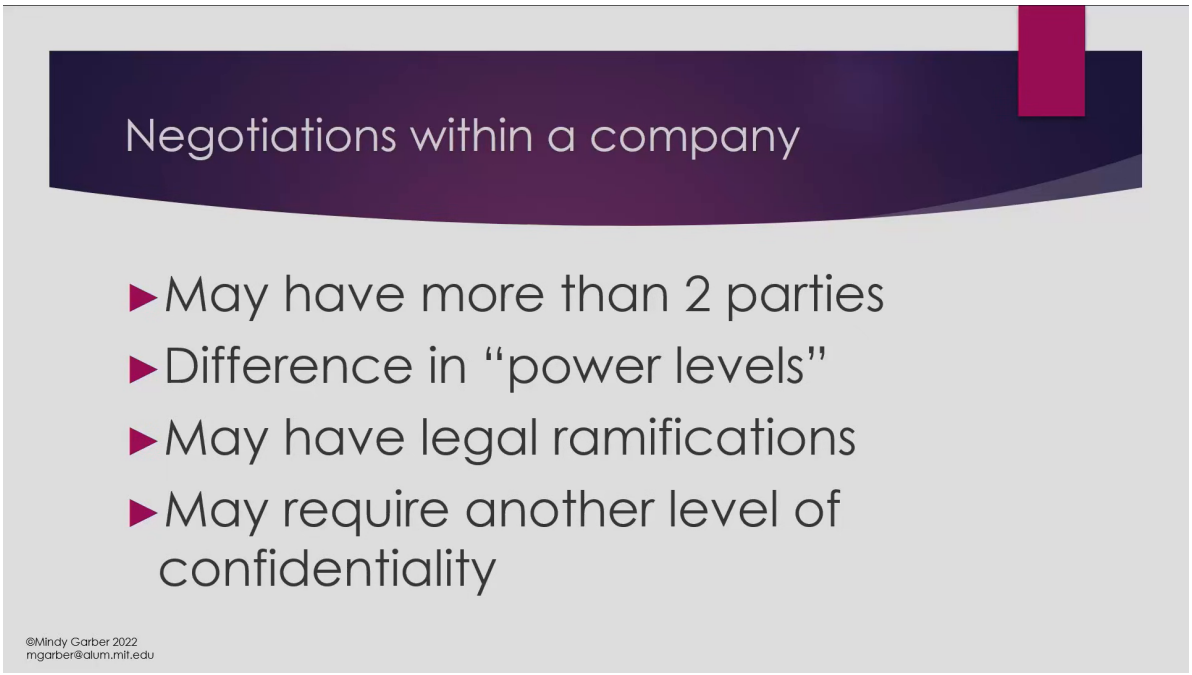

Figure 7.4: Internal-company negotiation complications.

The slide gives a clean checklist: there may be more than two parties, differences in "power levels," legal ramifications, and a further level of confidentiality. Editorially we may write

$$n > 2 \tag{7.23}$$

and

$$N_{\text{internal}} = (\mathcal{P}, \{I_p\}_{p\in\mathcal{P}}, O, R, L, K), \tag{7.24}$$

where L records legal implications and K records confidentiality constraints.

Garber immediately adds the lecture's recurring refrain: even here she still draws her little circle of interests. In other words, the model does not change; the complication set grows.

The Judy case. The first application is internal. Judy wants a promotion to marketing manager. Garber's method is stakeholder decomposition:

$$S_{\text{Judy}} = \{\text{Judy},\ \text{manager},\ \text{company},\ \text{teammates}\}. \tag{7.25}$$

Judy wants challenge, recognition, greater responsibility, supervisory opportunity, and a chance to use newly acquired skills. Her manager wants to appear capable of developing employees but also needs the existing work completed. The company wants revenue, effective leadership, employee growth, diversity, and successful projects. Teammates may want promotion opportunities of their own.

Garber then helps Judy move from diffuse aspiration to structured preparation. Judy receives a list of marketing functions, puts them into a spreadsheet, and assigns weights:

$$w_j \in \{1, 2, 3\}, \qquad \text{with larger } w_j \text{ indicating greater importance.} \tag{7.26}$$

The spreadsheet is not the theorem; it is the mechanism by which a vague ambition becomes an ordered plan. Judy then walks into the conversation not merely saying, "Promote me," but showing how she has already decomposed the role and thought through its execution. The promotion succeeds because the request is framed in terms of other people's interests as well as her own.

The Little Inc/Big Inc case. The last major example takes the same method outside the company. Little Inc has a fixed-cost contract with Big Inc. The contract is vague; Garber uses the episode to remind us that even a 300-page contract never covers everything. Big Inc later cancels the contract, not because Little Inc did poor work, but because Big Inc is having wider financial problems and is canceling contracts across the board.

Little Inc does not begin by arguing tactics around the table. Garber says that this got them nowhere. Instead the team stops and enumerates stakeholders before the meeting:

$$S_{\text{contract}} = \{\text{sales executive},\ \text{service manager},\ \text{CFO},\ \text{Big Inc contacts}\}. \tag{7.27}$$

The interests are different even on Little Inc's side. The sales executive worries about commission and follow-on business. The service side worries about competence, reputation, and the customer relationship. The CFO worries about cash flow, revenue recognition, fairness, and not being treated as a trivial target by a much larger company. Garber adds concrete asymmetry here: Little Inc is a roughly 20-person firm, while Big Inc has on the order of 10,000 people. The team also agrees on two negative interests: they do not want escalation to the CEO, and they do not want lawyers brought in.

Big Inc's side must also be modeled. The company may worry about legal exposure, reputation, internal blame, and the standing of the people who signed the original agreement. The crucial lecture point is that these are not guaranteed facts; they are hypotheses to be prepared in advance and then updated in the room.

This is the setting in which Garber introduces what she loosely calls a solution space or zone of solution:

$$Z = \{a : a \text{ is acceptable under the parties' interests and constraints}\}. \tag{7.28}$$

The derivation is practical and exact:

1. enumerate the stakeholders on our side and the other side;
2. identify substantive, reputational, legal, and escalation interests;
3. ask what outcomes remain fair, survivable, and relationship-preserving;
4. construct Z before the meeting begins;

5. evaluate the incoming offer against Z, not against the vanity of bargaining harder.

This is how Garber reaches one of the lecture's most memorable claims. Big Inc makes an opening demand. Little Inc's CFO immediately accepts it. The room is startled. Garber's point is that the real negotiation had already happened before the meeting. The first offer was better than expected, it lay inside the acceptable region, it preserved the relationship, and it avoided escalation. Therefore

$$a_{\text{first}} \in Z \Longrightarrow \text{accepting may dominate further bargaining.} \tag{7.29}$$

This is the practical content of the 3M maxim Garber quotes: sometimes the first offer is the best offer.

7.7.1 Question & Answer

Question. Should we explicitly share our interests with the other side?

Answer. Garber's answer is deliberate rather than naive. Interests can be weaponized if we disclose them foolishly. But the lecture's closing stories show that selective, problem-oriented disclosure can change the entire meeting. In the acquisition-style anecdote, one side begins by naming the other side's likely interests, and the conversation opens. In the IBM–3M story, progress begins only when one side finally says what it is trying to accomplish. Editorially we may write

$$I_{\text{shared}} \subseteq I_{\text{all}}, \tag{7.30}$$

where I_{shared} contains the interests that help define the joint problem, not the vulnerabilities that simply weaken our position. Garber's own rule is plain: share what helps move the problem forward.

The final exchange of the lecture adds a useful refinement. If unexpected people appear in the room, their interests may differ from the generic interests of the company they represent. In that case the right move is often to ask directly: what matters to you, what are you trying to accomplish, what would be helpful? If they will not answer, then the feasible agreement space becomes harder to map. If they do answer, one may discover exactly the kind of low-value/high-value asymmetry from which agreements are built.

7.8 Summary

Garber closes the lecture by compressing it back into a handful of durable practices. Understand the interests of your cofounder. Write decisions down and revisit them because startups change quickly. Prepare, prepare, prepare before important negotiations. Then listen in the meeting, because new interests may appear there that were invisible in advance.

That closing summary is entirely consistent with the chapter's internal logic. We began with the claim that startup failure is often a failure of relationships. We then saw the lecture's practical mathematics unfold step by step:

$$\text{parties} \longrightarrow \text{interests} \longrightarrow \text{options} \longrightarrow \text{agreement under relationship constraints.}$$

Hadzima's final anecdote reinforces the same lesson from the far side of the table: sometimes the other side is not yet internally organized enough to negotiate coherently. Interests must be surfaced there as well. That is why Garber's method belongs in a startup course. It is not merely about getting a better deal. It is about learning how the venture thinks before it breaks.

Chapter 8

Financial Projections

In Joseph Hadzima's MIT OpenCourseWare course, Steve Durzynski takes over just after the financing panel and asks the next unavoidable question: once we know the customer, the business model, the team, and the legal form, what resources does the venture actually need in order to live? This companion chapter, curated by LazyingArt LLC, follows that sequence closely. We move from the motivational turn – why projections matter at all – into venture arithmetic, the standard financial statements, the bottom-up construction of a model, the end-state four-year profit-and-loss view, the quarterly cash-flow recurrence, and only then the late pivot to equity, dilution, vesting, and control. The lecture's discipline is simple to state and difficult to evade: the model is not a ceremonial prediction, but a device for forcing assumptions, costs, timing, and risk into the open.

8.1 Why We Need Financial Projections

The lecture opens with a handoff, and the handoff matters. We have already discussed customers, business model, team formation, and legal structure. Only now do we ask what resources the venture will actually require. Durzynski also takes a moment to establish his stance. He is not presenting himself as a pure accountant; he is an engineer and repeat founder explaining why even an engineer must learn to love the model. This is important because the lecture begins from resistance, not from compliance.

That resistance is familiar. The founder says: I am too busy for this; the numbers will be wrong anyway; the investors know every chart goes up and to the right; why not let an accountant do it later? Durzynski does not wave this away. He reverses it. Precisely because the venture is uncertain, the founder needs the model. The model is not introduced as paperwork. It is introduced as a survival instrument.

His governing metaphor is exact and memorable: cash is the oxygen of the venture. We may survive without many things for a while, but not without oxygen. Once that point is accepted, the model ceases to be decorative. It becomes the scorecard for the firm's condition, the roadmap for future spending, and the vehicle by which outside capital is obtained. Investors, lenders, and financially trained listeners want more than upside. They want evidence that we understand the business well enough to describe assumptions, drivers, milestones, and risks.

A compact way to express the operating pressure is through burn:

$$\text{BurnRate} = \text{MonthlyOperatingLoss} + \text{CapitalExpenditures}. \tag{8.1}$$

This is not yet a complete accounting system, but it is already managerial arithmetic. Once we know what we lose each month and what capital equipment we must buy, we are much closer to the central practical question: how long can we keep breathing?

The lecture insists on something else as well. This understanding cannot be outsourced in any meaningful early-stage sense. One may later hire accountants and eventually bring in a CFO, but the founder has to develop a direct feel for the business. What is the average selling price? What is the gross margin? How expensive is the research team? What happens to spending once selling begins? Which milestone failure becomes a cash failure? Durzynski's point is that a good spreadsheet is valuable not because it predicts the future with precision, but because it forces us to ask the right questions with precision.

8.1.1 Question & Answer

Question. Why do financial projections matter if they are going to be wrong anyway?

Answer. Because the point is not to be right in every detail. The point is to know what would have to be true for the business to survive and grow. A model that is later revised can still be useful if it has already forced us to identify cost drivers, hiring needs, cash burn, milestones, and downside risk. A founder who refuses to model is not avoiding error; he is avoiding contact with the structure of the business.

8.2 Investor Arithmetic Before Accounting

Only after cash has been established as the binding constraint does the lecture pivot into the numbers investors care about. This is the right order. Before we learn the full bookkeeping machinery, we first need to understand what kind of business can support venture capital at all.

A venture fund has obligations to its own limited partners. That is why not every good small business is venture fundable. A company that grows too slowly may be an excellent business for its founders and still fail the venture test. The point is arithmetic rather than moral judgment. Durzynski gives the lecture's clean rule of thumb:

$$(1+r)^5 = 4 \qquad \Rightarrow \qquad r \approx 0.32. \tag{8.2}$$

A four-times return in five years corresponds to an internal rate of return of roughly 32%. The lecture uses this as a heuristic, not as a theorem, but the operational meaning is clear: if the slope of the future curve is too shallow, the company may simply be in the wrong financing category.

The time horizon is also part of the arithmetic. A series A investor typically wants a return in something like three to five years, while earlier investors often wait longer. The lecture is careful here: sectors such as medical devices may legitimately require longer plans because regulation and time-to-market are longer. But the underlying logic is unchanged. We are always being asked what eventual path makes the investment worthwhile.

The ownership formula is then stated in its simplest form. If an investor puts in I dollars at a pre-money valuation V_{pre}, the post-money valuation is

$$V_{\text{post}} = V_{\text{pre}} + I, \tag{8.3}$$

and the fraction sold in that round is

$$f = \frac{I}{V_{\text{pre}} + I}. \tag{8.4}$$

Worked example. If we raise \$1M at a \$4M pre-money valuation, then

$$f = \frac{1}{4+1} = \frac{1}{5} = 20\%. \tag{8.5}$$

The investor owns one-fifth of the company after the round, and the post-money valuation is \$5M.

Durzynski then turns immediately to dilution, because ownership in one round is only the beginning of the story. In his example,

$$f_A = \frac{5}{5+5} = 50\%, \tag{8.6}$$

$$f_B = \frac{10}{15+10} = 40\%. \tag{8.7}$$

The second round dilutes the holders from the first round. If the first investor held 50%, then after the later 40% round that stake becomes

$$\omega_{A,\text{after } B} = 0.50(1 - 0.40) = 0.30, \tag{8.8}$$

$$\omega_{\text{VC,total}} = 0.30 + 0.40 = 0.70. \tag{8.9}$$

So after the two rounds the venture investors together own about 70% of the business.

Now the lecture makes its most important arithmetic move. The problem is no longer only dilution; it is scale. If investors have put in \$15M and want something like 4× to 6× back, then their stake must eventually be worth about \$60M to \$90M. Using the lecture's 70% ownership figure, the implied exit value is on the order of

$$V_{\text{exit}} \approx \frac{60\text{M}}{0.70} \text{ to } \frac{90\text{M}}{0.70} \approx 86\text{M to } 129\text{M}. \tag{8.10}$$

The transcript is unstable around the exact spoken number, so one should read this as cautious reconstruction rather than exact quotation. The key point is that after such financing the company must plausibly reach tens of millions of revenue and a correspondingly large exit value. The more money raised, the higher the bar.

This is why Durzynski is suspicious of the celebratory tone around fundraising. A round is not merely a badge of honor. It is a claim on the future. He says this several ways: the bar keeps moving up; more capital means more has to be achieved; and the right amount of funding is therefore a hard balancing act rather than a simple maximization problem.

8.2.1 Question & Answer

Question. If raising money sounds like success, why does it also make success harder?

Answer. Because the round buys runway and reduces immediate risk, but it also installs a larger future obligation. New investors own a fraction of the company and expect that fraction to be worth a large multiple later. So the absolute size of the eventual outcome must rise with the capital raised. A press release may celebrate \$5M or \$15M raised; the lecture asks us instead what revenue scale and exit value now have to exist for that financing to have made sense.

8.3 Reading the Business Through Standard Financials

Only now do we get the standard set of financial statements. The order is pedagogically deliberate. Durzynski wants us to see formal accounting as the language in which the business model becomes visible, not as a separate discipline dropped on top of it afterward.

Income statement. The income statement, or profit-and-loss statement, is the part everyone thinks they know: show revenue, subtract cost of goods sold, and compute what remains. In the lecture's simplified operating language,

$$\text{GrossMargin}_t = R_t - \text{COGS}_t, \tag{8.11}$$

$$\text{OperatingProfit}_t = R_t - \text{COGS}_t - \text{DeptExp}_t. \tag{8.12}$$

Here R_t is revenue, COGS_t is cost of goods sold, and DeptExp_t is the rollup of departmental expenses: sales, marketing, R&D, G&A, staffing, and the rest of the operating structure.

Balance sheet. The balance sheet carries assets, liabilities, and equity. Cash is an asset, but so are patents, equipment, and other property. Debts sit on the liability side. Equity is the residual ownership claim. This becomes essential later, when the class turns to the cap table and Durzynski explains that equity ownership belongs to this side of the firm's description, not to the revenue line.

Cash-flow statement. The cash-flow statement is where money actually enters and leaves. Durzynski treats it more visually than doctrinally: money comes in, money goes out, and the level rises and falls. That is already enough to see why profit is not identical to survival.

The lecture then shifts from absolute dollars to normalized percentages. Boards and investors often compare firms through ratios because the scale of firms changes rapidly but the percentage signature of a business model is often recognizable. We therefore define

$$\text{GrossMarginPct}_t = \frac{\text{GrossMargin}_t}{R_t}, \qquad \text{OpMargin}_t = \frac{\text{OperatingProfit}_t}{R_t}. \tag{8.13}$$

Different businesses leave different ratio fingerprints.

This is why Durzynski shows old-school hardware and software examples. A hardware company with very low R&D and low operating profit tells one story; a software firm with higher R&D, higher SG&A, and higher eventual margins tells another. Joe Hadzima's intervention sharpens the point: if our numbers differ greatly from successful comparable businesses, then either we do not understand the business model or we must explain why ours is economically different.

The lecture's rules of thumb are therefore not identities but comparison standards, especially for the year-four view:

$$\frac{\text{R\&D}}{R} \approx 10\%\text{–}20\%, \qquad \frac{\text{G\&A}}{R} \approx 5\%\text{–}15\%, \qquad \frac{\text{OperatingProfit}}{R} \approx 15\%\text{–}20\%. \tag{8.14}$$

Durzynski is careful to say that these are not meant for the earliest chaotic quarters. Year one in a startup is often too distorted by ramp-up and loss. But by the year-four view the ratio structure should begin to resemble something a pattern-matching investor can recognize.

There is also a practical asymmetry in the lecture's treatment of margins. Operating expenses can sometimes be squeezed. Gross margin is harder to repair because it is usually baked into the product or service itself. That is why the lecture treats gross margin as a particularly revealing number.

8.4 Building the Model From the Bottom Up

Having taught us how to read the business, the lecture now turns to how we construct it. The pace becomes more practical. We are no longer looking at statements from above; we are asking what cells, hires, prices, and contracts have to be entered to make those statements emerge.

The model begins with very concrete questions. What exactly is the product or service? What price will it command? What is the unit cost of manufacturing or delivery? What overhead is required to support it? How is it sold: direct, through partners, through distributors, or through an internal sales force that does not yet exist?

One of the lecture's clearest small examples contrasts direct sale and distributor sale. If the nominal sale price is 100, then in the direct case the company books the full sale price as revenue:

$$\text{BookedRevenue}_{\text{direct}} = 100. \tag{8.15}$$

In the distributor case, the distributor takes 20 off the top, so the company books

$$\text{BookedRevenue}_{\text{distributor}} = 100 - 20 = 80. \tag{8.16}$$

The tradeoff is immediate. Revenue booked is lower, but so is the burden of direct selling. The spreadsheet is useful precisely because such strategic choices show up immediately in the numbers.

The lecture is equally plain that in the earliest phase the founders themselves are often the first salespeople. That is more than folklore. It means the model cannot jump from zero to scale by magic. We have to ask how the first customer is obtained, then the second, then the tenth. We have to ask what the sales cycle is, how much support must be promised, whether support contracts renew, and how long it takes a new salesperson to become productive.

Durzynski then opens the KPI layer. Churn, committed monthly recurring revenue, customer acquisition cost, lifetime value, and the ratio LTV/CAC are not presented as decorative dashboard numbers. They are mechanisms by which a business may decide whether injecting more capital into customer acquisition actually creates value. In one audience exchange, this becomes a useful two-box distinction. There is an operating box and an investment box. An operating dollar might be modeled very roughly as

$$1 \mapsto 1.10, \tag{8.17}$$

while a venture investor thinks in far larger equity multiples. The boxes are different, but linked: a better operating machine can make the equity stake much more valuable.

Staffing then enters as the dominant driver of departmental expense. Durzynski uses simple operating heuristics:

$$\text{Benefits} \approx 0.15 \times \text{Salary}, \qquad \text{DeptExp}_{\text{staff}} \approx 0.67 \times \text{TotalExpenses}. \tag{8.18}$$

Most early burn is people. That is why the lecture spends real time on salary determination, minimum viable founder salary, and market-based compensation for key hires. A founder should not starve, but should not use newly raised capital as personal extraction. The lecture's language is memorable: investors do not want you to starve, but they do want you to be hungry.

There is also an instructive exchange about whether a single product-level NPV or positive-value argument is enough. The lecture answers no. A positive product unit economics story is only a kernel. Around it must sit infrastructure, rent, hiring, marketing, manufacturing, and the rest of the corporation. In the very earliest sketch, average-cost approximations may be fine. But once real capital is being raised to be spent on real people, the model must become granular.

Durzynski is equally sharp about tools. Business-planning software is not recommended as a substitute for understanding, and AI is treated cautiously. It may help us brainstorm KPIs or strategic variants, but it cannot relieve us of the obligation to understand how the spreadsheet works. If we cannot explain why the cells are what they are, then the model is not really ours.

Finally, the lecture returns to the central forecasting discipline: build sales from the bottom up. Start with one sale, then two, then the cadence of a real sales process. Do not begin by declaring that there are billions of people in the world and we will take 1% of them. Yet bottom-up does not mean blind. There must still be a market-size reality check. If the bottom-up forecast implies owning more than the market, the model has failed in a different direction.

8.4.1 Question & Answer

Question. Why is a bottom-up sales forecast more credible than a large market-share story?

Answer. Because a bottom-up forecast exposes mechanism. It forces us to say who buys, at what price, through which channel, on what timetable, with what staffing, and at what cost. A large-market story usually hides all of that beneath a percentage of an enormous total. Investors do not fund market size; they fund sales machinery. A defendable forecast is one whose gears can be named.

8.5 From Assumptions to the Four-Year P&L

The lecture now makes one of its strongest structural moves: start with the end in mind. Instead of opening every worksheet tab first, Durzynski shows the finished four-year profit-and-loss output. This is not a shortcut. It is a way of fixing our attention on the visible outcome investors care about before we descend into the machinery that produces it.

The slide is a summary chart rather than a derivation board, and that is precisely why it is so useful. Revenue rises steeply across the four years. Operating profit begins below zero, approaches the axis by Year 2, and is clearly positive by the later years. The lecture uses this picture to condense several thoughts into one visual object: top-line growth, delayed profitability, and the finite amount of cash that must be survived before break-even.

This picture allows Durzynski to make a very useful inference. If we sum the negative operating results before break-even, we estimate how much cash the company needs in order to survive until profitability. The lecture gives a concrete arithmetic version:

$$\text{CashNeed}_{\text{pre-BE}} \approx 3.3\text{M} + 2.6\text{M} \approx 5.9\text{M} \approx 6\text{M}. \tag{8.19}$$

The chart suggests the same thing geometrically:

$$\text{CashNeed}_{\text{pre-BE}} \approx \int_{\text{OperatingProfit}(t)<0} -\text{OperatingProfit}(t)\, dt. \tag{8.20}$$

The lecture uses this only as intuition – an “integral under the curve” way of seeing cumulative early losses – not as a claim that the slide itself gives precise continuous data. But the conceptual move is exactly right. The losses before break-even are the cash burden that must be financed.

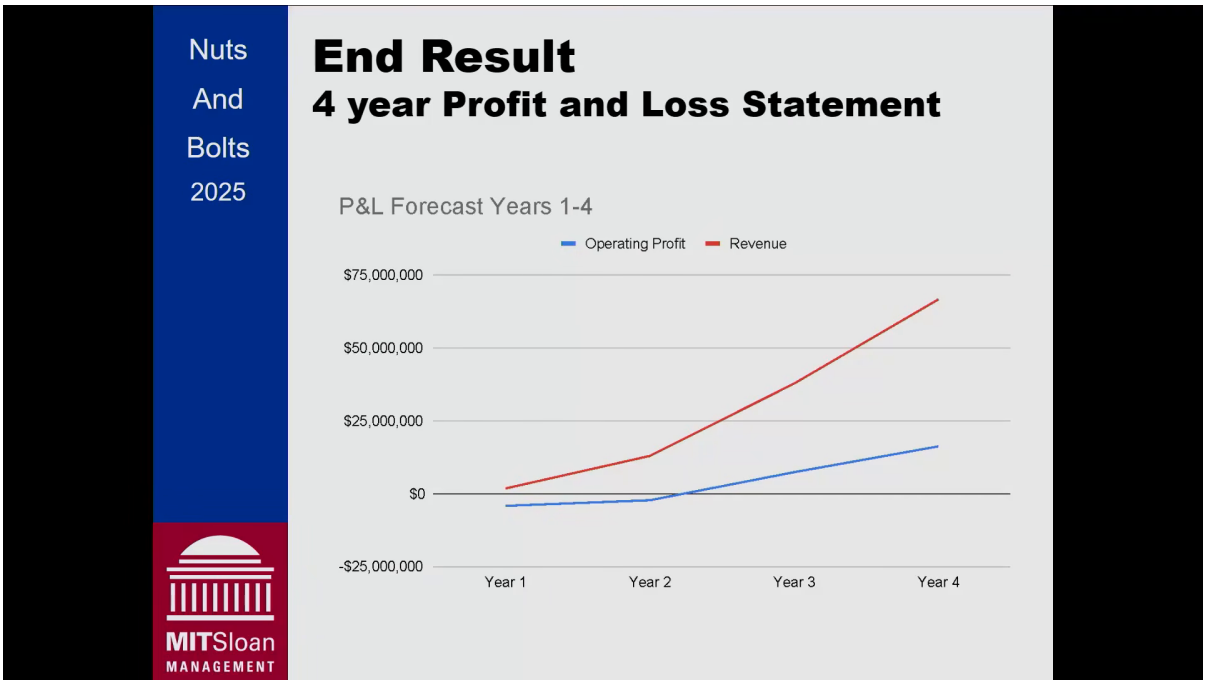

Figure 8.1: Four-year P&L forecast with revenue and operating profit.

This is why Durzynski does not stop at the minimum. If the curve suggests $6M of pre-break-even need, he still recommends raising more, perhaps $10M, so there is room for delays and mistakes. The model is not only an argument for future profitability. It is an estimate of how much uncertainty the venture can survive.

The lecturer also reads the slope of the red revenue curve for a rough sense of growth quality. The transcript is unstable around the term he uses there, and one should not formalize it too aggressively. The safe point is simpler: the chart is not read only for whether profit becomes positive. It is read for how quickly the business seems to scale once it begins to move.

8.5.1 Question & Answer

Question. If the company loses money in the first years, what exactly makes the model credible?

Answer. Not the smoothness of the graph by itself. Credibility comes from the path underneath the graph: units sold, selling price, channel choice, support retention, manufacturing cost, hiring plan, rent, overhead, and timing. Early losses are acceptable only if the route from those losses to later profitability can be explained in operational detail. The graph compresses the story; it does not replace the story.

8.6 Cash Flow, CAPEX, and the Timing of Survival

After the annual and quarterly P&L rollups, the lecture pivots into the cash-flow view. This is where the spreadsheet becomes dynamic. We are no longer asking only whether the business looks

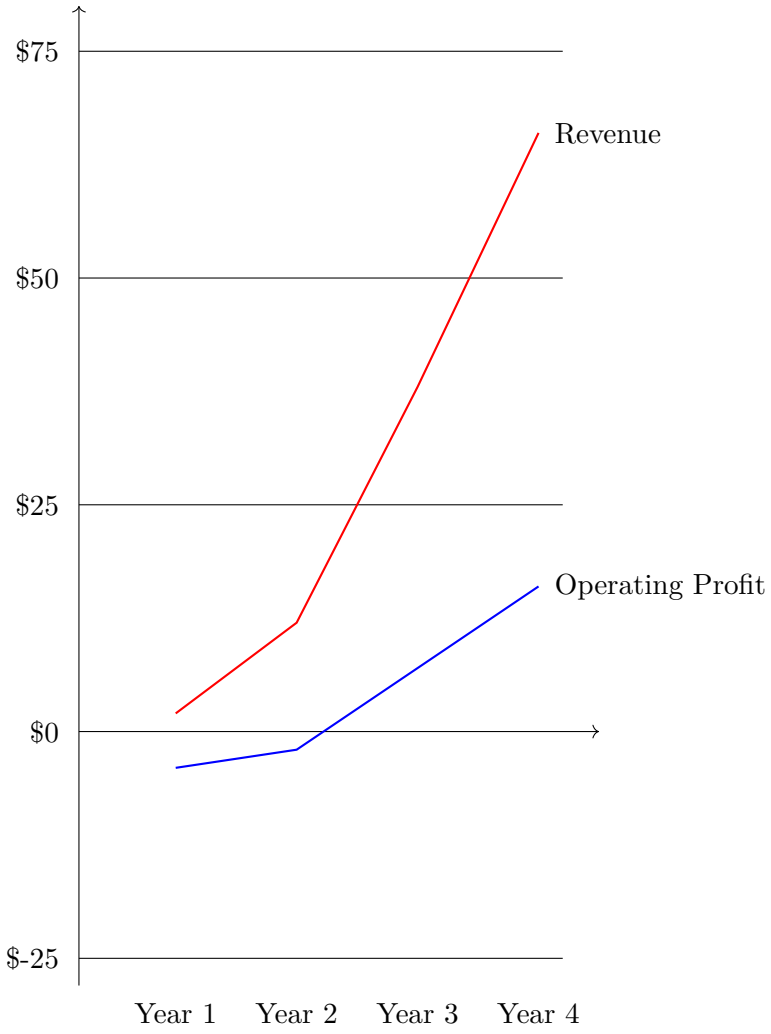

Figure 8.2: Qualitative redraw of the four-year forecast. The geometry is schematic rather than numerically exact.

profitable eventually. We are asking when cash actually runs low, when the next financing must happen, and how capital expenditure changes the path.

The slide teaches several things at once. First, it gives the row order of the cash model: beginning cash, investment, total revenue, COGS, departmental expenses, capital expense, change in cash, ending balance. Second, it shows that the workbook is modular: lower tables feed upper tables, and the color scheme indicates how the information travels. Third, it gives the reader a practical legend: red means user input, magenta means a link to another sheet, black means calculated value, and blue means a value pulled in from another sheet.

The bookkeeping spine of the slide can be written cleanly as

$$\Delta C_q = I_q + R_q - \mathrm{COGS}_q - \mathrm{DeptExp}_q - \mathrm{CapEx}_q, \tag{8.21}$$

$$C_q^{\mathrm{end}} = C_q^{\mathrm{beg}} + \Delta C_q, \tag{8.22}$$

$$C_{q+1}^{\mathrm{beg}} = C_q^{\mathrm{end}}. \tag{8.23}$$

This is the recurrence Durzynski is gesturing toward when he describes the cash level going up and

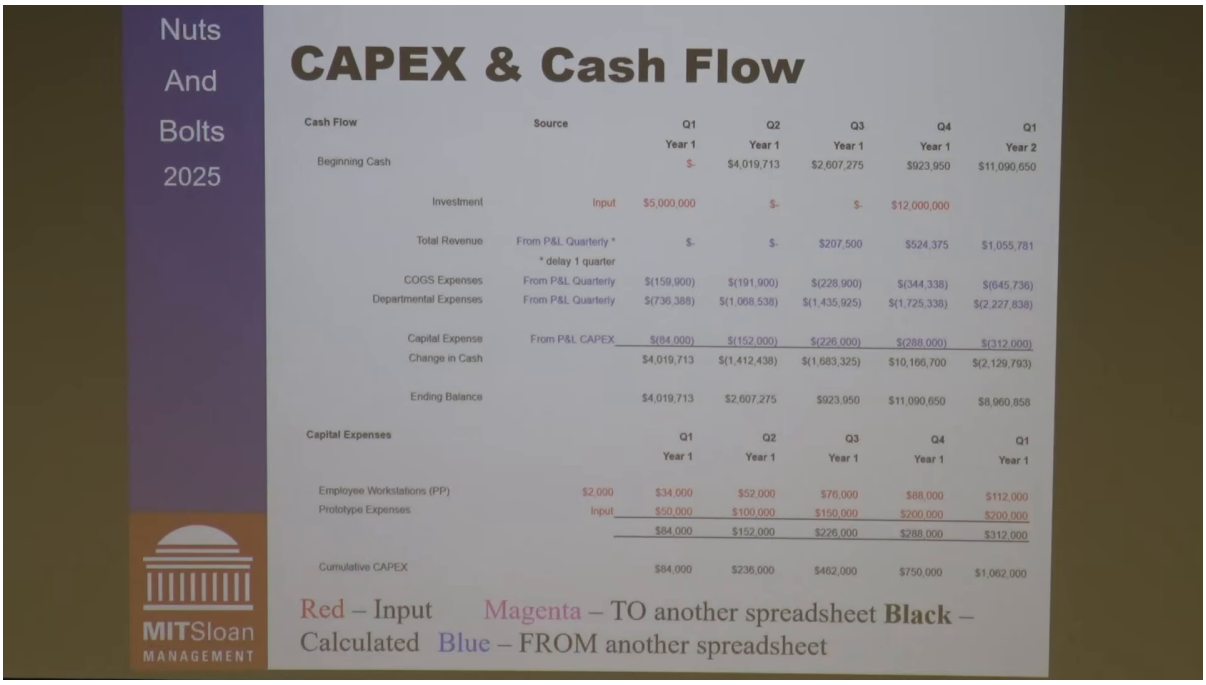

Figure 8.3: CAPEX and cash-flow spreadsheet layout.

down quarter by quarter. We begin the quarter with cash, add financing and revenue, subtract operating costs and capital spending, and carry what remains into the next quarter.

The lower capital-expense block explains where the capital line comes from. In abstract form,

$$\mathrm{CumCapEx}_q = \sum_{\tau \le q} \mathrm{CapEx}_\tau. \tag{8.24}$$

A few entries in the screenshot are legible enough to illustrate the running sum:

$$(\mathrm{CapEx}_{Q1}, \mathrm{CapEx}_{Q2}, \mathrm{CapEx}_{Q3}, \mathrm{CapEx}_{Q4}, \mathrm{CapEx}_{Q1'}) = (84, 152, 226, 288, 312) \times 10^3, \tag{8.25}$$

$$(\mathrm{CumCapEx}_{Q1}, \mathrm{CumCapEx}_{Q2}, \mathrm{CumCapEx}_{Q3}, \mathrm{CumCapEx}_{Q4}, \mathrm{CumCapEx}_{Q1'}) = (84, 236, 462, 750, 1062) \times 10^3. \tag{8.26}$$

These numbers should be read only as clearly visible examples from the slide, not as a full spreadsheet transcription. Their purpose is to show the cumulative logic, not to replace the original evidence.

The cash-flow slide also hints that revenue collection may lag the P&L by a quarter. The text is not legible enough to formalize that timing in full detail, but the existence of such lag is exactly why cash flow is not reducible to the annual income statement.

Durzynski's spoken walkthrough makes the recurrence concrete. Start with zero. Add \$5M of investment. Spend through the quarter. End with roughly \$4M. Carry that amount into the next quarter. Spend again. If the balance falls toward a dangerous floor, then the next raise has to happen before the floor is reached. In the lecture's example, that means raising more money before the company drops below a comfortable cash threshold.

This is why cash flow is operationally more useful than many founders first assume. Some conversations, especially high-level investor conversations, will focus on the income statement. But survival

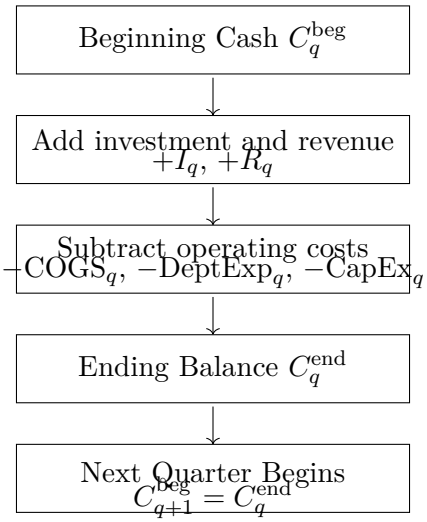

Figure 8.4: Cash-flow recurrence implied by the workbook structure.

is a timing problem, and the cash-flow table is where timing becomes explicit. The model is no longer merely asking whether the business can become profitable someday. It is asking whether the company can stay alive long enough to reach that someday.

8.7 Presentation Discipline and the Late Pivot to Equity

Once the model exists, the lecture turns to presentation. The first principle is that the model is iterative. It should be beaten up by friends, advisors, and investors. The point of that beating-up is not humiliation. The point is to discover which assumptions break first, and what has to be changed if they do.

At the executive-summary level, the lecture recommends an annual P&L with percentages so that investors can pattern-match the business quickly. In deeper due diligence, one may need the full structure: annual P&L, quarterly P&L, staffing plan, quarterly cash flow, and the assumptions behind pricing, hiring, profitability timing, and cash need. Large firms may bring in finance professionals of their own to push on the weak points of the model.

The lecture is equally disciplined about what not to show too early. One should know the biggest risks and have backup material ready. One should model what happens if milestones slip or key hires are delayed. But one should not lead the pitch deck with best-case and worst-case scenarios. The founder is being asked to show conviction in a base case, and then to defend that case under questioning. Risk scenarios belong more naturally in the appendix or in live discussion.

Only after this financial spine is in place does Durzynski pivot to equity. This ordering matters. The pie is not split in a vacuum. It is split after we understand what sort of enterprise is being financed and what kind of work still lies ahead.

The algebra of the cap table is simple:

$$\omega_i = \frac{s_i}{S}, \tag{8.27}$$

where s_i is the share count of holder i and S is the total fully diluted share count. The lecture then illustrates, through audience questions, how easy it is to recover valuation from a percentage.

If \$0.5M buys 10% of the company, the post-money value is about \$5M. If \$5M buys 50%, the post-money value is about \$10M. If \$10M is raised on a \$15M pre-money valuation, then the post-money value is \$25M. These are just late-stage reappearances of the same financing arithmetic from earlier in the lecture.

The lecture then adds vesting, because raw ownership without time commitment is unstable. A standard four-year schedule with a one-year cliff can be written as

$$v(m) = \begin{cases} 0, & 0 \le m < 12, \\ 0.25 + \dfrac{0.75}{36}(m - 12), & 12 \le m \le 48, \\ 1, & m \ge 48. \end{cases} \tag{8.28}$$

Nothing vests during the first year. At month 12, one-quarter vests. The remaining three-quarters vest monthly across the next 36 months. Durzynski's rationale is direct: much of the most valuable work is still in the future, so ownership has to remain attached to future contribution.

The cap-table examples are then read as a sequence of dilutions. Founders begin with a simple split. Key employees and advisors are added. An angel round arrives. An option pool is created, diluting existing holders even before the later investor shares are issued. Then larger venture rounds arrive and dilution accelerates. The lecture is concrete here: after enough rounds, a founding CEO may be down in the neighborhood of 10% to 15%, and if dilution becomes too severe the incentive structure can break. Boards sometimes respond by re-upping a founder's compensation package precisely because motivation is part of the value of the company.

The lecture also distinguishes advisory roles from board roles. Advisors may receive small grants – often around a quarter-percent or a half-percent, subject to negotiation and market conditions – in exchange for support and credibility. But a true board seat carries fiduciary duty and a continuing legal obligation to the corporation. That is why one often works with someone as an advisor before asking them to join the real board.

Durzynski then makes the control distinction explicit. There is equity control and there is board control. These are not the same thing. Preferred stock may carry board access and governance rights that do not appear in a naive percentage table. Hence the lecture's practical closing principle: do not think only about raw percentage. Think about who is at the table, who appoints the independent director, how decisions are made, and whether trust with investors is strong enough to support the company through later rounds. In that spirit he repeats a memorable line: in negotiations of this kind, one wants to be at the table rather than on the menu.

The final practical notes follow naturally. Outside VC financing often brings directors' and officers' insurance. Investor board seats do not normally receive extra equity beyond the investment itself. And because the market for talent changes over time, compensation for senior hires, advisors, and board members remains negotiated rather than fixed by formula. The arithmetic is clean; the actual bargain is not.

8.7.1 Question & Answer

Question. How much of the company can we give away before the founder stops having a meaningful reason to stay?

Answer. There is no exact universal threshold in the lecture, and the speaker refuses to pretend otherwise. But the warning is sharp. After enough rounds, a founder can discover that the ownership left on the table no longer feels like a real stake. When that happens, incentive and control begin to separate dangerously. The lecture's response is twofold: first, vesting and option-pool design should reward future contribution; second, boards sometimes have to repair a founder's package if dilution has gone too far. Equity, in this lecture, is never just arithmetic. It is arithmetic tied to motivation, trust, and continued work.

8.8 Summary

The lecture unfolds with more deliberateness than a bare list of topics would suggest. We begin from founder skepticism and answer it with a biological fact: cash is oxygen. We then pass into investor arithmetic, because financing changes the scale of success the business must eventually reach. Next we learn the grammar of the standard financial statements and the ratio language by which investors compare business models. From there we build the model from the bottom up: units, prices, channels, hires, support, overhead, KPIs, and sales timing. Only then do we step back to the four-year P&L, read its graph as a compressed argument about delayed profitability, and move into the quarterly cash-flow recurrence, where timing becomes the central problem of survival. Finally, and only finally, the lecture turns to equity, vesting, dilution, option pools, board control, and founder incentive. The model is wrong in detail because the future is unknown, but it is indispensable because without it we do not know what we are asking the company to become.

Chapter 9

Why Do You Want to Be an Entrepreneur?

This chapter follows Session 6, Part 1 of MIT OpenCourseWare's Nuts and Bolts of New Ventures, curated for these notes by LazyingArt LLC through Video2Book. Joseph Hadzima frames the session as a return to a personal question that had been present from the beginning: not merely how to start a venture, but whether this way of life is right for us. Bob Jones then turns that question into a disciplined inquiry. The mathematics is sparse, but it matters: a happiness heuristic, a failure-rate estimate, a hypothesis-testing loop, and a few back-of-the-envelope checks keep the chapter from becoming either startup romance or startup despair.

9.1 From Venture Mechanics to Founder Fit

The lecture begins by closing the course loop. Hadzima reminds the room that the earlier sessions have dealt mostly with factual and operational questions: whether a venture needs a corporation, how a founder finds a customer, how projections work, how negotiation and financing enter the picture. But one question was different. Is this right for you? What does entrepreneurship mean for the person who chooses it?

Jones does not answer immediately. He first asks the students what they expected from the course, what they got, and whether it had value. That exchange is not throat-clearing. It models the same customer-discovery discipline the course has been teaching. A course about customers asks its own customers what worked and what did not.

The answers also recap the course's practical orientation. Students mention distillation, negotiation, financial projections, pitches, financing sources, and the usefulness of seeing the whole process without assuming prior business training. Jones listens, draws out the distinction between high-level principle and hard execution, and then pivots: tonight's agenda is stories, the hard parts of entrepreneurship that rarely get discussed, some lessons, and a wrap-up.

So this chapter is not an inspirational appendix. It is the founder-fit chapter. The earlier lectures ask whether the venture can be built. This one asks whether we can live with what building it demands.

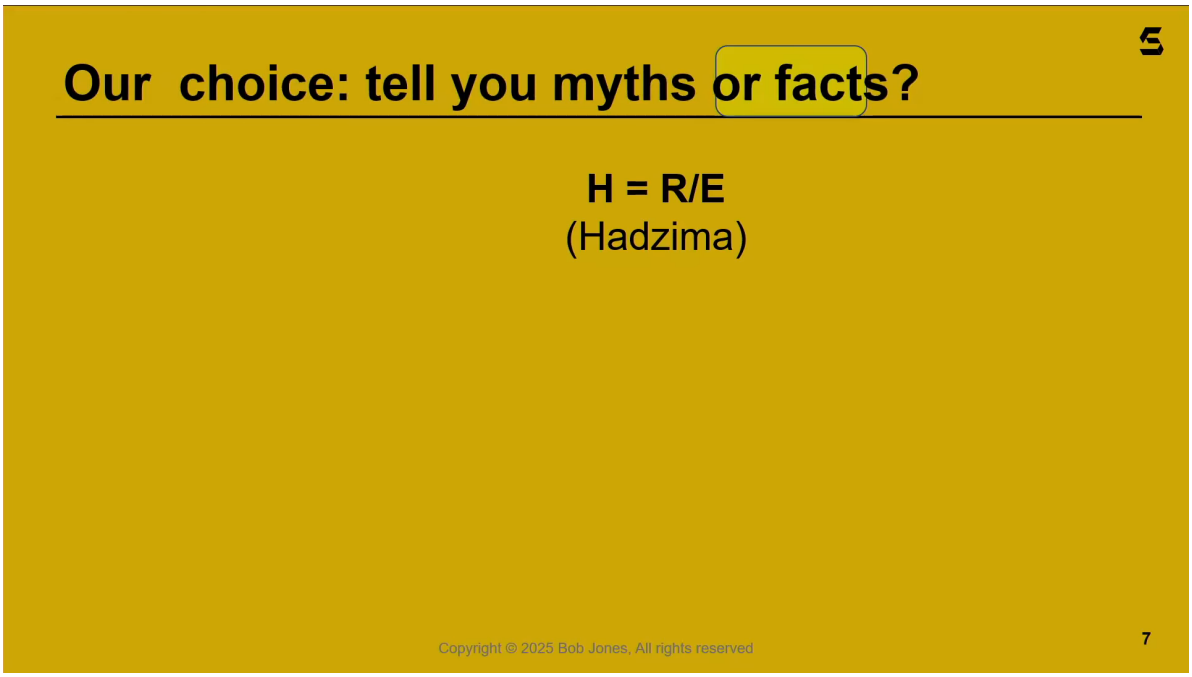

Figure 9.1: Hadzima's expectation heuristic as shown in the lecture.

9.2 Myths, Facts, and the Expectation Reset

Jones begins the substance of the lecture by naming the myths. There is a popular story in which money is widely available, investors shower a brilliant idea with capital, and a boot camp can supply a magic formula for success. Some organizations half-promise that story. Jones says the course faced a moral choice: promote the myths, or tell the facts.

The first mathematical object appears exactly there.

The slide displays the compact formula

$$H = \frac{R}{E}. \tag{9.1}$$

The transcript supplies the meanings of the symbols:

$$H = \frac{\text{reality}}{\text{expectations}}, \qquad H = \text{happiness}. \tag{9.2}$$

The slide itself shows the symbols and credits Hadzima; the spoken lecture gives the interpretation. We should read this as a heuristic, not as a psychological law. It is a way to say that disappointment is often caused not only by harsh reality, but by inflated expectations.

A small calculation captures the point:

$$H = \frac{R}{E}, \tag{9.3}$$

$$\text{with } R \text{ fixed, increasing } E \text{ decreases } H, \tag{9.4}$$

$$\text{with } R \text{ fixed, calibrating } E \text{ stabilizes } H. \tag{9.5}$$

That is why the lecture starts with the unpleasant facts. Jones is not trying to make entrepreneurship dreary. He is trying to remove false denominators.

Figure 9.2: The lecture's visible failure-rate framing. The statistic is used as a lecture claim, not as a sourced universal law.

The facts are blunt. Entrepreneurship is hard. Most startups fail. Most investors assume the founder will probably fail, even if the founder is talented and charming. Many entrepreneurs are competing for talent, capital, and other resources, and many will not get the money.

Later Jones returns to the same point with a slide that gives the lecture's starkest quantitative claim.

The slide's clean note form is

$$\text{survival after two years} \approx 10\%. \tag{9.6}$$

Jones also states the same general claim verbally:

$$\text{startup failure rate} \approx 80\% \text{ to } 90\%. \tag{9.7}$$

He immediately qualifies the measurement. In the food industry, one metric is whether the product is still on the shelf two years later; other industries use other metrics. So the statistic should be treated as expectation-setting, not as a theorem about all ventures.

Remark 9.1. The point of the numbers is not precision. The point is calibration. If the expected path is easy money, then ordinary entrepreneurial reality looks like catastrophe. If the expected path includes difficulty, rejection, and revision, then the same facts can be used for learning.

Jones then asks why he is talking this way. The answer is disclosure. People who enter demanding schools know they will take a beating somewhere along the way. Entrepreneurship is often promoted without the same candor. Jones wants the room to know that even smart founders fail, and that he does not know a single entrepreneur who got everything right the first time.

9.3 The First Flop

The lecture now moves from equation to story. Jones tells us about springboard diving with his fourteen-year-old daughter. The details matter because the analogy depends on the felt sequence: curiosity, fear, instruction, flop, pain, and decision.

A new diver first learns to work the board, gain height, and perform basic moves. Then comes the first real test. Jones's daughter is sent to the high board, told to turn backward with the water behind her, tuck, fall, wait for the command, come out of the tuck, find the water, and enter head first. The second time, this may seem manageable. The first time, it feels like death. She panics, fails to come out properly, and lands in a full back flop.

Jones then slows down the crossroads. She is hurt, embarrassed, and comforted. She can stop. Or she can climb back up. After a few minutes she looks at the boards, walks to the stairs, and says she is going to do it again and get it right.

The entrepreneurial analogy is direct:

$$\text{first serious attempt} \longrightarrow \text{painful flop} \longrightarrow \begin{cases} \text{stop,} \\ \text{try again.} \end{cases} \tag{9.8}$$

Jones does not say the second branch is always superior. For many people, stopping is the right decision. Entrepreneurship is not for everyone. But some people discover that the first flop did not kill them, and that discovery changes the next decision.

9.3.1 Question & Answer

Question: After the first hard failure, why would anyone continue?

Answer: Not because failure is pleasant, and not because persistence is automatically wise. The lecture's answer is narrower. Some people reinterpret the first flop as survivable. If it is survivable, it can become information. The diver goes back up the ladder not because the pain was imaginary, but because the pain did not settle the question. Likewise, the founder who continues does not deny the failure. The founder refuses to treat one failure as the final measurement of the person.

That is the first conceptual obstacle in the lecture. If failure is identity, the story ends. If failure is experience, the next question is what the experience teaches.

9.4 Failure, Rejection, and the Social Cost of Starting Up

Jones next gives a sequence of entrepreneurial cases. Trevor starts at twenty-six with newspaper attention, a car, status, and confidence. A year later he has lost the visible symbols of success: the car is repossessed, the relationship ends, and he is living in his parents' basement while working at Radio Shack. Jones does not sanitize the humiliation. Trevor says it took years to get the taste of failure out of his mouth.

The lesson is not that humiliation is noble. The lesson is that failure was not fatal. Trevor later succeeds, but only after becoming a perpetual learner. Jones states the trait conclusion sharply:

$$\text{entrepreneurial survival} \not\equiv \text{intelligence}, \qquad \text{entrepreneurial survival} \not\equiv \text{creativity}. \tag{9.9}$$

The trait he emphasizes is adaptability:

$$\text{key trait} \approx \text{adaptability}. \tag{9.10}$$

Laura's story adds the customer-discovery mechanism. She leaves a senior corporate role, assumes she is the smartest person in the room, and is shaken almost immediately by hard questions. One idea is wrong. Another consumes substantial personal money and fails. Then she realizes that the prospects are not going to use the product.

The compressed rule is one the course has been repeating:

$$\text{validate} \longrightarrow \text{talk to customers} \longrightarrow \text{learn whether they want it}. \tag{9.11}$$

If customers do not want it, enthusiasm cannot force a sale.

Laura's later success comes through reaching out. She has coffee with a real-estate executive and does not initially understand the scale of his network. When he says he wants the product for all his agents, the scale becomes concrete:

$$\text{one useful conversation} \longrightarrow \text{11,000 possible users}. \tag{9.12}$$

This is not a universal distribution formula. It is a local demonstration of why private reasoning is too small. Markets are discovered through people.

Jones then turns to loneliness. The Jack-and-Jill parable is partly garbled in the transcript, but its mechanism is clear. Startup uncertainty, normal inside the venture, may sound terrifying at home. When Jack says he does not know what he is doing, Jill hears mortgage, children, savings, and risk. Jack learns that bringing self-doubt home may not be safe. But leaving it at work isolates him.

The rejection is not occasional. Customers say no. Partners say no. Investors say no. Potential employees take safer jobs. A founder who has been rewarded for intelligence and effort may enter a world where effort does not prevent repeated refusal.

Jones adds a harsher warning from another entrepreneur's story: not everyone has the founder's best interests at heart. The sidebar is direct:

- Do not raise money from the wrong people.
- Protect intellectual property, especially a secret algorithm, formula, or defensible business method.

The emotional companions of this phase are anxiety, self-doubt, and imposter syndrome. A founder may present as the capable captain of the ship while privately wondering whether that is true. This is where the lecture needs a method, not another anecdote.

9.5 Failure as Evidence, Not Identity

Jones asks whether there is hope for us. His first answer is community. Ordinary friends may not understand why the founder is not simply working in a stable job. They may see the founder's choices as irrational. So Jones recommends finding a community of similarly wired people: people for whom uncertainty, ambition, repeated attempts, and bruising failures are intelligible.

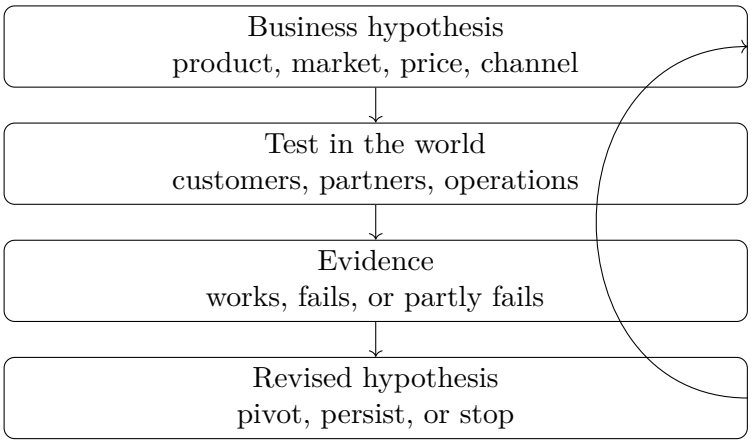

Figure 9.3: A transcript-backed reconstruction of Jones's hypothesis-testing loop. It is not copied from a slide; it renders the spoken method compactly.

Then he shifts gears. The word "pivot" can sound like a polite synonym for not knowing what we are doing. Jones accepts the embarrassment and then changes the frame. Scientists articulate hypotheses and test them. If a hypothesis proves wrong, they do not say that the scientist is worthless. They say it is time for the next hypothesis.

The lecture's cleanest formal structure is therefore

$$\text{hypothesis} \longrightarrow \text{test} \longrightarrow \text{failure/evidence} \longrightarrow \text{revised hypothesis}. \tag{9.13}$$

We can write the update rule as a cautious notation. Let B_n be the current business hypothesis at attempt n, and let D_n be the evidence produced by testing it. Then

$$B_{n+1} = \text{Revise}(B_n, D_n). \tag{9.14}$$

This is not an algorithm that computes the next company. It is a discipline of interpretation. The practical identity to reject is

$$\text{failed hypothesis} = \text{failed person}. \tag{9.15}$$

The lecture's replacement is

$$\text{failed hypothesis} \neq \text{failed person}. \tag{9.16}$$

9.5.1 Question & Answer

Question: If failure is common, how should we think about it?

Answer: We should not make it smaller than it is. Failure can cost money, time, status, confidence, and relationships. But in Jones's scientific frame, failure is also evidence. It tells us that one product, market, pricing assumption, channel, team configuration, or timing choice did not work. The useful response is not denial. The useful response is revision.

This is why the jockey analogy works. We may be a capable jockey on the wrong horse. The answer is not to insist that the horse is fine. The answer is to find a better horse.

9.6 Supports: Community, Advisors, and Self-Management

The hypothesis loop is not meant to be run alone. Jones recommends a business advisory board, and he distinguishes it sharply from a board of directors.

A board of directors represents shareholders. It has fiduciary duties. It may be empowered, and in some cases obligated, to remove the founder. That is often not the room in which to say, "I have no idea what to do next."

An advisory board is different. It has no fiduciary control over the company. It is a place to assemble people who know things we do not know, want the venture to succeed, and can help in narrow but crucial domains. Jones's beverage example makes the point concrete. He and a scientist friend had a 2.5-ounce drink and needed a manufacturer who could bottle it below an unaffordable scale:

$$\text{needed production scale} < 2{,}000{,}000 \text{ bottles}. \tag{9.17}$$

Two days later, an advisor had solved the problem because the advisor knew that narrow arena.

The equity anecdote is a small cap-table lesson. Jones thought about giving each advisor one point:

$$\text{proposed advisor equity} = 1\%. \tag{9.18}$$

The advisors said that was too much and recommended half a point:

$$\text{revised advisor equity} = 0.5\%. \tag{9.19}$$

This is not a general rule that all advisors should receive 0.5%. The point is that good advisors can improve not only operations but also the decisions that future investors will read as evidence of sophistication.

Jones also describes the right mix. Some advisors should be encouraging enough that we can say we fell down and hear, "You will get up." But one or two should be able to give ten sensible reasons why the current plan may not work. If seven are already covered and three are new, those three are valuable.

The support structure then turns inward. Many founders believe the business must consume every hour. Jones argues that this destroys judgment, and judgment is one of the founder's scarce assets. The fixed constraint is

$$W = 168 \text{ hours/week}. \tag{9.20}$$

No founder gets more. Spending all of W on the company is not a proof of seriousness. It is a path to self-destruction.

Jones gives a sleep comparison from literature he encountered while working in the sleep business:

$$5 \text{ hours/night of sleep} \sim \text{cognitive impairment of three drinks}. \tag{9.21}$$

We should keep the comparison in exactly that spirit: a lecture claim used to make the decision-quality point vivid. Frayed-out entrepreneurs make bad decisions. They may not die, as in the motorcycle analogy, but the business may.

The recommendations are plain: eat decently, exercise, walk, sleep, take care of loved ones, and do things that recharge the battery. Jones's Halloween story supplies the weight. His children remembered that he missed Halloween long after he remembered the business reason for missing it. Business issues are usually still there in the morning.

9.7 Why Do It Anyway?

After making the path difficult, Jones returns to the question that gives the lecture its title. Why would anybody do this?

He answers with several motives, and he does not reduce them to one. Some people feel they have to create something better. Some want to be their own boss badly enough to take a pay cut. Some discover that corporate life simply does not fit. Some want to improve the quality of life for others. Some want to create wealth for themselves and their families. Some love the buzz of technology or the act of building.

The service motive appears in two stories. One entrepreneur in health-care software seems, at first, to be building toward investor appeal and acquisition. But when he talks about an orphanage in Venezuela, his energy changes. Jones helps him see that the real objective may not be to grow as fast as possible toward sale. It may be to build a durable company that gives him room to support the work he cares about.

Another entrepreneur brings internet access to rural Nebraska. The business does not sound glamorous until he describes a mother who has been taking her children to a library parking lot so they can use Wi-Fi for homework. When internet access reaches her home, she can cook dinner with her children for the first time in years and return to normal work hours. That is not mythology. It is a concrete improvement in daily life.

A compact way to record the motive set is

$$\begin{aligned}\text{entrepreneurial motive} \in \{&\text{autonomy}, \text{misfit with corporate life}, \text{service}, \text{wealth creation}, && (9.22)\\ &\text{love of building}, \text{technology excitement}, \text{creative discipline}\}. && (9.23)\end{aligned}$$

Jones's music analogy gives the last term its shape. A musician needs scales, modes, time, and framework, but also creative life inside that framework. Entrepreneurship is similar:

$$\text{good entrepreneurship} \approx \text{creativity} + \text{discipline}. \tag{9.24}$$

9.7.1 Practical Timing and Runway

The final audience questions return the lecture to judgment. If it may take several attempts to build something valuable, how does a founder create enough runway?

Jones gives two answers. First, failed entrepreneurship does not necessarily make someone unemployable. Some companies want change agents: people who are not afraid of a blank sheet of paper and can bring creative thinking into established organizations. Second, a founder can weave in and out. Work a normal job for a while, learn, earn, calm the family situation, and later start again.

The next question is whether to keep a day job while testing the business, or to go all in. Jones argues against going all in before doing at least some market arithmetic. He refers back to the first session's back-of-the-envelope analysis:

$$I_{\text{needed}} = \text{income the venture must produce}, \tag{9.25}$$

$$V_{\text{customer}} = \text{economic value of one customer}, \tag{9.26}$$

$$N_{\text{required}} \approx \frac{I_{\text{needed}}}{V_{\text{customer}}}. \tag{9.27}$$

The next check is whether the market can plausibly supply that many customers:

$$N_{\text{reachable}} \geq N_{\text{required}}. \tag{9.28}$$

If we have not done this sort of check, Jones says it is not the right time. If we have, it might be.

The small-step version is

$$\text{side test} \longrightarrow \text{mistakes while still paid} \longrightarrow \text{validation} \longrightarrow \text{stronger full-time pitch.} \tag{9.29}$$

This is where passion must be governed. The right level of passion helps us work on Saturday and get back up after being knocked down. Too much passion blinds us to the obvious. Jones's rule is to deploy passion wisely.

9.8 Summary

This lecture is the personal counterpart to the operational sessions of Nuts and Bolts. The earlier material teaches incorporation, customers, projections, negotiation, and financing. This session asks whether we can live with the uncertainty, rejection, fatigue, and repeated revision that those mechanics imply.

The first tool is Hadzima's heuristic:

$$H = \frac{R}{E}.$$

The chapter is largely about correcting E, expectations, before they distort judgment. The visible failure-rate slide and Jones's spoken estimate that eight or nine out of ten startups fail are not universal laws, but they are strong expectation-setting claims.

The most useful method is the hypothesis loop:

$$\text{hypothesis} \rightarrow \text{test} \rightarrow \text{evidence} \rightarrow \text{revised hypothesis}.$$

Failure remains painful, but it need not become identity. To survive that loop, founders need community, advisors, health, sleep, family, and a motive that survives contact with reality. Entrepreneurship has rewards for certain people, but it is not for everyone. That is not an afterthought. It is one of the course's central facts.

Chapter 10

Common Startup Mistakes

This chapter follows the final practical movement of MIT OpenCourseWare's *Nuts and Bolts of New Ventures*, curated by LazyingArt LLC through Video2Book. Marina Hatzopoulos is not giving a generic startup checklist. She is reconstructing, from the Z Corporation case, how a protected technical possibility becomes a product, a company, and a market. Bob frames the session as part of Joseph Hadzima's course: the speakers were not scripted, but the same themes kept recurring. We should therefore read the repetition as evidence from practice.

10.1 A Deep-Tech Case Before the Rules

Bob's handoff is important. Marina has noticed that some of her warnings overlap with earlier speakers. Bob tells the room that the course speakers did not coordinate their slides. If the same patterns recur, the repetition is not a slogan; it is convergence. That is the right frame for the chapter. We begin with a case, not a doctrine.

Marina begins with biography because it explains the mixture of skills the lecture will require. She studied math and music, then worked in banking, where she learned accounting, selling, and negotiation. At Thermo Electron, now Thermo Fisher, she discovered that evaluating a technology company was not just reading financial statements. It required judgment about the technology itself. That missing piece took her to MIT for a master's degree in mechanical engineering, not to become an engineer, but to manage engineers and speak their language.

The turning point was the Technology Licensing Office. MIT had patented inventions made by faculty and students, and the office looked for people and companies to commercialize them. There Marina encountered an early 3D-printing technology, before "3D printing" had become a common term. Professor Ely Sachs's group had been working on an ambitious printing project. Two engineers then turned the problem around: instead of designing a printhead around a chosen material system, they designed a material system around an off-the-shelf inkjet printhead. The first material story was deliberately crude: water, or ink, and cornstarch.

That reversal is already the deep-tech pattern. The company did not start with a finished product. It started with protected technical knowledge, a proof that the core science could work, and a large uncertainty about what market would care.

Z Corporation was founded in 1994 in Kendall Square. The company delivered its first product

within about two years, reached profitability two years later, and eventually grew to

$$\text{revenue} \approx \$30\text{M}, \qquad \text{employees} \approx 125. \tag{10.1}$$

It did this with no venture capital. Marina is careful about that point: the absence of venture capital was not proof of superior wisdom. Nobody wanted to invest in a hardware business while the Internet market was taking off. The constraint nevertheless became part of the operating discipline.

The first time scale is therefore sober:

$$\text{deep-tech venture horizon} \sim 10 \text{ years}. \tag{10.2}$$

For these notes, that is a boundary condition. Patents, proof of principle, market search, pricing, team formation, fundraising, and founder psychology all unfold inside a long hardware clock.

10.2 Intellectual Property and Proof of Principle

Marina next asks why someone creates a startup. The motive may be love of a technology, desire to solve a problem, independence, teamwork, community, or money. Her point is not moral ranking. The founder has to know the motive because the definition of success depends on it. A founder trying to improve brain-tumor diagnosis may make different choices from a founder whose primary aim is financial return.

Then the lecture narrows: when we are doing a deep-tech startup, the starting point should be intellectual property. A patent is presented as a bargain. We disclose how to do something, and in exchange receive a time-limited monopoly:

$$\text{patent protection} \approx 20 \text{ years}. \tag{10.3}$$

This is not decorative law. In Z Corporation's case, Marina says a large customer took apart their machine, but did not launch a competing product until after the MIT patents expired. The patent period bought time to invest, educate the market, and grow.

The second rule concerns disclosure order:

$$\text{patent first} \quad \longrightarrow \quad \text{publish second}. \tag{10.4}$$

The lecture states this as an operating rule for researchers and students. Publishing gives credibility, but premature publication can destroy patent protection. We do not need to expand this into a full legal doctrine here. The venture rule is enough: do not publish away the boundary that makes the technical advantage defensible.

The proof of principle is the next object. It is not the product. It is not even a usable prototype. It is the ugly machine that removes the central scientific doubt:

$$\text{proof of principle} \neq \text{finished product}, \qquad \text{proof of principle} \Rightarrow \text{core technical risk reduced}. \tag{10.5}$$

The early Z Corporation machine required manual powder handling and multiple computers. Its value was not elegance. It proved that a paper-printer-derived process could build a three-dimensional part.

Remark 10.1. For the rest of the lecture, intellectual property functions as a North Star. When the market is unclear, when customers disagree, and when the team is pulled in several directions, Marina returns to the same question: what makes us special that we can protect?

Figure 10.1: Market disruption sources. The slide supplies the categories; the photograph is illustrative.

10.3 From Stealth to Market Disruption

The first operational mistake is part-time entrepreneurship. Marina accepts Bob's earlier qualification: one should validate before quitting a job. But once the venture is validated and moving forward, time matters. A startup worked on for a few hours on Saturday morning will not move fast enough. Researchers need entrepreneurs, and entrepreneurs need researchers.

The next instinct is stealth. Founders fear that if they share the idea, someone will steal it. Marina's answer is practical: if all we have is an idea, then anyone can copy it once it reaches the market. The solution is not silence about the problem. The solution is patents, trade secrets, protected know-how, and disciplined judgment about what not to reveal. Customers do not need the inner workings. They need to know whether the product solves their problem.

The first market questions are therefore:

$$\text{painful problem?} \qquad \text{large enough market?} \qquad \text{worth entering?} \tag{10.6}$$

The slide gives a compact set:

$$D \in \{\text{regulations},\ \text{demographic changes},\ \text{new enabling technologies}\}. \tag{10.7}$$

The transcript adds the important caveat. Some disruptions are too transient. Marina mentions corporate DEI spending and COVID-related startups as examples of openings that may not last long enough to support a durable company. Thus the test is not merely whether the world changed. It is whether the change persists long enough for company formation:

$$\text{startup opportunity} \approx \text{durable disruption} + \text{startup ability to move}. \tag{10.8}$$

10.3.1 Question & Answer

Why does disruption help a startup when large companies have almost every advantage?

Large companies have manufacturing, R&D, sales, marketing, distribution, money, experience, brands, market knowledge, and customer relationships. Marina says they have nearly everything except the ability to pivot quickly. The startup edge is not that incumbents are stupid. The edge is maneuverability:

$$\text{startup edge} = \text{ability to pivot under incomplete information.} \tag{10.9}$$

This matters because innovative markets do not present themselves as clean problem sets. We are not given all the variables. We are making decisions in fog. The startup must try, learn, and redirect.

10.4 Market and Technology Development Interact

The next movement is the lecture's main formal model. A deep-tech company develops two paths at once. One path is technical: run tests, try technology directions, isolate the highest-risk component, and de-risk it without building the whole system. The other path is market-facing: articulate what is special because of the IP, consider a wide range of industries, speak with experts, and test possible markets quickly.

The transcript gives a rough target for the early search:

$$N_{\text{market ideas}} \approx 20. \tag{10.10}$$

This is not a statistical sample size. It is an antidote to premature narrowing. Marina calls the search speed dating, not dinner. We are looking for fast evidence, not perfect certainty.

The visible slide notation is:

$$\text{iterative process} \Rightarrow \text{information from one drives the other.} \tag{10.11}$$

For the notes, the cleaner reconstruction is:

$$\text{market learning} \longleftrightarrow \text{technology development.} \tag{10.12}$$

The important constraint is that product-market fit cannot drift away from protectable differentiation. A market may be attractive, but if the company cannot protect the advantage there, the fit is strategically weak.

10.4.1 Question & Answer

How do we validate a deep-tech idea before there is a finished product?

Marina's answer is staged. First build the proof of principle to reduce the core technical risk. Then build a crude prototype, storyboard, or sample part. The artifact does not have to be beautiful. It has to be concrete enough that customers can react to something closer than an abstraction.

The update rule is:

$$\text{crude artifact} \longrightarrow \text{customer feedback} \longrightarrow \text{technical revision} \longrightarrow \text{better artifact.} \tag{10.13}$$

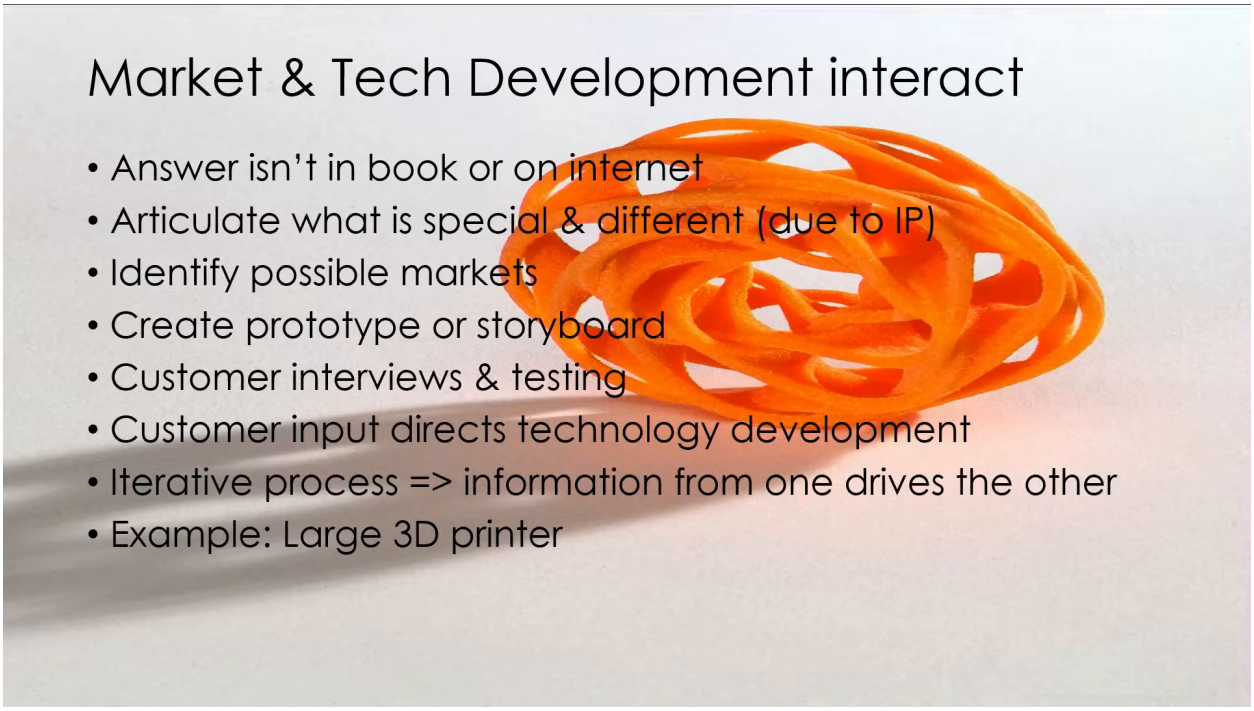

Figure 10.2: Market and technology iteration. The slide's arrow is informal; it marks a practical feedback loop.

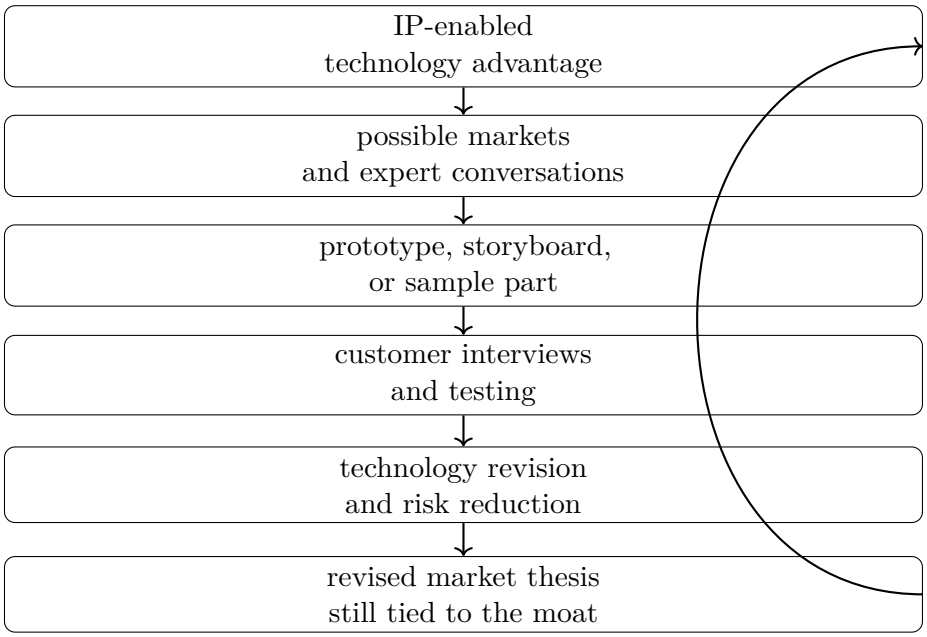

Figure 10.3: A pocket-safe reconstruction of the deep-tech feedback loop. The loop is narrow by design: the market and the technology revise one another, but the loop keeps returning to the moat.

In the Z Corporation story, Marina took sample parts to a trade show without even having a booth. Much of the feedback was negative: the parts were crude. But the conversation exposed the dimension that mattered. The parts were not perfect, but they could be made quickly.

The evidence hierarchy is:

$$\begin{aligned} \text{compliment} &< \text{reaction to idea} && (10.14)\\ &< \text{reaction to prototype} && (10.15)\\ &< \text{near-real product use} && (10.16)\\ &< \text{customer money.} && (10.17) \end{aligned}$$

The closer the customer is to using the product as intended, the more useful the feedback becomes. Payment, deposits, and R&D contracts are the strongest signals in the lecture.

10.4.2 Question & Answer

What is deep tech in this lecture?

Marina defines it pragmatically. Deep tech is technology where the core science exists, often from university research, but the product will take a long time and substantial development to build:

$$\text{deep tech} \approx \text{core science exists} + \text{long development path.} \tag{10.18}$$

Often it is a solution looking for a problem, not a problem already matched to a product. That is why the two paths must run in parallel. We do not know exactly where to take the technology until we learn who cares.

She is also careful about scope. When asked about software, AI, open-source models, and low-cost prototyping, she says that her experience is patentable hardware and that she is not the right person to answer the software/open-source version. These notes should preserve that boundary rather than inventing an IP framework she did not give.

10.5 Differentiation, Customers, and the Pricing Pivot

The lecture now shifts into its recurring rhythm: mistake, then instead. One mistake is to seek a popular market because everyone is talking about it. The founder imagines a crowded market must be promising, assumes incumbents are slow and stupid, and thinks the company can always lower price.

Marina rejects that logic. Lowering price is not the answer. In deep tech, the state of the art is the starting point, not the endpoint. The aim is to be better on a dimension the customer cares about and that the company can defend:

$$\text{protectable differentiation} \Rightarrow \text{possible premium price.} \tag{10.19}$$

Competition is a double-edged sword. For Z Corporation, other 3D-printing firms helped educate the market. Once customers understood why they might need a 3D printer, Z Corporation could sell against alternatives by explaining that its machine was faster, cheaper, and capable of color. The worst competitor was often not another machine. It was doing nothing.

The central quantitative story begins with the initial market. High-end rapid-prototyping machines sold for roughly

$$\$100{,}000 \text{ to } > \$1{,}000{,}000. \tag{10.20}$$

Z Corporation planned to enter at the low end with an office-compatible machine:

$$P_0 = \$20{,}000. \tag{10.21}$$

A year into development, before launch, two incumbents announced office-compatible machines at

$$P_{\text{incumbent}} = \$60{,}000. \tag{10.22}$$

At the trade show, Z Corporation learned that its technology was approximately

$$\text{Z Corp speed advantage} \approx 10\times \tag{10.23}$$

faster than the incumbent technologies. The surrounding transcript is partially garbled, but the repeated clear claim is speed. The parts had shortcomings, but customers cared that they could hold something quickly instead of waiting.

Walter Bornhorst, the chairman, proposed remaking the price decision:

$$\$20{,}000 \longrightarrow \$60{,}000. \tag{10.24}$$

10.5.1 A Worked Decision Update

We should write the pricing story as an update, not as a universal rule. Let P be the planned price and let E be the new evidence collected before launch:

$$E = \{\text{incumbents at } \$60{,}000,\ \text{speed advantage} \approx 10\times,\ \text{customers care about speed}\}. \tag{10.25}$$

The old plan was:

$$P_{\text{old}} = \$20{,}000. \tag{10.26}$$

The revised decision was:

$$P_{\text{new}} = \text{revise}(P_{\text{old}} \mid E) = \$60{,}000. \tag{10.27}$$

The mathematical point is the conditional update:

$$\text{new strategic evidence} \Rightarrow \text{remake the strategic decision}. \tag{10.28}$$

This does not mean "triple the price." It means that price should track the value evidence, the competitive reference point, and the company's survival economics.

Marina adds a practical asymmetry:

$$\text{lowering price after launch is easier than raising price after launch}. \tag{10.29}$$

The higher price mattered because production, sales, and market education were more expensive than expected. Educating the market took about:

$$\text{market-education cycle} \approx 1 \text{ year}. \tag{10.30}$$

10.5.2 Question & Answer

When should we listen to customers, and when should we ignore them?

Marina does not give a magic rule. Entrepreneurship constantly requires both listening and not listening. Customers may dismiss a new thing because they cannot yet imagine it. Industry veterans may reject anything unfamiliar. But feedback becomes stronger as it gets closer to real use:

$$\text{feedback value} \uparrow \quad \text{as} \quad \text{use context approaches intended product use.} \tag{10.31}$$

There is also danger on the opposite side. Power users can pull a company into a narrow, expensive direction. Z Corporation built a larger machine because many early customers asked for it. Later, Marina says the market was limited and the technical challenges were larger than expected. The project pulled against the long-term path of making machines less expensive and easier to use. The test is:

$$\text{customer request useful?} \iff \text{request lies on the long-term strategic path.} \tag{10.32}$$

This is why the North Star returns. If architecture customers cared about beautiful surface finish, but Z Corporation's protected advantage was speed, then architecture was not the right early market. The better question was: who cares about speed more than surface finish?

10.6 Teams, Conflict, Priorities, and Cash

The next mistakes are internal. Hiring people one trusts sounds reasonable, but loyal friends and family may not supply the missing skills. A famous board chair who is too busy may be less useful than a less famous person who is engaged. Z Corporation's founding team mattered because it combined business strategy, business execution, material science, software, and hardware.

The team rule is:

$$\text{founding strength} \approx \text{complementary skills} + \text{mutual respect} + \text{ability to make hard changes.} \tag{10.33}$$

Marina adds a sharp practical warning: do not hire someone you cannot fire. If possible, work with someone before hiring. A one-week paid trial can reveal more than an interview.

The cultural mistake is to be merely nice. Marina distinguishes courtesy from avoidance. Z Corporation argued about everything, but without taking it personally and without grudges. That structure matters:

$$\text{diverse views} + \text{nonpersonal conflict} \longrightarrow \text{better decisions.} \tag{10.34}$$

Negotiation is treated the same way: not a personality trait, but a skill to practice.

The next trap is confusing work volume with strategic work. A founder can get enormous satisfaction from checking off easy items. But the company needs the first item, not the seventeenth. Marina separates decisions:

$$\text{strategic} = \{\text{business},\ \text{team},\ \text{product},\ \text{pricing}\}, \tag{10.35}$$

$$\text{non-strategic} = \{\text{offices},\ \text{furniture},\ \text{software choice},\ \text{conference rooms}\}. \tag{10.36}$$

The strategic decisions must be remade when new information arrives:

$$\text{new information} \Rightarrow \text{remake strategic decisions.} \tag{10.37}$$

No venture capital was not originally a chosen strategy for Z Corporation, but it shaped the company. Bootstrapping forced revenue, profitability, and cash discipline. Marina summarizes the case scale as:

$$\text{capital} \approx \$2.5\text{M}, \qquad \text{revenue} \approx \$30\text{M}, \qquad \text{employees} \approx 125. \tag{10.38}$$

These are company facts, not a benchmark. The lesson is that financing constraints alter operating behavior.

This connects to two later mistakes. One is always staying positive: ignore bad news, believe the schedule, believe the cost estimate, and assume this time will be different. Marina's correction is to learn from failure and assume time and cost may be at least double what people say. Another is focusing only on growth. Growth matters, but companies die when they run out of cash:

$$\text{survival} \Rightarrow \text{track cash, time to revenue, and time to profitability.} \tag{10.39}$$

A strong company may become an acquisition target. The immediate task is to build the strong company.

10.7 Pitching, Funding, and the Venture-Capital Trap

Investor communication is another place where founders put on a show. Marina warns against acronyms outsiders do not understand, complicated diagrams, tiny-font paragraphs, unrealistic projections, exit talk, confidentiality demands, and adding AI no matter what business one is in. The correction is simplification.

The visible presentation rule is:

$$\text{words per slide} \sim 10\text{–}20. \tag{10.40}$$

The spoken pitch template is:

$$\text{We make } \_\_\_\_\_ \text{ for } \_\_\_\_\_ \text{ to } \_\_\_\_\_. \tag{10.41}$$

The blanks are product, customer, and problem. If the listener cannot say what the product is, who the customer is, and what problem is being solved, technical detail will not rescue the pitch.

The funding map is broader than venture capital. Friends and family often come first. Angels can range roughly from:

$$\text{angel investment} \approx \$25{,}000 \text{ to } \$1{,}000{,}000. \tag{10.42}$$

Government grants can be substantial in deep tech. Strategic investors may be suppliers, customers, or distributors. R&D contracts are non-dilutive and validating. Customer deposits can also fund operations:

$$\text{customer deposit} = 20\%, \qquad \text{delivery} \approx 3\text{–}6 \text{ months}. \tag{10.43}$$

Venture capital and venture debt are part of the map, but they are not the universal destination.

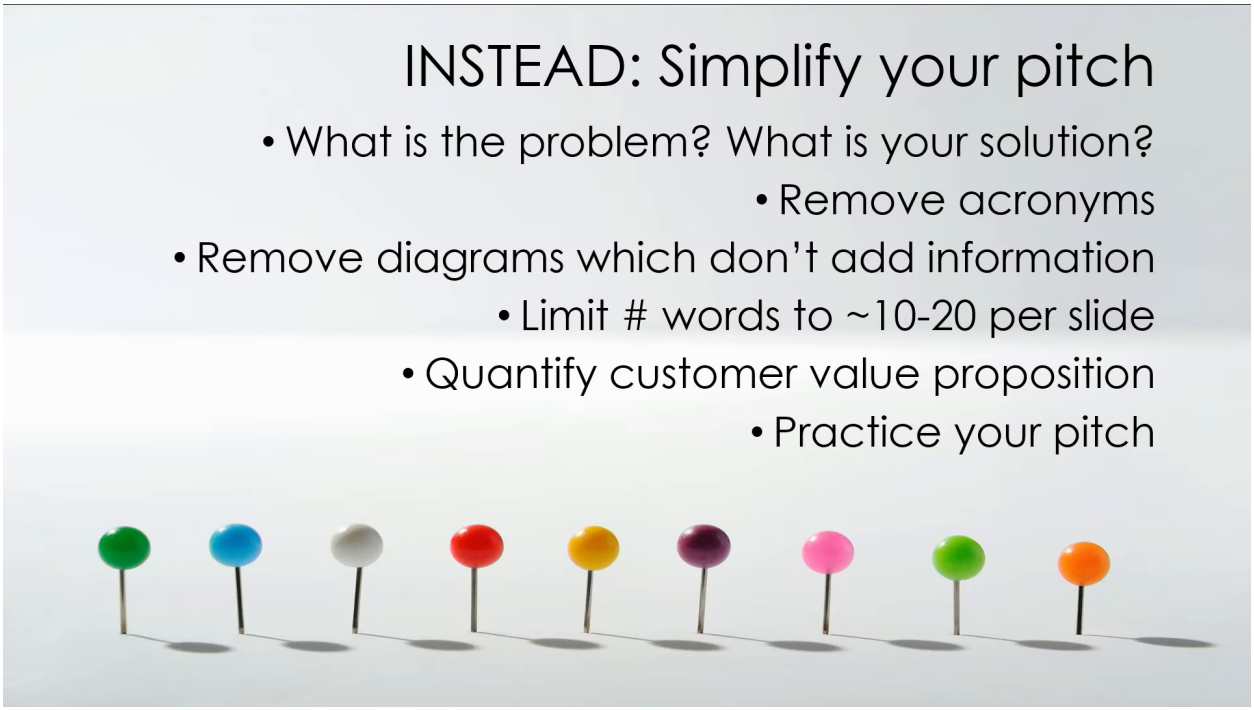

Figure 10.4: Simplifying the investor pitch. The slide contains the visible guideline of roughly 10–20 words per slide.

10.7.1 Question & Answer

Is venture capital always the right path?

No. Marina gives both sides. Venture capital can provide cash, advice, connections, and validation. Some companies need it. A two-sided global market, or a satellite company with high capital costs, may require large and fast financing.

But venture capital carries risks. Too much money can allow a company to move ahead before product-market fit is proven. Some markets cannot grow at the pace venture investors expect. Founders can lose control. High valuations raise the return bar. The more precise relation is:

$$\text{founder outcome} \neq \text{headline valuation}. \tag{10.44}$$

A very profitable \$10M company that the founder owns entirely may be a better founder outcome than pushing for a billion-dollar business and failing. The correct financing path depends on the business.

Fundraising itself is work. Research investors before pitching them. If an investor funds only AI software and the company is selling 3D printers, do not pitch and call it optimism. Marina's rule is to reject yourself before collecting needless rejection. Every failed pitch should produce a strategic takeaway:

$$\text{failed pitch} \longrightarrow \{\text{wrong investor},\ \text{unclear pitch},\ \text{flawed business model}\}. \tag{10.45}$$

Even rejection can be useful if it teaches whether the list, the communication, or the business model is wrong.

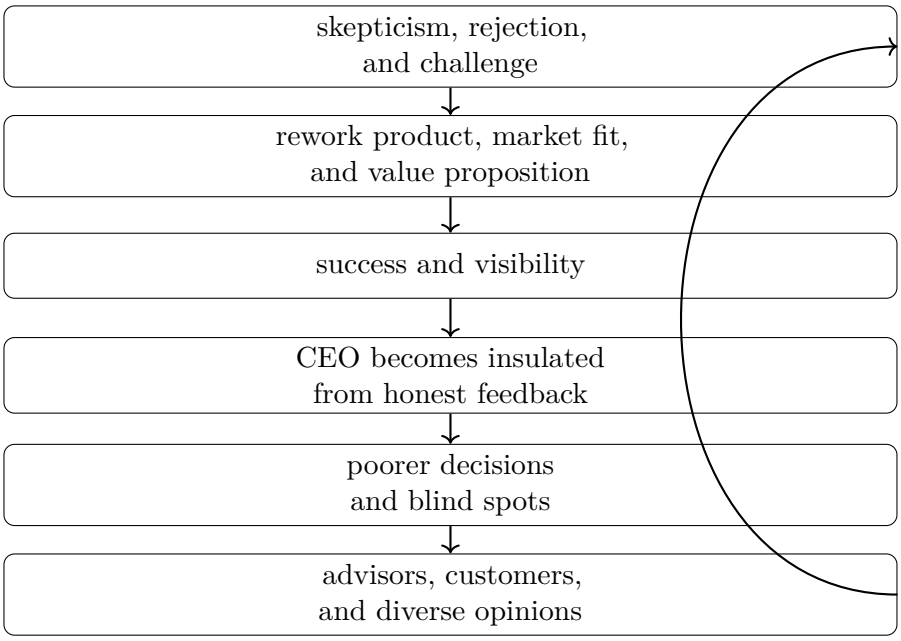

Figure 10.5: A conceptual reconstruction of CEO syndrome and its antidote. The challenge loop must be deliberately restored.

10.8 CEO Syndrome and Course Closure

The last founder mistake is psychological. Early CEOs are challenged constantly by engineers, customers, employees, and investors. That challenge forces rework: redesign the product, redefine product-market fit, and sharpen the value proposition. Then success arrives. People begin to call the CEO a visionary. Visibility and awards accumulate. New employees never saw the early uncertainty. The CEO starts to feel invincible.

The danger is a feedback loop in which success removes the challenge that produced success.

The antidote is not humility as decoration. It is structure: an advisory group willing to challenge ideas, customers close enough to the product to give real feedback, and a culture that asks questions and listens. Small rejections feel unpleasant, but they are the mechanism that improves decisions.

Marina closes her portion by restating the operating rules for deep-tech hardware: defensible product differentiation, a strong and diverse team, debate and healthy culture, customer contact, prototypes in customers' hands, revenue and profitability in the beachhead market, and decisions remade whenever new information arrives.

Bob then returns to the course as a whole. The course was not meant to be theory; it was meant to increase the probability of success. Entrepreneurship is a full-contact sport, not an academic exercise. The closing slide points students toward MIT's entrepreneurship resources.

The visible roadmap includes participation counts:

$$50{,}000+, \quad 5{,}000+, \quad 2{,}000+, \quad 350, \quad 72. \tag{10.46}$$

The small labels are dense, so the notes should preserve the screenshot and summarize the large movement: inspiration, exploration, fundamentals, application, and acceleration. Bob also names

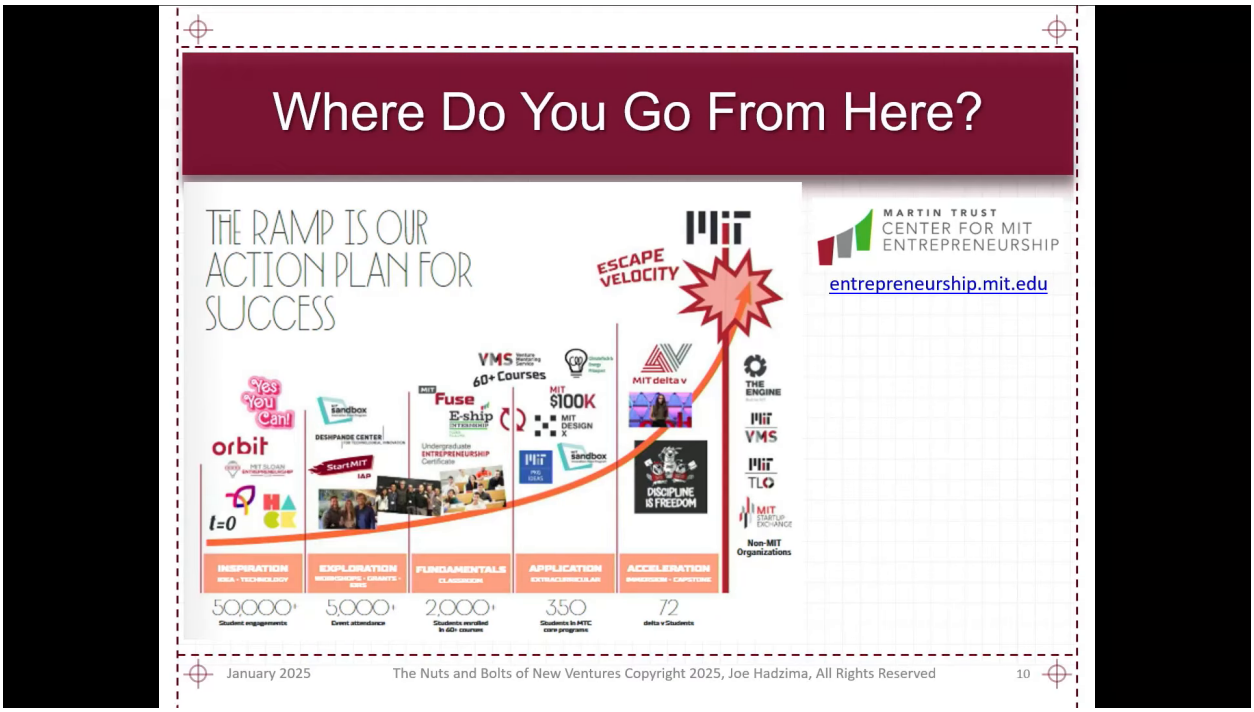

Figure 10.6: MIT entrepreneurship roadmap. The detailed logos are best preserved as screenshot evidence rather than redrawn.

the Venture Mentoring Service as a resource where entrepreneurs receive homework, mentor teams, and regular meetings.

The final formula is Joseph Hadzima's course heuristic:

$$H = \frac{R}{E}. \tag{10.47}$$

Here H is happiness, R is reality, and E is expectations. Bob's interpretation is practical: ask whether we can create a reality that exceeds the expectations people have. The same fraction applies to customers, investors, employees, and founders themselves.

10.9 Summary

The lecture begins with a case and ends with a course-wide rule. In between, the argument is consistent. Deep-tech entrepreneurship starts with protected differentiation, then tests market reality without giving away the core secret. We should expect fog, incomplete information, and wrong turns. The work is to build enough proof to learn, enough customer contact to steer, enough team conflict to decide well, and enough financial discipline to survive.

The central update is the Z Corporation pricing decision:

$$(\$20{,}000,\ \text{low-end entry}) \xrightarrow{\text{new evidence}} (\$60{,}000,\ \text{premium speed value}). \tag{10.48}$$

That is the lecture in miniature. A startup is not a fixed plan. It is a protected technical possibility learning its way through the market.

Chapter 11

Winning SBIR Contracts

This lecture is about the interval before traction: the awkward, expensive region where a technical invention is real enough to matter, but not yet reduced to a product that customers can buy, test, or even fully understand. Steve Derezinski presents SBIR not as a magic pool of government cash, and not as a form-filling exercise, but as a funding instrument with a particular logic: incentives, deadlines, agency variation, eligibility rules, and compliance obligations. The lecture belongs in the practical sequence of *Nuts and Bolts of New Ventures*, credited to Joseph Hadzima and MIT OpenCourseWare, with this edition curated by LazyingArt LLC through Video2Book.

11.1 Why This Lecture Exists

The opening promise is deliberately attractive: free government money. But almost immediately the lecture narrows the claim. The subject is not how to fill out SBIR forms. The subject is how SBIR awards are actually won.

Derezinski explains that he built the lecture because he saw a gap at MIT. There was administrative advice, and there were consultants, but not enough direct discussion of the practical winning process. He then gives the students a small piece of evidence from the prior year: the applicants he knew about won awards, and one company instead raised a $5 million venture round. That anecdote is not a theorem. It is a way of saying: this topic can matter materially to a young technology company.

His background supplies the operating lens. He was trained as a mechanical engineer at MIT, worked as a roboticist, patented an algorithm for reducing vibration in the space shuttle robotic arm, went through Sloan, helped create a faculty venture studio at Georgia Tech, and later worked in Kendall Square forming companies from technology. The transcript gives two different federal-contract figures, roughly $40 million and a $23 million list, so we should not combine them into a single precise number. The safe point is that he has substantial direct experience as a PI or business lead on federal contracts with agencies including NASA, NSF, DOE, ARPA-E, and DOD.

The first rule of the lecture is therefore not about SBIR specifically. It is about money generally:

$$\text{money is not just cash; it is a constraint system.} \tag{11.1}$$

Different money has different color. It brings different motivations, different reporting duties,

different time constants, and different failure modes. Once we see that, the rest of the lecture becomes an exercise in matching the stage of the venture to the right kind of money.

11.2 Disruptive Technology And The Funding Gap

The first substantive distinction is between incremental and disruptive technology. If the product is incremental, ordinary customer discovery can work rather directly. We ask customers what they want, and they answer in the coordinates of what they already have: cheaper, faster, higher quality. Their present product gives them the language for the next product.

Disruptive technology has a different problem. Customers may not know what to ask for until they can see it, touch it, test it, or imagine its use. The lecture invokes the familiar examples of the faster horse and of customers who do not know what they want until they are shown. The point is not to celebrate mythology about visionaries. The point is a practical one:

$$\text{disruptive technology} \quad \Longrightarrow \quad \text{customer understanding often requires reduction to practice.} \tag{11.2}$$

Now the funding gap appears. Universities and government research programs can support discovery, especially when the application is not yet clear. Private capital can support expansion when a company has customers, traction, and financial history. Between these two regions lies the hard interval: the work needed to turn a paper, laboratory result, or prototype into something the market can evaluate.

11.2.1 Question & Answer

Question. If customers cannot yet understand the product, how can a startup validate or fund it?

Answer. The lecture's answer is that we should not pretend that ordinary customer traction already exists. The first task is to reduce the technology to practice far enough that customers can respond intelligently. That is the stage where Derezinski places SBIR-type funding: not as a universal best source of startup capital, but as a practical instrument for deep-technology work before repeat customer traction.

The local obstacle can be written as

$$\text{paper or lab result} \not\Rightarrow \text{customer-ready product.} \tag{11.3}$$

There is a missing middle term: reduction to practice. The lecture argues that this missing term is better matched to grants and non-recourse funding than to recourse capital.

11.3 The Research-To-Market Boundary

The central slide of the lecture turns the funding gap into a stage map. The visible process is

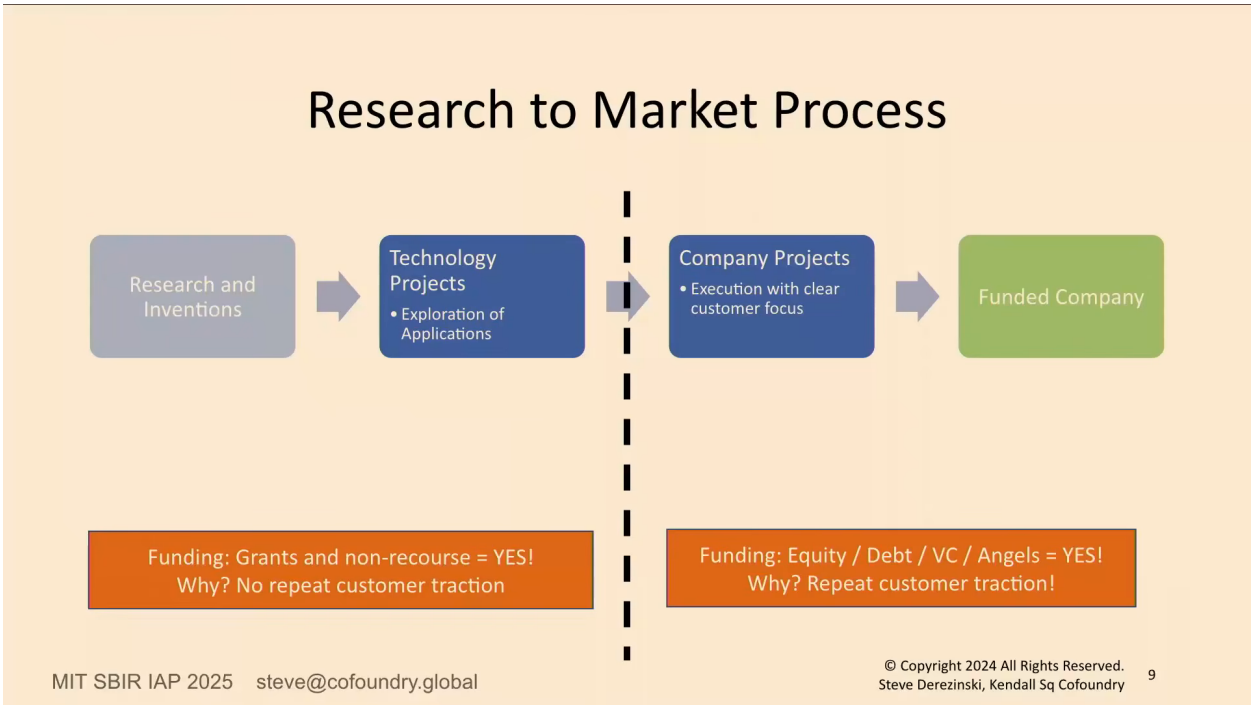

Figure 11.1: Research-to-market funding boundary. The screenshot is the lecture's direct visual evidence for the distinction between technology exploration, customer-focused company execution, and the funding types appropriate on each side of the traction boundary.

$$\text{Research and inventions} \to \text{Technology Projects} \to \text{Company Projects} \to \text{Funded Company}. \tag{11.4}$$

The two blue boxes carry the lecture's central distinction. A *technology project* is still exploring applications. It may contain serious research, but there is not yet a real customer who understands it, has engaged with it, knows how to buy it, or knows how to interact with it. A *company project* comes later: there is a demo, something people can use, a clearer customer focus, and a need to refine, debug, package, and execute.

The dashed boundary in the slide is not labeled, but the transcript supplies its meaning. It is the boundary between no repeat customer traction and repeat customer traction. A compact reconstruction is:

The funding rule is the payload:

$$\text{no repeat customer traction} \Rightarrow \text{grants and non-recourse funding}, \tag{11.5}$$

$$\text{repeat customer traction} \Rightarrow \text{equity, debt, VC, angels, or growth capital}. \tag{11.6}$$

This is not a universal law of entrepreneurship. It is a stage map. The warning is strongest on the left side. If we take recourse funding while we are still exploring applications, we create obligations before we know the direction. The company owes performance to somebody, but the market has not yet told us what performance means.

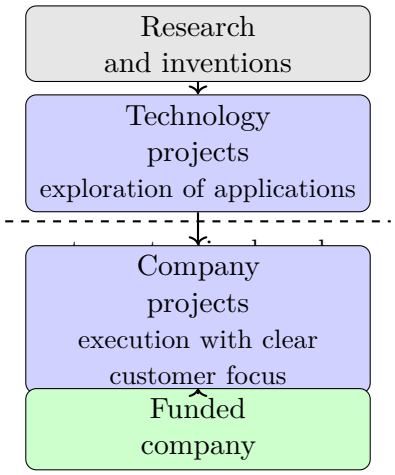

Figure 11.2: A narrow reconstruction of the lecture's research-to-market process for pocket layout.

On the right side, the situation changes. When the pump is already working, when customers are present and repeat traction exists, growth capital can be the right fuel. The same dollar can be sensible or dangerous depending on which side of the boundary the company occupies.

11.4 Incentives: Why SBIR Is Not Venture Capital

Once the boundary is in place, the lecture quantifies SBIR. The annual SBIR pool is described as roughly

$$F_{\text{SBIR}} \approx \$4\text{B per year.} \tag{11.7}$$

Derezinski then makes a scale comparison to venture capital. His heuristic is that venture funds deploy about one fifth of committed capital per year. If a government program distributes about $4 billion per year, the venture-capital-equivalent fund size is

$$F_{\text{VC-equivalent}} \approx \frac{\$4\text{B}}{1/5} \tag{11.8}$$

$$\approx \$20\text{B}. \tag{11.9}$$

This calculation is deliberately simple. Its purpose is not to say that SBIR is venture capital. Its purpose is to make the size visible before we compare incentives.

The lecture then gives the incentive comparison. Angels invest their own money and can decide quickly. Venture capitalists invest limited partners' money, so they are pushed toward high returns, meaningful equity ownership, and time-sensitive exits. The return pressure can be written schematically as

$$\text{VC pressure} \sim \text{large return over limited time.} \tag{11.10}$$

If the time stretches, the required financial outcome must become much larger. That is one reason venture capital can be a poor fit for pre-traction technology exploration.

SBIR has a different structure. In the ordinary grant-like sense, the government does not expect repayment. It also does not require the company to become commercially successful in order for the Phase I work to have been legitimate. The expectation is narrower: do the proposed work, use the funds properly, and report the result. Winning an SBIR may later help private diligence, because external investors may treat it as a signal that technically serious reviewers found the work credible. But the SBIR itself is not venture money.

11.4.1 Question & Answer

Question. What is the first motivation of an SBIR program manager?

Answer. The class guesses science, job creation, and proper use of funds. Derezinski accepts those as real motivations, but gives the first motivation as: do not get fired.

This is the lecture's clearest tension-and-resolution moment. It explains why the bureaucracy is not merely decorative. If a proposal is excellent but violates an instruction, the program manager may not be able to repair the defect after submission. The safe administrative move is rejection. Thus formality follows from incentives.

We can summarize the practical ordering as

1. Do not create a compliance problem for the program manager.
2. Do not embarrass the program manager by appearing careless or unserious.
3. Then make the technical case that the work advances the agency's agenda.

This is why instructions, eligibility, formatting, and the official solicitation PDF matter. A proposal is not only a technical argument. It is also evidence that the company can be trusted with public funds.

11.5 Agencies, Phases, And Program Shape

The next movement is from funder incentives to agency architecture. SBIR has one name, but not one uniform implementation. Derezinski notes that NSF says SBIR, while DOD speakers may pronounce it SIBR. The pronunciation is trivial; the underlying lesson is not. Same name, different agency behavior.

He then splits agencies into two practical buckets: market-driven and mission-driven.

Market-driven solicitations may sound broad: do something important in optics, energy, software, or another technology area, and show that the market cares. Mission-driven solicitations can be much more specific: a battery with a particular pack constraint, operating temperature, duration, weight, and field condition. The specificity is both burden and opportunity. It is harder to match, but it may also reduce the number of serious competitors.

Agency type	Where the customer sits	Practical implication
Market-driven	Outside the agency, in the marketplace	The proposal must make a credible case that someone will buy, adopt, or validate the technology. NSF and DOE are given as examples.
Mission-driven	Inside the agency or its contractor ecosystem	The solicitation may specify a narrow operational need. DOD and NASA are given as examples.

Table 11.1: The lecture's operational distinction between market-driven and mission-driven SBIR agencies.

11.5.1 Question & Answer

Question. Is the market-driven versus mission-driven split just the lecturer's interpretation?

Answer. The transcript answers that the framing is tied to America's Seed Fund and NSF material, but the chapter should treat it as an operational taxonomy rather than a statutory classification. It tells founders where to look for the customer, how broad the solicitation may be, and how much matching work the proposal must do.

The phase structure then adds the next layer of operating rules:

$$p(\text{Phase I win}) \approx 10\%\text{–}15\%, \tag{11.11}$$

$$p(\text{Phase II win} \mid \text{Phase I completed}) \approx 50\%. \tag{11.12}$$

Phase I is the entry point. One must win Phase I to qualify to apply for Phase II. Phase III is commercialization tracking; in the lecturer's framing, there is no ordinary federal Phase III funding. The exception he emphasizes is a Phase IIb-type matching program.

Let P be qualifying private recourse investment and $M(P)$ be the federal match. The lecture states a two-to-one private-to-federal match, capped at \$500,000. The compact rule is

$$M(P) = \min\left(\frac{P}{2}, \$500{,}000\right). \tag{11.13}$$

The examples are immediate:

$$P = \$1{,}000{,}000 \Rightarrow M(P) = \$500{,}000, \tag{11.14}$$

$$P = \$250{,}000 \Rightarrow M(P) = \$125{,}000. \tag{11.15}$$

The private investment must be real recourse funding. The lecture explicitly rejects self-loan games: one cannot lend money to oneself, collect the match, and repay oneself. Investor letters and investment structure matter because they prove that the private money is genuine.

Timing also enters here. The federal fiscal year closes on September 30, and a new one begins October 1. Many processes are driven by the need to allocate awards before the year closes, even if

some activity spills over. Mission-driven agencies may also use pre-release windows. During those windows, founders may be able to speak more freely with program managers; once the topic is locked down, questions may have to be posted publicly.

11.6 Eligibility, Multiple Applications, And SBIR Versus STTR

Eligibility is not merely a checklist. It is a timing problem. The basic constraints in the lecture are

$$\text{employees} < 500, \tag{11.16}$$

$$\text{U.S. citizen or permanent-resident ownership} \geq 51\%. \tag{11.17}$$

The company must be for-profit, and both startups and established companies can apply. Derezinski adds a practical observation from older DOD winner data: one-person companies can win, but the more natural early-company shape is often two to nine people, with some mixture of technical and business capability.

The timing rule is central:

$$\text{eligibility is determined at time of award, not merely at submission.} \tag{11.18}$$

That matters for students, postdocs, and founders still inside a university. One may describe a plan to form or join the company by the time of award, but the team, employment, and budget commitments must become true when the award is made. The PI need not necessarily have a PhD or MD, but must have the technical expertise to oversee the scientific and technical work. On citizenship and visa questions, the lecture is explicitly cautious. The company ownership rule is clear; the PI visa details should not be turned into legal advice.

11.6.1 Question & Answer

Question. Can we apply to multiple agencies, or does that create a duplicate-award problem?

Answer. The distinction is the scope of work. If one proposal funds software development and another funds a different technical component of the same eventual product, those may be distinct scopes. If the same work is copied from NIH to DOD with only the agency name changed, that is the problem.

A useful notation is

$$\text{scope}(P_i) = \text{the work actually funded by proposal } P_i. \tag{11.19}$$

The duplicate-award constraint is

$$\text{scope}(P_i) = \text{scope}(P_j) \quad \Longrightarrow \quad \text{do not accept both awards.} \tag{11.20}$$

Different scopes may be disclosed and separately pursued. The same work cannot be funded twice. The lecturer treats this not as a clever probability game, but as a compliance boundary. The loose intuition that two 15% chances may feel like 30% should not be formalized as independent probability.

Question. Do key team members have to be inside the company?

Answer. If a person is named, budgeted, and used as part of the credibility of the proposal, then the company must be able to deliver that person at time of award. Generic later hires are different. A budget line for unspecified software developers is not the same as a named expert whose biography helped the proposal win.

Question. How should founders handle incorporation and university conflict of interest?

Answer. The lecture deliberately does not give a universal answer. Registration often requires some entity or registration path, but university conflict-of-interest issues depend on the lab, department, institution, and technology rights. The right note is caution, not invented legal advice.

The SBIR/STTR distinction adds another layer. SBIR permits a research institution partner; STTR requires one. The lecture gives the allocation arithmetic as follows:

$$\text{STTR small business share} \geq 40\%, \tag{11.21}$$

$$\text{STTR research institution share} \geq 30\%. \tag{11.22}$$

The minimum required shares sum to

$$40\% + 30\% = 70\%, \tag{11.23}$$

leaving roughly 30% of the budget flexible within program and agency constraints. This is one reason STTR can sometimes be strategically useful: more hoops may mean fewer applicants. But the hoops are real.

The overhead example gives the hidden cost. If a \$300,000 Phase I award sends 30% to a university, then

$$0.30 \times \$300{,}000 = \$90{,}000. \tag{11.24}$$

Derezinski warns that overhead may reduce the amount available for actual research to something like

$$\$90{,}000 \Rightarrow \text{roughly } \$45{,}000 \text{ or less for the researcher.} \tag{11.25}$$

This is not a formal accounting identity. It is a practical warning. A university partner can be worth the cost when equipment, expertise, or credibility are essential, but university participation is not free.

11.7 Proposal Mechanics, Letters, Quad Charts, And Closing Discipline

The final movement of the lecture is operational. A proposal is typically about 25 pages, and for NSF Derezinski points especially to the project description and letters of support. The project description is described as 15 pages and includes an elevator pitch, commercial opportunity, company and team, technical solution, and technical-merit R&D plan. For Phase I, the center of gravity should remain the technical solution and the R&D plan.

The review hierarchy is simple:

$$\text{technical merit first} \quad \text{then} \quad \text{commercial merit and firm qualifications.} \tag{11.26}$$

A young company may not yet have much firm history. That is why smaller awards or team accomplishments can matter. They show that a group can act as a cohesive unit, not merely as unrelated people under a new name.

Letters of support reduce reviewer uncertainty. A strong customer letter says, in effect, that the technology would be valuable if reduced to practice in a particular way. A strong investor letter shows serious understanding of both the market opportunity and the technical risk. A strong partner letter shows a plausible path to market access or later investment. Weak letters come from people with little substantive connection to the work: paid consultants who obviously benefit from the award, investors with no technical understanding, or political supporters with little relevance.

The reviewer has many proposals to read. The easiest answer is no. Therefore the proposal should make yes easy:

$$\text{review burden high} \quad \Rightarrow \quad \text{make credibility visible inside the proposal itself.} \tag{11.27}$$

The lecture then turns to the quad chart or project pitch. This is the recommended first outreach instrument. Before spending too much money on company formation or administrative machinery, prepare a compact description of the technology, objectives, market opportunity, and team. NSF has a project-pitch process; other agencies commonly respond to a quad chart. Derezinski calls the quad chart a kind of love language for the SBIR community because it packages the opportunity in a format program managers can digest.

11.7.1 Question & Answer

Question. What kind of feedback can a program manager give before submission?

Answer. The program manager may not give tailored advice that makes the proposal more competitive. The legal boundary varies by agency and stage, and the lecture does not try to define it precisely. What the founder can seek is fit feedback: does this look like the kind of work the agency or topic is meant to fund? A positive answer is useful. A lukewarm answer may still be useful. A fast no is also valuable, because it prevents wasted proposal effort.

The workflow is sequential:

1. Identify agencies and topics that could fund the technology.

2. Search for previously funded awards similar to the project.
3. Find the official solicitation PDF and the relevant program contact.
4. Prepare the quad chart or project pitch.
5. Ask the program manager for fit feedback, not unfair proposal coaching.
6. Ask who else should be consulted.
7. Then proceed with registration, company administration, and full proposal work.

The last question, who else should we talk to, is not small talk. The person named may be a future reviewer, an internal advocate, or a technical gatekeeper. The lecture includes a cautionary anecdote: ignoring such a referral can cost the proposal.

The timing arithmetic is part of the discipline:

$$T_{\text{broad proposal}} \approx 2 \text{ months}, \tag{11.28}$$

$$T_{\text{specific mission fit}} \approx T_{\text{broad proposal}} + 1 \text{ month}, \tag{11.29}$$

$$T_{\text{registrations}} \approx 30 \text{ days}. \tag{11.30}$$

Registrations can run in parallel with proposal writing, but they should not be ignored. The official solicitation PDF is the ground truth. Consultants, summaries, templates, and advice may help, but the legal and procedural authority is the solicitation itself. If the solicitation does not answer the question, ask the program manager.

Two final cautions close the lecture. First, AI can help draft, but it does not replace judgment. Derezinski's rough grading metaphor is

$$\text{AI} : F \to B, \qquad \text{human judgment} : B \to A/A^{+}. \tag{11.31}$$

That is, AI may improve a bad draft into a serviceable draft, but funded proposals still require technical specificity, relationship-building, budget sanity, and real credibility. People fund people.

Second, diligence increases as the money increases. For Phase II, the company may need to show an accounting system, timecards, proper allocation of labor, proof, and sign-off. The transcript's audit acronym is uncertain, but the operational requirement is not:

$$\text{government money} \Rightarrow \text{allocation evidence, time records, accounting control, and sign-off.} \tag{11.32}$$

Fast-track applications receive the same practical treatment. They can be appropriate when commercial partners are already lined up and the solicitation fit is unusually strong. But because many applicants want to skip directly toward the larger money, competition can be tougher. This is exactly the sort of choice to test with the program manager.

11.8 Summary

The lecture begins with the promise of free government money and ends with homework. That arc is the point. SBIR is not free in the sense of effortless. It is non-recourse funding with rules,

deadlines, incentives, and compliance obligations.

The central map is the customer-traction boundary. Before repeat customer traction, a deep-technology company may need grants or non-recourse funding to reduce the invention to practice. After repeat traction, recourse capital can become appropriate growth fuel. SBIR sits in the gap, but each agency implements it differently, and each solicitation must be treated as its own ground truth.

The practical prescription is disciplined rather than glamorous: understand the agency, understand the phase, disclose related applications, avoid duplicate scopes of work, be careful with university partners, write for technical merit, use meaningful letters of support, prepare the quad chart, speak to the right program manager, and treat a fast no as useful information.