Working Paper 25-031

Charting (and Updating) the Path: A Bayesian Perspective on Entrepreneurial Learning

Joshua L. Krieger

Charting (and Updating) the Path: A Bayesian Perspective on Entrepreneurial Learning

Joshua L. Krieger
Harvard Business School

Working Paper 25-031

Funding for this research was provided in part by Harvard Business School.

Charting (and Updating) the Path: A Bayesian Perspective on Entrepreneurial Learning *

Joshua L. Krieger

December 1, 2024

Abstract

Abstract: This chapter explores two distinct modes of entrepreneurial learning: assessing venture viability and choosing between alternative development paths. It introduces a framework for decomposing venture viability into technological feasibility, commercial potential and development costs, while highlighting how entrepreneurs and investors may have divergent beliefs about these components. The chapter then presents a novel approach for visualizing and analyzing multiple potential development paths, incorporating factors such as independent project characteristics, learning spillovers, and development distances between paths. By integrating these elements, the framework allows entrepreneurs to make more informed decisions about initial paths and subsequent pivots, moving beyond simple heuristics to a more comprehensive understanding of tradeoffs and relative values across different development strategies.

*This chapter was prepared as part of the book *Bayesian Entrepreneurship* (forthcoming from MIT Press), edited by Ajay Agrawal, Arnaldo Camuffo, Alfonso Gambardella, Joshua Gans, Erin Scott and Scott Stern. I am grateful to Jim Matheson and the students from the Harvard Business School course Tough Tech Ventures for helpful comments.

1 Introduction

Entrepreneurs have a bias for action because entrepreneurial learning requires generating new information. The entrepreneur creates new evidence through actions—e.g., building prototypes, talking to suppliers and customers, running scientific experiments, surveying potential customers, etc.—that update his or her theories around an opportunity and/or level of conviction. However, the type of information produced and its subsequent interpretation vary greatly depending on the nature of the entrepreneur's *motivation to learn*. While skilled entrepreneurs excel at systematically resolving uncertainty and persuading others to believe in them, what evidence they seek depends on their personal convictions as well as their beliefs about the audiences they aim to influence. The venture's early experimentation will differ based on whether the key uncertainties lie in the validity of the venture's underlying (contrarian) theories vs. the venture's ability to outperform competitors within the boundaries of more commonly held beliefs. Like a navigator charting the path in through unseen terrain, the entrepreneur not only has to choose the initial path, but must determine what signals will warrant further commitments to continue, shift directions or turn around entirely.

To examine how entrepreneurs update their beliefs about their venture, we first must consider the entrepreneur's motivation for gathering information. This chapter groups those motivations into two broad categories: de-risking the venture's *viabiliity* and choosing from *alternative development paths*. As argued below, those motivations and the strength of conviction behind the entrepreneur's prior beliefs shapes the type of activities on the venture's roadmap. Furthermore, the inevitable asymmetries between the venture team's beliefs and other stakeholders' beliefs alter the nature and sequencing of information generating activities.

By comparing the distinct motivations of entrepreneurial experimentation, this chapter aims to provide a common language and framework for venture investment and sequencing choices. Too often, entrepreneurs and investors justify venture development paths based off simple heuristics that prioritize a single element of commercial potential (e.g., total addressable market, attractive "beachhead" markets, experimentation costs, etc.). Rather than replacing those heuristics with a new one, the framework described below uses the concepts of Bayesian entrepreneurship to integrate these elements, highlight the tradeoffs between them, and quantify the relative value of different development paths.

2 Defining Opportunities and Venture Beliefs

In articulating their expectations around any new venture, entrepreneurs and investors will often bundle beliefs of varying specificity and criticality (to the venture) into broad categories of risk and rewards.[1] Further, entrepreneurs may conceive of "opportunities" quite differently. Some might describe broad applications of a new technology (hammer with plausible nails) or as a problem "pain point" to be solved with all available tools (nails looking for hammers), while others might define their venture with detailed descriptions of particular business models (including value proposition, technology and operations, go-to-market, and cash flow plans). However, the process of learning about the venture depends on underlying assumptions about which and whose beliefs the entrepreneur aims to challenge and update through information gathering.

To fix ideas, I discuss an *entrepreneurial opportunity* as specific solution-problem pairs (e.g., robotic hands for packaging cleaning supplies). Thus, opportunities must embody a set of both technological and commercial hypotheses—though the degree of uncertainty may vary independently. Entrepreneurs, investors and potential customers might also vary in their baseline beliefs about the "fit" of any particular solution-problem pair.[2] "Adjacent opportunities" take the same or similar technologies and apply them to a distinct commercial application. The differences between and two adjacent applications may be subtle or obvious, but what makes them distinct is some divergence in either in their technological development needs or commercialization steps.[3]

Thus, in discussing learning venture viability, the chapter will focus on updating beliefs over a particular opportunity, (probability of success (p_i), cost ($CostDev_i$), and profitability conditional on success (V_i)), as well as the venture's overall viability (p_{overall}, $CostDev_{\text{overall}}$,

[1]A central challenge of generalizing about entrepreneurial experimentation is that the core beliefs behind a venture vary greatly across industries and technologies, and entrepreneurs rarely articulate those beliefs as precise hypotheses to be tested.

[2]For example, an individual may see merit in addressing a given problem and may even think the proposed solution would work, but still not find the match of the two compelling (e.g., "willingness to pay for solving that warehouse problem is sky high, but I've seen robotic systems that look more effective than this design.").

[3]For simplicity in this chapter, I abstract away from (the many) business model choices that might exist within the given opportunity or across the set of related adjacent entrepreneurial opportunities—e.g., degree of vertical integration, payment models, licensing, and sales channels—and assume that each opportunity is tied to a particular dominant business model. However, one could also represent "adjacent" opportunities as similar problem-solution pairs that differ mainly in their business model hypotheses, if those business model distinctions require a different set of commercial development activities.

potential V_{overall} across all target applications).[4] The venture's "overall" expected values reflect either the sum of its various opportunities, if all are viable, or the expected value of the subset of opportunities that the entrepreneur expects to pursue. Meanwhile, the choice of alternative development paths involves the subset and sequence of i's that maximize the value of the opportunity, including learning spillovers and option value.

3 Venture Viability

The venture's "viability" refers to beliefs about whether pursuing the venture is a bet with positive net present value. Those beliefs are distinct from the question of whether a given entrepreneur—with his or her own skills, opportunity costs, access to resources and risk preferences—*should* chose to pursue that venture in the first place. Aside from that personal choice, individuals might vary in their forecasts about the venture's viability due to factors like asymmetric information about the a particular opportunity and overall venture potential, or different understanding of how key resource providers might view an opportunity or overall venture. Thus, a venture's viability is in the eye of the beholder, but these beliefs can still be decomposed into distinct components, each of which is subject to Bayesian updating.

3.1 Components of Venture Viability

Viability encompasses multiple subsidiary beliefs, which can be grouped into three main groups: technological, commercial, and development cost.

Technological feasibility is the belief of whether or not the proposed problem-solution pair will "work." That feasibility threshold might depend on whether a novel technology will achieve it's most basic goals (e.g., does the flying car actually take off and land?), or whether the technology or service can achieve a minimum performance threshold (e.g., will the customer service chatbot answer questions as quickly as a human employee?). p_i summarizes the technological feasibility belief as the likelihood that the venture can deliver on this minimum performance threshold for a particular opportunity, with the assumption that the entrepreneur commits to maximizing the venture's technological potential for that given opportunity. The "strength" of this belief is reflected in the distribution of p_i.

[4] V here is a payoff or profit measure, conditional on the venture achieving a reasonable threshold of "success" in terms of both it's initial technical hypotheses and ability to gather sufficient resources to develop the venture.

Commercial viability implies a similar set of questions about the economic potential of the venture, but assuming all technological uncertainty has been resolved favorably. Will the willingness to pay exceed marginal costs? What is the break-even scale of production? What is the ultimate market size of the best version of the opportunity ($\max_{i \in \{1,2,\ldots,n\}} V_i$) or cluster of adjacent opportunities ($V_1 + V_2 ... + V_N$). Again, the entrepreneur may be more or less certain about the commercial potential, as reflected in the spread of the distribution(s) of each V_i.[5] The entrepreneur (and investors) must consider the time until the venture can receive payoff V, such that the expected payoff today is $\text{NPV} = \frac{V}{(1+r)^t}$, where t is the expected time until the payoff can be realized and r is the discount rate.[6]

Development Costs ($CostDev_i$) refer to the company building efforts necessary to pursue the venture "in earnest" (i.e., with full commitment to at least one particular opportunity and to the best of the entrepreneur's abilities). Even if technological and commercial feasibility beliefs are strong, the entrepreneur must still evaluate the cost of capital and time needed to set up the fully operational venture. Holding the nomical $CostDev_i$ constant, the timing and shape of a venture's development cost curve can vary (e.g., front-loaded capital expenditures vs. steady increase in development costs over time), such that the true cost of development is expected to be $CostDev_i = \sum_{t=0}^{N} \frac{c_t}{(1+r)^t}$, where c_t is expected development cost in period $t = 0$ for a development path of N total periods. Long timelines or and large opportunity cost of locking up significant capital might limit the set of potential investors, even if those investors share the entrepreneur's beliefs about the potential for the mature version of the venture. For example, if a scientist that invents a compelling early diagnostic test for dementia might face a very long road to commercialization, with multi-year clinical trials and regulatory scrutiny. Even a venture capital investor convinced of the large market opportunity might pass on the investment due to capital requirements and fund timeline considerations.

To summarize, a venture opportunity's viability can be broken into three main components of viability that together give a rough sense of the opportunity's Net Present Value:

$$VentureValue_i = NPV(p_i \times V_i - CostDev_i) = p_i \times \frac{V}{(1+r)^t} - \sum_{t=0}^{N} \frac{c_t}{(1+r)^t} \tag{1}$$

[5] The shape or skewness of the distribution of V_i likely differs by industry. For example, some fields might tend to have "longer tailed" payoff distributions if they tend to have a small number of "blockbuster" winners.

[6] V itself could also be considered a time-discounted expected value based on the forecasted stream of cashflows produced by the venture starting in time t.

3.2 Testing Viability

Learning about the venture's viability may involve focused experiments to resolve uncertainty about a particular component of viability for one (or more) opportunities, or gathering information that informs multiple elements of viability.

For example, the founders of startup Rest Devices wanted to understand the commercial value of embedding their thin flexible sensors within baby sleepwear to sleep and breathing patterns. The application was attractive, but was not the team's natural passion or area of expertise. Lead by three MIT graduate students, none of which had children or experience in the baby gear market, they had initially believed the entrepreneurial opportunity was in sensors for sleep apnea, before they learned that that commercial path would require a long and expensive regulatory gauntlet (Ghosh and Payton, 2016). The feasibility of the baby wear sensors venture would rely on a number of factors: could they produce accurate sensors and accompanying monitoring software that works reliably (technical)? Could they convince parents of infants to purchase their product at a profitable margin (commercial)? And assuming it was technically feasible, could they build out inventory and customer service systems to properly serve retailers and mass market customers (cost of development)?

The Rest team's early testing involved different activities to address the various components of the venture's viability. They created prototype versions of the baby wear onesie outfits, equipped with sensors and a connected smartphone app, and tested it with the permission of some early-adopter parents. On the commercial side, they conducted an online survey and follow-up phone interviews of new parents to gauge interest and willingness to pay, and they brought their prototype to trade shows to get feedback and sales opportunities with retailers. In other words, their early evidence generating activities were multi-pronged, testing multiple components of their viability expectations.

In other cases, testing can be more concentrated in a particularly critical component of viability. For example, Vaxess Technologies had a design for a new type of shelf-stable microneedble patch for vaccine delivery (Quelch and Rodriguez, 2014). The idea was rooted in published research and intellectual property from Tufts University and MIT. Despite academic research to support their approach, the viability of the venture would primarily depend on two sub-elements of technical feasibility: producing the tiny silk protein-based tips that painlessly penetrated the skin and slowly released the vaccine, and showing their safety and efficacy in clincal trials. Early experimentation focused on engineering and optimizing the

3D printing the microneedle tips, and then testing their efficacy in mouse and monkey studies. The focus on technical feasibility did not mean that Vaxess had no commercial questions. Indeed, even if the vaccines proved effective and manufacturable at scale, the company would have to navigate partnerships, distribution, pricing, marketing and negotiations with insurers and government health agencies. But the greatest uncertainties (weakest or most dispersed priors) were associated with the technology and the team had confidence that a vaccine market would still be there and provide sufficient demand for high-performing vaccine delivery technologies.

3.3 Viability Belief Differences Across Key Parties

In addition to forming one's own priors over p, V, and $CostDev$, the entrepreneur also tries to understand the comparable beliefs held by key stakeholders and resource providers. That group of other players includes potential investors, customers, employees, suppliers, and regulators. I follow Chapter 1 and focus on investors as the representative external party.

Entrepreneurs and investors might have divergent beliefs about the key components of the venture's viability. Those may arise from differences in their previous experiences, as well as different access to information or expertise. By virtue of *choosing* to pursue (and persuade others about) their venture, the entrepreneur is relatively more optimistic than others about the opportunity. Frontier technology startups that need to prove viability related to both their technology and commercial prospects face an even more daunting challenge in convincing investors to take on multiple types of uncertainty and evaluate technologies more likely to lie outside of their expertise (Lerner and Nanda, 2020; Arora et al., 2024; Nanda, 2020).

That gap in beliefs is not a problem if both parties still believe that the business is viable—i.e., sufficiently high $pV - CostDev$ to warrant investment. However two possible scenarios emerge if the investor's viability beliefs are worse than the entrepreneur's and lower than the investor's threshold for investment.

First, if the entrepreneur holds "strong" priors (e.g., narrow distributions around their expected values of p, V, and $CostDev$), then the entrepreneur's main motivation in generating updating/learning is *persuasion*. See Figure 1 for a chart describing the the distinct "modes" of entrepreneurial learning. The process of generating new information in service of that goal may or may not align with the development roadmap the entrepreneur would have followed without any resource constraints. To the extent that persuasion activities diverge from that optimal roadmap, the strong prior entrepreneur will want to take on experiments that directly address the component of viability where the entrepreneur and external party have their greatest disagreement, and do so with minimal cost (e.g., "lean startup" methodologies).[7] However, persuasion activities need not take the company off it's efficient roadmap. If persuasive activities are aligned with the natural company building roadmap, then the strong prior entrepreneur will want to maximize the size/scope of experiments (e.g., experimenting at scale, larger prototypes, expanding manufacturing earlier), subject to resource and investor patience constraints.

The second situation is that the entrepreneur holds "weak" priors (e.g., wide distributions around their expected values of p, V, and $CostDev$). Despite holding more optimistic viability beliefs than the investor, the two parties now share the common goal to efficiently retire uncertainty and maximize the option value of their investments. Practically, that means striving to run cheap experiments focused on increasing precision over the venture's overall viability, targeted at efficiently maximing learning for cost of that experiment/activity. Updated beliefs may result in either a quick termination of the venture, or the entrepreneur strengthening his or her convictions and the investor updating positively.

Figure 1: Entrepreneurial Learning Modes

	Experimenting for... Viability	**Experimenting for... Optimal Path**
Strong Priors	**Persuasion**	**Updating Development Path**
Weak Priors	**Learning to Kill or Commit**	**Pure Exploration / Messing around**

[7]Notably, persuading the investor is not merely about moving mean expectations of viability, but may also be about increasing the investor's expected variance (e.g., probability of long-tail payoff outcomes) such that the investor is willing to embrace the venture's option value in it's portfolio.

4 Choosing from Alternative Development Paths:

Learning that a venture is a worthwhile pursuit to initiate, continue or finance is only one “mode” of entrepreneurial learning. Success in venture building also depends on ability to navigate development towards the most promising directions possible for a given venture. The previous section highlighted distinct motivations for learning about a venture’s *viability*, and how those motivations map to different information gathering and persuasion needs. But many entrepreneurial activities aim instead to rule in/out *different development options for the given venture* (i.e., adjacent opportunity paths, i’s). Charting a venture roadmap with multiple potential development options is central to entrepreneurial choice, especially when the entrepreneur has already established conviction around the venture’s overall viability and committed to pursuing the venture in some shape or form.

Choosing paths is especially important for novel technologies with potential utility across multiple application domains. Demonstrating viability along intial paths might be a prerequisite for gaining access to adjacent opportunities and spuring complementary investments as “general purpose technologies” (Bresnahan, 2010). But initial path failures might also quell such opportunities for follow-on exploration. Learning about any one path has implications (spillovers) for learning about adjacent paths. In other words, the entrepreneurs is choosing between multiple *initial development paths* as well as possible *sequences* of development paths for adjacent opportunities.

To select the sequence of development paths, I propose a framework that addresses the independent viability of distinct paths, as well as the degree of development spillovers (adjacency) between paths. The framework requires estimating three project features:

1. **Independent project-path characteristics:** Each potential project path i has a development cost c_i, a probability of success p_i and a potential profits conditional on development success V_i.

2. **Learning spillovers between paths**: Any pair of two possible paths have a correlation in their probability of success (e.g., the probability of success on Path B, conditional on success in path A is represented as $S_{B|A}$), which is increasing in their R&D overlap and decreasing in their difference in relative cost/complexity ($c_B - c_A$).

3. **Development “distance" between paths**: The distance is the additional cost ($c_{A \to B}$) of developing the technology for the second application, after developing it for the

(given) first application. This distance mechanically follows #2 above. Knowing the relative direction and cost of both paths is sufficient to know their distance. Project pairs with tighter correlations ($S_{B|A}$ closer to 1) will have shorter development distances ($c_{A \to B}$, $c_{B \to A}$).

Visualizing paths: These features can be visualized as a set of vectors, where each vector represents a distinct project path (see Figure 2). The length of each vector represents the total cost (c_i) of developing the the venture for that particular opportunity/application. Those costs includes labor, capital expenditures, R&D, experimentation costs, legal, and commercial development (e.g., sales, marketing, demonstrations, etc.). At the end of each vector is a node, whose size represents the expected enterprise value (V_i) conditional on successfully developing the venture along that path.[8] Since initial estimates of V_i might be highly uncertain, one can also visually represent the nodes as a range of potential outcomes with concentric circles for different percentiles in the expected payoff distribution.The degree of development overlap between any two paths is reflected by the angle between their two vectors.

Cost of switching across paths: The cost of developing for a subsequent application (after completing development for an initial application) is simply the adjustment cost (distance) from the end of one path to another. For example, if the venture team initially develops along Path A and succeeds, then the cost of pursuing Path B next becomes $c_{A \to B}$, represented as the the straight line distance from the end of the first path to the end of the adjacent pathby, where $c_{A \to B} < c_B$. That cost can be further discounted by a cost-reduction factor γ if the first path is more complicated than the subsequent path ($c_A > c_B$).

The framework also allows for intermediate milestones along any given path (see Figure 2, Panel B). The distance to/between milestone points may be naturally defined by the cost of an opportunity (e.g., testing a prototype aircraft in a wind tunnel), imposed by regulation (e.g., a phase I dosage safety trial for a new cancer drug), or a choice that reflects that venture team's learning motivations and viability priors (see prior sections). Upon reaching a milestone, if the firm receives bad news about the active development path, the team can either terminate development (option to quit) or pivot to its preferred adjacent path. The

[8]More specifically, V_i is the net present value of cashflows of the venture, after it has established its viability. The level of experimentation and company-building activities needed to establish that viability will vary across settings and opportunities. The venture's products or services may still evolve after establishing viability, but here I assume that the entrepreneur has a prior over the expected captured value (i.e., serviceable obtainable market) that the mature venture can achieve along that given application path.

Figure 2: DEVELOPMENT PATHS FRAMEWORK

A. Selecting Vector Sequence

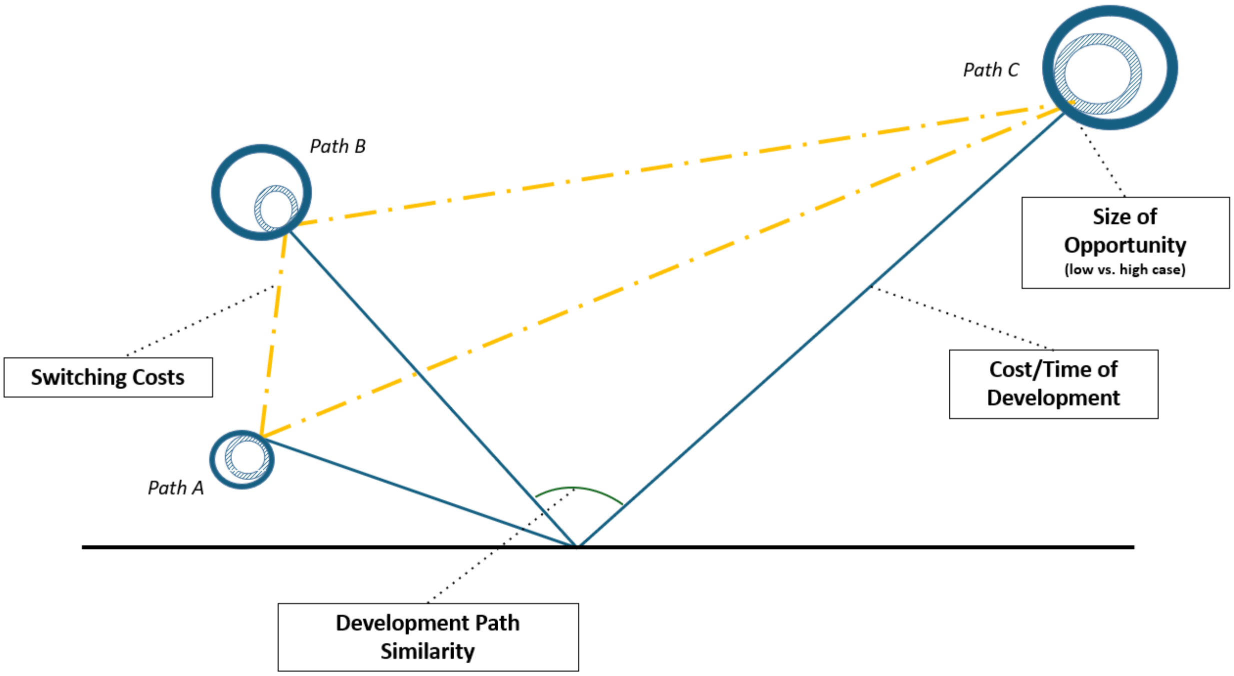

B. Including Intermediate Milestones

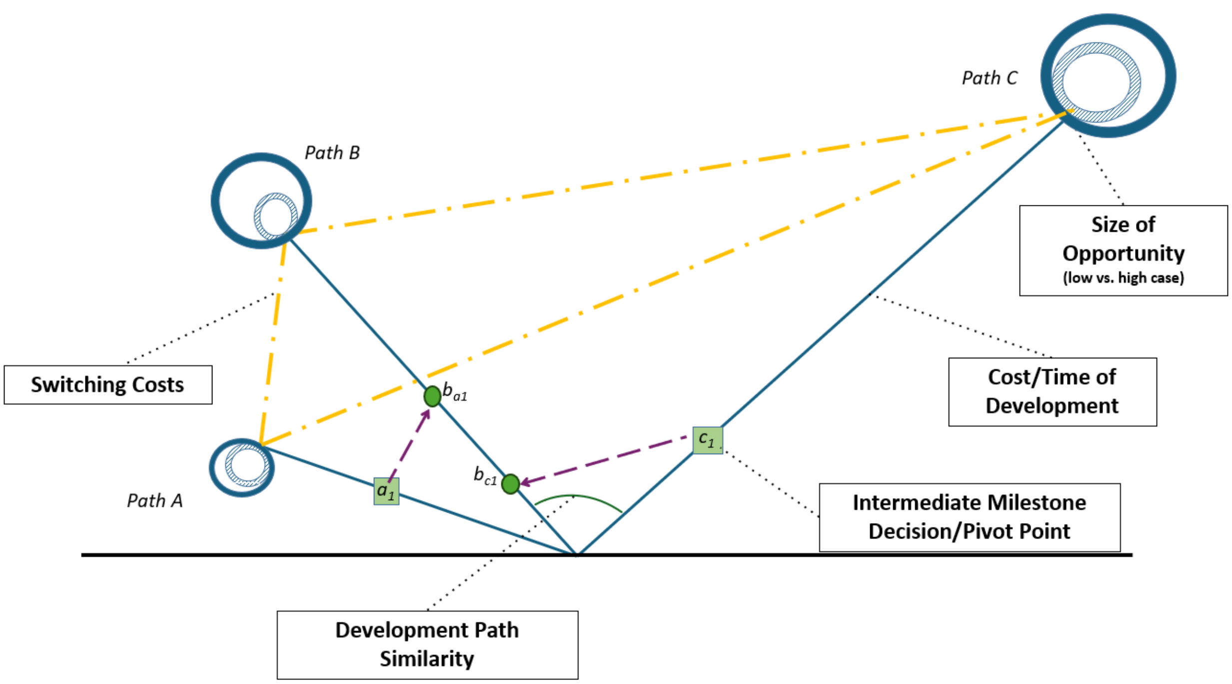

pivot cost would be the length of the line from the intermediate milestone to the point that minimizes the sum of switching costs and cost of development to the next milestone (or completion) on the new line—i.e., move towards point b_{a1} on path B, where the location of b_{a1} is chosen in order to minimize the term $c_{a_1 \to b_{a1}} + (c_{b_1 - b_{a1}})$.

Learning and updating across adjacent paths: By moving down any given path, the venture learns valuable information about the viability of that specific path and updates on adjacent paths. If development continues all the way to the end of the path, the firm will resolve all uncertainty about the technology, development costs and commercial potential for that given path. The degree of adjacency between related venture opportunities captures the level of *spillover learning* across paths.

For both success and failure on the initial path (or milestone), the venture can update the probability of success for adjacent paths based on their development overlap. That overlap is the pair's path success correlation— denoted as the probability that path B is viable given success on path B ($S_{B|A}$). That correlation is determined by two components: the angle between the two paths ($\theta_{A,B}$), and the relative lengths of the two paths (cost/complexity differences: c_{A-B}).[9]

To calculate the updated probability of success on adjacent lines, one simply applies Bayes rule as follows:

$$P(B|A = success) = \frac{P(A|B) \cdot P(B)}{P(A)} = \frac{S_{(A|B)} \cdot P(B)}{P(A)} \tag{2}$$

$$P(B|A = failure) = \frac{(1 - P(A|B)) \cdot P(B)}{(1 - P(A))} = \frac{(1 - S_{(A|B)}) \cdot P(B)}{(1 - P(A))} \tag{3}$$

Thus, the expected value of any given path includes the the risk-adjusted value of that given path, plus the probability weigthed values of the expected *updated* value of adjacent paths following the initial path outcome. Conceptually, one can calculate the expected value of all plausible paths, including all intermediate milestones and switching pathways. In practice, the number of possible scenarios increases greatly as one adds paths and milestones, so the Bayesian entrepreneur is may limit the set of possible paths to the most likely, and limit

[9] In practice, an estimate of the percentage of project overlap direction can be converted into the angle of overlap, or vice versa. For example, 80% development overlap can be converted to a 36° degree vector angle because $(1 - (36°/180°) = 80\%)$. Even though $\theta_{A,B} = \theta_{B,A}$, $S_{A|B}$ and $S_{B|A}$ need not be equal, since the amount of learning can also be scaled by the relative complexity of the two paths—e.g., if $c_B - c_A > 0$, then any updating on p_B could be scaled by that difference, $S_{B|A} = \frac{c_A}{c_B} \times (1 - (\theta^{\circ}_{A,B}/180°))$.

intermediate decision points to the most significant milestones (rather than any opportunity to update).

Any venture that purports to develop a multi- or general purpose technology implicitly makes this optimization decision. But instead of articulating the sequencing tradeoffs of this development paths framework, they often determine or justify their choice of initial paths by using heuristics like total addressable/obtainable market size, targeting "beachheads" with high likelihood of success, or establishing "basecamps" with great learning opportunities. However, different heuristics can sometimes point to different development sequences.

For example, RightHand Robotics spun out of Professor Robert Howe's Harvard engineering lab with a gripper technology that had beat competition in in a DARPA competition for robotic hands picking up a variety of objects without pre-programming. The venture considered a wide range of potential commercial applications including produce harvesting, bomb disposal, meal kit preparation, warehouse automation and recycling (Eisenmann and Straaberg, 2013). RightHand could justify any one path by emphasizing a single selection heuristic or a simple weighted combination of factors.

Going straight for the largest addressable markets would point to produce harvesting or warehouse picking/packing, however both would require complex and capital intensive development to meet the high performance, computer vision, and mobility needs. Recycling might allow for a shorter path to market and more fault-tolerance, but the profit potential was potentially much smaller with only moderate learning applicable to other applications. Meal kit preparation was a growing market with a clear need for stationary robotics to pick and scoop a small set of repeat items, but required application-specific engineering (e.g., food safe grippers, operating in refrigerated rooms, identifying poor quality ingredients) with little overlap to their other target applications. The development paths framework incorporates the risk and profit opportunities from each opportunity with the learning across applications to highlight the optimal *initial path* and allow the entrepreneur to analyze how sensitive that choice is to any given within- or across-path variables.

In addition to choosing initial paths, the development paths frameworks also applies to evaluating new opportunities for existing ventures. Young ventures often begin by pursuing a single market application, but then learn about adjacent market opportunities or even consider new technologies to serve their existing target markets. Properly evaluating these new options requires assessing their learning and real option value, as well as their individual path expected values ($p_{newpath} \times V_{newpath} - c_{newpath}$). When new opportunities arise, the same approach as the initial path selection can be applied—visualizing the decision point as either a new origin or as an unexpected intermediate milestone.

The team from weather analytics company Tomorrow.io faced multiple such moments of reassessing their path. One occurred early on when they were focusing their commercial efforts on serving the airlines and live events (e.g., professional sports and concerts). The car maker Porsche offered them an opportunity to collaborate and use sensor data from their connected vehicles as new inputs into their weather models (Krieger et al., 2019). The decision to take on the pilot project required estimating implications that all feed into the path selection framework: 1) the time and cash costs of pursuing the pilot project (c_i) including opportunity costs to other current R&D efforts, 2) the probability that they win a bigger contract with Porsche based on the pilot project, 3) the value of the Porsche opportunity conditional on succeeding, 4) the potential value of adjacent automotive industry opportunities, conditional on a successful relationship with Porsche , and 5) the learning and spillover value of both success and failure on the Porsche pilot project on adjacent paths in automotive, shipping, airlines and military applications. The company did pursue the collaboration, only to once again reevaluate it's path choice again with the emergence of another new (set of paths) around developing its own satellite constellation (Krieger et al., 2022).

5 Conclusion

This chapter highlights two primary "modes" of entrepreneurial learning: determining viability and optimizing path direction. However, the two modes are not mutually exclusive. Even if the entrepreneur has enough conviction to start down a particular development path, the necessary activities will still generate new information that leads to updating about venture viability (for the team and invesors). Sophisticated entrepreneurs might select initial paths and experimental milestones that resolve questions of overall venture viability while also trying to adjudicate choices between adjacent market applications paths.

Whether committed to pursing the venture and trying to navigate the best path or still trying to learn whether the venture is viable, the type of experimentation and information gathering depends on the venture's greatest uncertainties. This approach in this chapter emphasizes three key expected values for each opportunity: probability of success, cost of development, and profitability (conditional on success). However, the framework also shows that ranking projects based on their individual expected values is not sufficient for charting the venture's learning direction. The optimal sequence of paths also depends on the strength of the priors (i.e., spread of the distribution) for each of these variables, the ability to (cost-effectively) update on those priors, and the spillover learning value for adjacent opportunities.

This approach not only helps entrepreneurs navigate the uncertainties inherent in venture creation but also aligns their strategies with the expectations of key stakeholders. The visualization of development paths and the incorporation of learning spillovers offer a robust method for making informed decisions about initial and subsequent venture trajectories. Ultimately, by moving beyond simplistic heuristics, this framework empowers entrepreneurs to systematically evaluate and pivot their strategies, thereby enhancing their ability to build successful and innovative ventures.

References

Arora, A., Fosfuri, A., and Rønde, T. (2024). The missing middle: Value capture in the market for startups. *Research Policy*, 53(3):104958.

Bresnahan, T. (2010). General purpose technologies. *Handbook of the Economics of Innovation*, 2:761–791.

Eisenmann, T. and Straaberg, S. (2013). Righthand robotics: Choosing the first market. *Harvard Business School Case Study*.

Ghosh, S. and Payton, C. (2016). Rest devices. *Harvard Business School Case Study*.

Krieger, J. L., Nagaraj, A., and Barnett, J. (2022). Tomorrow.io goes to space. *Harvard Business School Case Study*.

Krieger, J. L., Stanton, C., and Barnett, J. (2019). Forecasting climacell. *Harvard Business School Case Study*.

Lerner, J. and Nanda, R. (2020). Venture capital's role in financing innovation: What we know and how much we still need to learn. *Journal of Economic Perspectives*, 34(3):237–261.

Nanda, R. (2020). Financing innovation "tough tech". *The Global Innovation Index*, pages 113–119.

Quelch, J. A. and Rodriguez, M. (2014). Vaxess technologies, inc. *Harvard Business School Case Study*.