

Lecture Text

Jan W. Rivkin

Where Do Great Strategies Come From?

(edited for clarity)

Introduction

I want to use the Ryanair case as a key input to a broader discussion of a question that I think is of fundamental interest. The question is, where do great strategies come from? And I think a natural candidate answer is, great strategies come from great strategic planning processes.

And what I'm going to argue is first, that there are pathologies in the traditional planning process that actually make it fairly unlikely that the process is going to produce great strategies. And I'm going to advocate an alternative approach that I think captures fairly cleanly the gist of what many great strategists actually do when they go about their business and try to lay out explicitly what that process looks like.

Now before I can do any of that, I have to be very explicit about what I think constitutes a great strategy. So that's the plan. So I'm going to say what I think the tests of an effective competitive strategy are, and then use that to offer a critique, I hope a respectful critique, of the traditional strategic planning process. And then I'll offer an alternative approach, with a special emphasis on where creative options come from, on how you can be more creative and come up with new strategies.

What is Strategy?

I need to be clear about what constitutes a great strategy. Before I can even go into that, I have to come clean and say what I mean by a strategy, because there are a gazillion definitions out there.

The definition I've put up here is one that my colleagues and I have found to be fairly effective. It's straightforward. It says, "A strategy is an integrated set of choices that position a firm, in an industry, to earn superior returns over the long run." And this, to me, is useful because it really captures the three things that, in my mind, distinguish strategy in the business.

What's different about strategy? And the first thing is this comment about "in an industry." The strategists are the people who are responsible for seeing the overall external environment, in which a firm competes, and seeing it broadly and dealing with it.

The second thing is "an integrated set of choices," that strategy is about looking across the parts of the company holistically—everything internal to the firm—and trying to understand how the pieces fit together to form some greater whole.

And the final thing is this "over a long run" bit; that the strategists are uniquely responsible in an organization for looking out for the long-term health of an organization in the face of competitive dynamics. All right? So external, internal, and dynamics.

And corresponding to those three things that make strategy distinct are the tests of a good strategy that I'm going to advocate for. And I'll argue, for instance, that the Ryanair story can be seen through this lens pretty effectively.

Three Consistency Tests

So what do I mean by these three consistency tests? The first thing, external consistency: Does the strategy tap the opportunities and neutralize the threats offered by the outside world in a unique way? The converse of this would be either ignoring what's going on in the outside world, or simply being "me, too," not being unique.

Second, internal consistency: Do the parts fit together to form a whole greater than the sum of the parts? The converse of this would be having pieces of a strategy that don't quite make sense together. Having a low-price offering and adding on extra amenities at the end.

Dynamic consistency: Does the strategy call on a company to do today what is required to succeed tomorrow? Does it, for instance, correctly anticipate what competitors will do and deal with that before they cause problems for the firm?

Internal consistency

Now of these three, I want to spend a little time on internal consistency, because this is the one that, frankly, has gotten the least play in old strategy thinking. You all had, for instance, the "five forces" imposed on you when you were here. That's about looking at the external world. We've got tools for thinking about competitor analysis. But the internal consistency part, I think, is relatively new.

The first part of the internal consistency element that I want to focus on is the fact that effective strategies are typically embodied in a large number of choices that the company makes. And this goes back to what Connie was saying earlier, that, really, the entire value chain of the company is aligned at Ryanair to contribute. There are all these ancillary streams of revenue, and so forth. This runs against, actually, a lot of thinking in the management literature in the 1990s. There was a lot of focus on the core competence of the corporation, which typically was an exercise in trying to find out the one thing in a company, the "silver bullet," that really made the company successful and, I think, by and large, was not a successful venture—the whole idea of looking for a single silver bullet.

In fact, back about a decade ago, Mike Porter and I tried to go through a set of about fifty highly successful companies in search of a core competence, and frankly, in forty-nine out of the fifty, we couldn't find any single thing that we could attribute success to. Instead, we saw things throughout the value chain of the company that contributed. And so the value chain is simply a tool for cataloguing all the things that a company has to do to design, produce, deliver, sell, market, and service a product. And here I've got the value chain for Ryanair, and if you're in the back of the room, or even in the middle of the room, you don't have a prayer of actually reading this slide.

And that's exactly the point, in fact. There's so much going on here, so many things contributing to what makes Ryanair successful, that it's difficult to pin it on any one thing. Individual pieces here don't really necessarily make sense. Having no meals, for instance. Well, if you have short-haul, point-to-point flights, that may make sense. If you've got transatlantic flights and you offer no meals? Doesn't make a ton of sense, right? And so when you also try to understand the connections across the pieces here, you start to see all

sorts of reinforcement. The individual elements don't stand alone, but instead they make sense because of everything else the company is doing. And that's what raises the possibility of internal consistency or inconsistency. And sometimes you probably have seen big diagrams like this that kind of lay out the core activities, not in the value chain but in a more of a geometric way, and put the connections together. So for instance, the twenty-five-minute gate turns at Ryanair are achieved in part because you don't have to slow down the flight in order to load up meals. You don't have air-bridges that might break down—Jetways. You don't have seat assignments. You've got gate crews that are very familiar with your fleet, which is entirely 737s. Pete is glad: They're all from Boeing, and the next one hundred will be also, and a hundred after that, and a hundred after that. But that all allows the twenty-five-minute gate turns, right?

The effects of connections

Now, to talk about this further, I wanted to say something a little abstract about the effects of those connections, those pervasive connections across decisions. And it is a little bit abstract. And what I find is that for some people, the imagery I'm about to introduce makes a ton of sense and really clarifies things. For other people, it clouds matters. If you're in the latter category, don't worry. The clouds will clear in about three minutes, I promise.

So what are the effects of the connections across decisions? I'm going to argue that they make the problem that a management team faces in trying to find a good strategy that is internally, externally, and dynamically consistent extremely difficult. It makes this a non-trivial challenge to discover a great strategy.

So what I have done here is imagine a company that has to make two choices, and those two choices produce some performance outcome. And in the first situation, there aren't connections across the decisions in the following sense: that the marginal effect of Decision 2 on performance is not affected by what you decide about Decision 1—a cross-section, with respect to Decision 2. The same regardless of what you do in Decision 1. And you've got a nice, smooth landscape with a single peak.

Over here if you change Decision 1, the marginal effects of Decision 2 vary dramatically, and you wind up with a very rugged surface. Now, in your job as manager, you're in a much higher dimensional space than this. You've got a ton of decisions to make, and they affect one another, so you've got a very, very difficult problem to solve.

And in fact, many of the things that I'll talk about can be seen right in the slide. Over here you've got kind of one right way of doing business. It's a very clear, one-peak business. Over here, there are lots of different ways of doing business. There are different internally consistent ways of doing business. There's a Ryanair way, and a British Airways way. There is a UPS way, and FedEx way. There's a Wal-Mart way, and a Kohl's way, and a Nordstrom's way. And everything I know about business says reality is actually over here.

Over here, frankly, there's no need to worry about active coordination across the parts of the company. If you just let each part do its own thing, you come up with the one best way of doing business. Reality's over here. In fact, there's a need for active coordination. Over here, how do you actually start? Where you start and how you search is not really relevant. You're going to find your way just by tweaking your way uphill to a good combination of choices. Over here, where you start and how you search is dramatically important to where you wind up.

And the upshot is that if you've got lots of connections across decisions, finding your way to this kind of thing is very difficult. Actually discovering a strategy that works well is difficult because, how do you get to this? Is it a process of design? Is it a process of discovery? You've got that kind of payoff surface in which you're searching.

Ryanair's performance viewed through consistency tests

Now heading back to those tests of internal, external, and dynamic consistency, to my mind a lot of the Ryanair story can be seen right through these three tests. So we go back to Ryanair in the old days. Think about the external consistency. They're going head-to-head with a "me, too" strategy against an entrenched competitor without a clear source of advantage, no unique position. It really exploits the opportunities in the outside world. Moreover, they're radically undermining the structure of the industry by giving away—basically by doing the same service but offering a lower price. Doing nothing more than that.

We talked about the internal inconsistencies in this. You've got what looks like a higher-end product but you're lower price, and we don't really know what we're doing. We're trying to be everything to everybody by the end of the day.

And then, frankly, the real killer to me is the failure to anticipate what Aer Lingus and British Airways are likely to do as a response. I think that they frankly did not anticipate what was going to go on. They provoked it and they certainly did not defend against it.

In contrast, since 1991 and the turnaround, you've got a dramatic difference. You've got a company that is really taking advantage in the changes in the business: the deregulation, the increasing mobility of Europeans. You also are finding a company that has found a way to overcome the threat posed by the full-service airlines. For instance, they've got such a low cost structure—negative two pounds Irish marginal cost. That's something that the full-fare airlines just cannot threaten. Everything about this fits, as we talked about. A very different way of doing business. And, as we mentioned before also, they know what they're doing and, to Jeff's point, they know what they're not doing.

Then finally, there's the dynamic consistency. Whereas in the past, they kind of sat back and watched and waited for the attack, here they recognize that in the low-cost airline business, there's only room for one low-cost player on each route. And as a consequence, you'd better be the first one on that route. And part of the dramatic acceleration in their expansion, which has occurred, basically starting in 1999—the couple hundred planes ordered from Boeing, for instance—is an effort to get out there and claim those routes first.

How Do I Get There?

What I want to move to is the question of, how do you actually get this? Because I find that with most executive teams, they're fine. They look at this. They say, "This is great. I understand how this works. I understand the test. I understand this is beautiful. My God, what do I do? How do I find one of those?" And, "I can't go online and order one. What can I do?"

Now, a natural candidate for this is, maybe the strategic planning process can yield this. Another natural candidate is, maybe it's a matter of luck. And I would be remiss if I didn't say that, in fact, whether I've done my research on where strategies actually come from, there is a large element of luck involved here. And there are actually studies showing what

portion of companies actually use formal strategic planning or formal business plans when they start up, among successful entrepreneurial companies, and it's relatively small.

But I want to hold that possibility—just luck—aside, because I think I can argue by the end that, in fact, there is a structure to what effective strategists do when thinking through problems, even if it doesn't take the form of formal planning.

Traditional Strategic Planning Process

Now I want to offer a critique of the role of traditional strategic planning processes in discovering things like this. And the critique I'm going to offer is, frankly, a caricature. And I desperately want you to push back on me on this, and share your experiences, whether this gels with them or not.

It is a caricature that I offer with some apologies, but not completely apologetically, because I think it captures what many of my colleagues and I have seen, and in many very prominent firms. Also, frankly, my colleagues who run these processes have encouraged me to publicize this caricature. Here's how it looks. Traditional process.

Driven by calendar

First of all, this process is typically driven by the calendar. So, we are engaged in strategic planning, not because we currently face a strategic issue, but because it's April, or May.

Focus on one preferred option

And when we do this, the first thing we do is we very quickly focus on a preferred option. That preferred option often comes from the sum of individual parts. The individual functions, or product lines, or regions submit their plans. We kind of staple them together, and then we make some adjustments. The CFO goes through it and says, "Uh-uh, 5 to 10 percent low. Everything's got to go up." Right? I can see from some of the smiles that, actually, you've been there and done that. There are occasionally some alternatives given, but basically they're nonviable strawmen, and there's a preferred option.

Analysis

Now early on, a lot of the work of doing this gets delegated to staff members. So the senior team kind of goes to the staff and kind of whispers the right answer in the ears of the staff. The staff then goes out and conducts an exhaustive search for confirming data, for the right answer. The world is awash in data. If you search hard enough, you will find confirming data. And sure enough, it comes back with large binders of analysis that basically confirm what we thought was the right thing.

Preferred option ratified

There is then a kind of a coronation ceremony where the senior management ratifies the preferred option. It's often very close to what the company was already doing, and a lot of the analyses that went into this don't actually get looked at; they get filed away.

Now have you been on the receiving end, or the giving end, of those analyses? Both. This sounds familiar to you, it sounds like.

Viewed through consistency tests

Okay, now I just want to view the output of this through the lenses of those tests that I mentioned before. First of all, external consistency. This is driven by the calendar. It is

relatively internally focused; there are very few external inputs given to this process. And therefore, changes in the world and differences in the world don't tend to be reflected here. Moreover because it's basically the sum of the functional plans, what happens is relatively incremental change.

There's never a step in this process where someone can say, "You know what? What if we changed everything about our airline—how we price, what services we offer, where we fly, how we train people, how we compensate them, how we promote them? If we changed everything, maybe there's a better way of doing business." There's never an opportunity for that here. In addition, because you're basically getting the sum of the functional plans, internal inconsistencies tend to creep in.

And I've got to say, a somewhat dubious privilege of the job is the opportunity to read a lot of strategic plans. And they are good, bad, and ugly. But one thing you note with dramatic frequency is internal inconsistencies, where a company is—you know, the marketing folks are promoting the new product just as the production folks are cutting back on capacity. Or the product line is being expanded just as the sales force training gets cut back. Lots of internal inconsistencies.

In addition, there's a real problem with dealing with change. Typically, the changes are suppressed. And because you're summing up across the functional plans, there's a tendency to try to look at a world, to assume a world, that looks fairly close to what today's world looks like.

Is the critique accurate?

And that's a fairly harsh critique, I understand. I just want to do a quick "gut check" with you and to get a sense from you of how accurate this looks, based on your own experiences. I want you to think about the activity systems that you saw; you know, the Ryanair thing. And think about the strategic planning process in the companies that you've worked with and worked for. Are they capable of producing things like what we saw? Or are they incapable? Any reflections? Is the critique actually on-point? And you gave a suggestion that it actually—that you've seen this. Can you open it up?

___: Yeah, I have—well, I guess one of the things that I've seen that ties in with Ryanair is that a lot of times what really helps is to have a strategic plan that fails. And in your mind, and you're heading along, and it seems like a great idea, and all of a sudden it fails. And before too long, the reason that you are where you are, which turns out to be a better place, is because you ran out of cash.

PROFESSOR RIVKIN: Now certainly one candidate for where great strategies come from is near-death experiences. I don't want you to all go back and say, "You know, at HBS they taught us we need to go bankrupt in order to succeed." But, the fact is, that is how companies suddenly wake up and see this. And the alternative approach that I'm going to push toward a little is one that I see in a number of places. And one of them that clearly comes up again and again is when companies reach a crisis. So I've systematically gone through and asked, "When crisis hits and you start throwing out this traditional process, what do companies really rely upon?" And the essence of that, I think, is distilled, to some extent, in this alternative approach that I'll turn to in a moment.

Alternative Strategic Planning Approach

And the alternative—I'm going to call this an alternative approach, but actually I think of it, frankly, as a complementary approach.

___: That's the scary part. We're going to have twice as much work now, huh?

PROFESSOR RIVKIN: I'm actually going to argue that the alternative is going to cut back on work, and here's why. There are a couple of things that are concerning about this beyond just these tests. One is, there's a huge amount of waste here. There are lots of analyses done that are not targeted, analyses that would never have a prayer of actually affecting what a company does, but are done, nonetheless.

And the other thing that concerns me about this is, to the extent that there is change that's made here, it's far divorced from implementation. You can get way down this process with a new idea without ever saying, "What does it actually mean on the ground?" And so I think the alternative actually addresses both of those issues.

Disclaimers

So I just want to put forward a couple of disclaimers about this. There are three, actually. The first is, I view this complementary to the existing process. The second note I want to say upfront, is that I'm going to talk about it as if it's in a large company, but actually I think that the essence of what's going on here, particularly the mindset involved in this alternative approach, can be applied at the smallest unit of analysis, in the individual level, and can be applied by very small companies. I've applied it with entrepreneurial companies a number of times. I'm on the board of a little student start-up. We used this when there were four people involved; we're using it now when there are sixty people involved, we're going to use it when there are a thousand. I actually think it does scale. I should also say, and this always worries me whenever people in academia make some claim of universality. I think actually it applies even more broadly. I have students now who tell me that they've made dinner plans, changed their dating strategies and their educational plans for their children, using this kind of approach, so you'll have to kind of reel me in occasionally about this.

The third thing I want to say before we plunge into this, is that I think the mindset changes here are as important as the process. And I'll come back to that in just a moment. Now what is the alternative I'm going to put forward? Basically it's kind of changing the sequencing and turning things on its head.

Driven by actual strategic issues

And it starts with a process that is driven by actual strategic issues. We are engaged in strategic planning, are making strategic choices, because we actually face a strategic issue. So it's because a Japanese competitor has entered our market, or because our market is undergoing deregulation, or something along those lines that we are actually engaged in this process.

Cross-functional management team

And early on in the process, a cross-functional management team comes together and creates a set of well-integrated options for the company. So instead of saying, "We're doing this," and quickly getting to a preferred option, the team actually says, "You know, we can

be a feeder through Shannon. We could be a no-frills airline. We could add business class, or we could get out of this business."

And to the earlier point, and I think Mark forced us to this, we have to actually go through and specify, what do we mean within no-frills? What is the full value chain around that? So it's really a detailed concrete alternative put forward early on.

Agreed-upon burden of proof

Then comes the hard part. You've got to agree on the burden of proof. What would have to be true? What would you have to believe to prefer one option over another?

Analysis and choice

And then you can conduct the analyses, and the team, then, having conducted the analyses, chooses the options this burden of proof has met.

Now, one important point about this is that I don't want to claim that this is an original invention. I actually believe that this is a distillation of what many effective strategists do. And certainly if you go to the strategy consulting firms, you see this kind of process laid out fairly cleanly. So McKinsey has hypothesis-driven consulting. And Bain's got the mantra of "answers first, analysis second." And Monitor Company has what they call "strategic choice structuring process." And BCG and Mercer all have their own labels, and they all capture some version of this.

Viewed through consistency tests

Let me say a word about why I think this actually ends up with a better outcome. First, with external consistency: It is driven not by the calendar, but by outside threats and opportunities. And this is, when you've got a big merger and acquisition, or something like that, you sidestep the original process, the traditional process, and go with this instead. In this process, there actually is an opportunity for systemic change. Right from the beginning you have to consider the alternatives to what we're doing. And what would it mean to change everything about what we're doing? And we have to flesh out, throughout the value chain, what are the effects of a given change? The fact that you have to have multiple options is crucial. If you start this and you only have one option, you're not making a decision. You're simply not. And finally, you're explicitly recognizing what the company is *not* doing. So there's a real focus on the outside world and the possibility of doing something unique here.

The second thing about this is because you start by laying out the alternatives in their full glory right from the beginning, you're building in internal consistency from the beginning. And finally, because you are working backwards from options instead of working forwards from preferred option, you're able to typically have a reaction to and deal with discontinuous change. You can accept the possibility of discontinuous change because you actually have the option of something discontinuously changing within the company, rather than assuming that the world doesn't change any faster than your company can, which is a typical and typically bad assumption.

Shifts in the mindset

And it's actually the shifts in the mindset that are the most important here, rather than the process, the linear process. I just want to emphasize what those are, in my opinion; there are a few of them.

The first is a shift from, what's the right answer? I think a lot of us, myself included, have an instinct to jump to the answer, right? We want answers; we want solutions. Here you have to start not with what should we do, but what might we do. And that means that you consider possibilities you would not otherwise.

The second is a shift from what do I believe about the world to what would you have to believe. And this is allowing the unusual alternatives to live a little bit, to breathe, to be built up before the inevitable pit bulls in organizations come in. And you know the pit bulls are there, right? The kind of, "Grrr, bad idea!" They're always there. And it's always easy to crush an idea before it gets developed. This allows them to develop it a little bit.

The third shift is from what is the right answer to what is the right question.

Creating Integrated Strategic Options

Now, the very last things I want to say, and then I want to stop broadcasting and get your thoughts, is this: How do you come up with creative options? Where do they actually come from? One guide is your gut, and I actually think the managerial gut comes in for a real beating at Harvard Business School, and inappropriately so. Often creative options come from the feeling that a manager has. The problem is, if you want to have multiple options out on the table initially, that's often very hard to get from one individual's feeling. So there are some techniques and some probing questions you can ask to try to generate additional options. And in looking across kind of the "option generation" I've seen in companies, in case writing, and in the MBA classroom, there are some very clear patterns that come up; some styles of thinking that come up again and again. I just want to highlight three.

Inside-out logic

The first is the kind of inside-out questions. I pose these as probing questions. The inside-out logic says, what is this company really good at? What does Ryanair have, really? And then ask, what would it mean to deepen or to broaden it? So this is Michael Eisner at Disney in the 1980s, saying, "You know what? Everyone says Disney is about family entertainment. Wrong. Disney is about managing creativity." And that's a much broader field. Now there's a danger. This is also Enron saying, "We're good at trading. And I don't care if we're trading natural gas, or broadband capacity, or weather derivatives. We're good at trading." Right? So we've got to be a little cautious about that.

Outside-in logic

I tend to be more partial to the outside-in logic, looking at the external environment and trying to find the white spaces. Where are there underserved customer needs? Saturn recognized that customers of automobiles actually want a pleasant buying experience. And that's not served by conventional auto dealers.

Often the greatest strategies are finding overserved customers. Ryanair is a great example of this. Not everyone needs a meal on a short-haul flight. And often, this involves looking for space. And I think in an age of consolidation, this is actually a critical one. People underappreciate how much space gets opened up by consolidation. That when the beer business consolidates, there's room for microbreweries. When the newspaper business consolidates, there's room for smaller community newspapers.

Far-outside-in logic

The final style of logic that I see occasionally is the far-outside-in approach. And the way this works is by looking beyond your industry, and whenever you see a competitor that you admire, ask yourself, what would it mean to be the Ryanair of my business? What would it mean to be the Wal-Mart or the Intel of my business? Now there are dangers here. I find that every time I teach the Dell case, for instance, some MBA student wanders in to my office and he or she says something like, "Ah, I love the Dell case! I want to be the Dell of the pizza-delivery business!" And I say, "Oh, that's fascinating. Sounds great. Very exciting." And then I say, "What do you mean? If you mean you're going to take orders directly from customers, make to order, and then ship directly to the customer, we kind of have that in the pizza business. So what does that mean?"

Parting Shots

Well, let me just get to a handful of parting shots. So here's what I want to say. And the first thing I want to emphasize is that, to my mind, coming up with a creative new option is just the most creative act in business. There is nothing higher than coming up with the direct model of Michael Dell, for instance. It's hard to imagine. That's being able to piece together a business in an entirely different way. And as a creative act, it is an art, not a science. And I actually think it will always be, in large part, an art, if only because anything that you can actually make scientific and codify, also diffuses and ceases to be a source of advantage. And so I think it will remain an art, but I don't think it's a black art. There are actually steps in the process that one can capture, and one can distill, that improve the odds of the outcome being effective and meeting those tests of internal, external, and dynamic consistency.

And to my mind, a process that is responsive to the outside world, that actually builds in internal consistency from the start and reverse-engineers from the options the analyses you need to do, is far more likely to produce those kinds of outcomes than a process that kind of adds together the individual parts and pays less attention to the outside world, and then hopes, at the end of the process, that something creative pops out.

The final thing that I just want to emphasize is the personality traits and the habits of the management teams involved here. I've shortchanged that because of the short period of time we have together here. But I actually think, and this goes to an earlier point, that there are a bunch of contradictions that a team has to manage here. The ability to say, how might the world be, and to be able to recognize how the world actually is. To be able to be creative, and yet be cold, hard, and analytical. It's very difficult to manage. And the second trait that I find extraordinary in the individuals that seem to pull this off is the courage to do something different, to resist the gravitational pull of the "me, too" strategy and do exactly what everyone else is doing.

And I can't claim that the process alone is going to generate that, but I think the process can actually start to pull out these traits. These are traits we can aspire to. And anything in the processes that actually pulls out these kinds of traits, I think is going to contribute to the creation of effective strategies.

So I want to wrap up on time. I want to thank you for being such a great group this afternoon. I'll hang around and talk to anyone who's interested. I will talk, you can tell, about these things until you're blue in the face. I'm a strategy junkie. And I really appreciate all the help and work we've done together. Take care.