

ELEVEN

The Safety Net

An Independent Devil's Advocate Review

“More often than not, failure in innovation is rooted in not having asked an important question, rather than having arrived at an incorrect answer,” Clay Christensen has observed.¹ Unfortunately, the internal process safeguards that we proposed in chapter 10 are no guarantee that the important questions will be asked about unfounded assumptions, unattainable forecasts, untreated deal fever, or any one of the many other individual and organizational tendencies that can lead to ill-conceived strategies. In the heat of the moment, even the most experienced executives and the strongest safeguards can fail.

Steve Hilbert, the CEO of Consecoco who acquired Green Tree Financial, is a testament to this hazard. Hilbert had successfully built Consecoco through dozens of acquisitions. He had a rigorous acquisition process and a crack team, both honed through almost two decades of deals. Yet, in Consecoco's biggest deal, Hilbert's experience failed him. His vision of taking Consecoco “to the next level,” to create a broader financial-services company—a financial-services Wal-Mart for middle America—led him to pay an exorbitant premium for a company that was clearly in trouble. He was willing to gamble more than \$6 billion based on his vision and just seven days of due diligence. Given that the investor and analyst skepticism was immediate and overwhelming, he surely faced internal skepticism as well. But that skepticism was muted. It is well nigh impossible for someone inside an organization to tell the CEO that his rousing vision is flawed and that he has fallen victim to deal fever.

To protect against the most dramatic strategic failures, we call for an additional safety net—a series of rigorous “second-chance” hearings by an independent devil’s advocate.

Whenever an organization is embarking on a high-stakes strategy, it should subject that strategy to an independent devil’s advocate review. Ideally, as we described in chapter 10, this is a formal part of the strategy development process. It is one of our “first, decide how to decide” process safeguards: it should be well understood by all involved before the strategy itself emerges. Just the anticipation of the review, we believe, will lead to better outcomes. In addition, strategy developers would benefit from the questions raised by the devil’s advocate throughout the process, rather than see them as attacks on their hard-earned positions. The CEO would also get feedback and challenges along the way, before making any sort of public commitment and risking embarrassment if he later backed away from the strategy. This independent devil’s advocate review would culminate in a final “last-chance” review, which should be done toward the end of the strategy-formulation process but before it’s too late, before the momentum is so great that almost nothing could stop the strategy. As one executive put it, “Once the lawyers enter the room, it’s too late to turn back.” In the event that, for whatever reason, the devil’s advocate review is not initiated during the planning process, we believe CEOs should still commission a formal review at the end of the strategy-setting process.

We find that a devil’s advocate review can build consensus among a management team and can convince a board that it doesn’t need to probe too deeply into a strategy; if a formal review from an independent viewpoint can’t find a problem with a strategy, then it’s hard for the board or executive team to object. Some CEOs have commissioned reviews even after a strategy is set to identify the problems that might surface and let everyone be better prepared if they do. We still believe the review should happen before the big bet is placed, but we won’t argue with success.

Whenever a review happens, and for whatever particular purpose, here is, perhaps, one of the most sensitive aspects of our recommendation: The independent devil’s advocate process and, especially, the result of its last-chance review, should be completely transparent to the board of directors. Based on our research and consulting experience, many CEOs are hesitant at adopting this suggestion. They suspect that such

transparency might disrupt their carefully managed board interactions. Our experience is that, properly orchestrated, such transparency raises the level of constructive discourse between management and board members.

The key is to establish the right guiding principles, which we will lay out later in this chapter. One ex-CEO, who retired a few years ago after a long run at the top of a business with more than \$20 billion in annual revenue, helped us sharpen the design of this safety net in this key aspect. While many CEOs are loath to do anything that would cede any authority, he said that he'd accept a formal devil's advocate review as long as we maintain a key distinction: that between governance and management. He said he'd resist any attempt to take over the management of the business from him. But he accepted that, in the interest of good corporate governance, the board had every right to understand and react to his decisions.

This is what we've been saying all along: that we're not trying to take decision-making authority from the CEO or anyone at another level of the business who has earned the right to decide about a strategy. But we thought it might be worth putting the issue in a CEO's terms: The devil's advocate review is about governance, not about usurping management authority.

As for how to organize a devil's advocate review: The review should be orchestrated by someone with credibility, objectivity, and some level of inoculation from organizational politics and pressures (acknowledging, of course, that no one is completely immune). This might be an independent director or some recently retired senior manager with no ax to grind. The reviewers should consist of, in large part, credible, objective outside reviewers with no stake in the outcome. The reviewers should be told to argue the "no" side, with license to raise any question whatsoever to ensure that the strategy is in the best long-term interests of the organization.

Skillful orchestration is an imperative. The closest analogy to the independent devil's advocate is probably the formal due diligence conducted in the course of acquisitions and joint ventures. In particular, our reviews conceptually overlap with "strategic due diligence," the stage that probes the logic of the transaction to ensure that it is aligned with the strategy of the acquirer.

Yet, by the time the strategic due diligence is reached, it is usually too

late to address the misaligned goals and flawed strategies that doom most acquisitions, irrespective of the quality of the acquisition target. This happens, in part, because by the time the due-diligence stage is reached, the corporate strategy motivating the transaction is pretty much assumed, and not an open topic of discussion. The due-diligence team is assembled to assess the quality of the target and to assign the right price to it, not to assess the motivating strategy or the target's fit with that strategy. Even if the motivating strategy were subject to inspection, due-diligence teams are usually not equipped to address it. Due diligence requires extensive investigation into every aspect of the target, including legal, financial, accounting, tax, operational, human resource, technology, and organizational issues, just to name a few. Team members are chosen for their expertise and doggedness at ferreting out the details and skeletons in one or a few of those areas. They rarely have the context, skills, or time to relate the findings back to the original strategy. As a result, flawed strategies, and even the lack of a strategy, are usually not flagged in the due-diligence process. For example, a study by Bain of 250 senior executives involved in major acquisitions found that more than 40 percent had no clear strategy of how the acquisition would boost profits and market value. Of those who claimed that they did, half admitted that they discovered that their approach was wrong.

The devil's advocate process that we propose occurs earlier, when it might still be possible to head off the issues that would otherwise not be addressed until strategic due diligence.

Formal due diligence is less common for strategic moves that do not involve outside transactions. For example, Green Tree's mortgage loan strategy, FLYi's decision to become an independent carrier, or Kodak's decision not to pursue digital photography in earnest would normally not have sparked a due-diligence review. Yet these sorts of decisions would clearly have benefited from an independent devil's advocate review.

The exact design of the independent devil's advocate review will depend on the context and particulars of the strategy in question. But, based on our research and experience, here are five general principles to follow:

1. Establish a clear and limited charter.
2. Organize for success.

3. Focus on the strategy, not the process that produced it.
4. Deliver questions, not answers.
5. Come to closure.

1. Establish a clear and limited charter

The license and charter of the independent devil's advocate review should be open and widely communicated. This allows the panel to conduct its review effectively, especially as it pursues sensitive information and explores uncomfortable questions. The charter also establishes a transparent social contract about the scope, conduct, and ultimate use of the findings. This keeps the review panel from ranging too far, addresses potential concerns about ulterior motives, and lessens the chances that uncomfortable findings will be buried. It also raises the likelihood that insiders can voice concerns they might otherwise keep to themselves if they thought that the review were merely a rubber-stamping exercise.

The scope of the review should be clearly limited to the strategy being reviewed. Depending on the strategy itself, the review might well open up a wide range of questions relevant to the context and history around the strategy. But the review team should resist straying further than it needs to in order to do the review. This constraint is intended to ease organizational concerns that the panel is some sort of parallel management team or that it was set up to evaluate people rather than the strategy. Ranging into other domains would undermine the effectiveness of the panel in carrying out its core function.

The conduct of the review should be restrained. The charter to "argue the 'no' side" is not a charter for indulgence. The devil's advocate is not an inquisition. Instead, explorations should be based on facts, not feelings, intuition, or emotion. The review should focus on pivotal aspects of the strategy on which success or failure will ultimately depend, not on minor defects. It is important that all understand that the purpose of the review is to uncover important questions about the strategy, not to suggest alternative approaches. This is so that those who formulated the strategy won't feel any more challenged than they already do and so that the panel will not enter the review with its own agenda. Panelists must

continually separate themselves from their own preconceptions of the “right” answer.

Panelists also need to stay in the real world, which is one of alternatives. Panelists can’t merely compare strategies against some model of perfection, because they’d dismiss all strategies as not measuring up. The review panel needs to focus on whether a strategy is the best one possible, even if it has warts and a boil or two, not whether the strategy is perfect. The panel also needs to remember that, as we’ve shown, doing nothing can be a bad strategy, too.

2. Organize for success

In general, the credibility and effectiveness of the devil’s advocate review hinges on the credibility and position of the review’s leader. Our experience is that this leader should be outside the management hierarchy directly affected by the proposed strategy. A corporate executive or executive from a different business unit might lead a business unit-level devil’s advocate review. An independent board member or some other seasoned outsider familiar with the organization, such as a retired executive, might lead a corporate-level devil’s advocate review. Ideally, the review leader should not have been involved in the strategy development and, without doubt, should have no strong interest in the outcome of the review. The review leader should also not be the chairman, the chief executive officer, or the lead outside director. Those leaders must grant license to the devil’s advocate review and orchestrate the proper use of its findings. Therefore, they should not be conducting the review themselves.

Beyond the leader, there are two other important components of the review team—staff support and the rest of the review panel. The staff team will help gather information, organize the review process, and explore issues raised by the panel. The staff team should be experienced in the range of issues surrounding the proposed strategy, including business, operational, financial, organizational, technology, and human resource issues. The team mirrors in some ways the skills that would be found on a due-diligence support team, though the level of resources should be much less, as the review staff team should be able to leverage prior work. (Our experience is that one warning sign that a strategy may

fail is when the devil's advocate staff team needs to do significant primary analysis—in other words, that the normal process hasn't already done the work.) It is actually a plus if the staff team members were involved in the planning and analysis of the strategy. They should not, however, have had leadership roles. They also need to understand and accept the important change in point of view that the review team members must adopt. Given these provisos, devil's advocate reviews are wonderful learning opportunities for young managers and analysts.

The review panel is probably the most difficult organizational task, and it is a key determinant to success or failure. One important aspect to keep in mind is that the review should provide a different point of view, not merely replicate existing expertise. The strategy development process probably already engaged a range of operational and strategic experts. The review team is not intended to be smarter versions of those experts. Instead, the team should collectively offer a divergent perspective, assembled to surface and test critical issues associated with the proposed strategy. But it would be a mistake to see the devil's advocate panel as merely a panel of experts. Our experience is that expertise is necessary but far from sufficient. More important is the ability of the participants to step back, listen, ask questions, be contemplative, and take a broad, systems-oriented approach.

To give you a flavor of the kind of person we think should be on a review panel, we'll mention two of our favorites. One is Vince Barabba. Vince has held senior strategy and market-research positions at General Motors, Kodak, and Xerox. He was twice director of the U.S. Census. He has written a number of books based on his experience. Thus, Vince always brings deep, practical experience in market research and strategic development. But Vince is not just an advocate for ideas and positions shaped by his immense experience. He has a natural tendency to search for the connected parts of any business problem and to ask questions that tease out the systems-level assumptions, issues, and consequences of any design. David Reed is another person we often call to participate in reviews. David was the chief scientist at Lotus and at various times a computer science professor at MIT and the famed MIT Media Lab. He has conducted advanced research at various world-class laboratories, including Interval Research and HP Labs and is credited as one of the original architects of the Internet. David clearly brings immense technological

expertise, a key underpinning to many strategies. (Mike Hammer, the prominent consultant and author, taught a legendarily difficult graduate class in computer science, and says he was stunned to find a freshman in his class. The next semester, he had that freshman teach the class. As you've guessed, that freshman was David Reed.) David also brings a propensity for looking at the broad consequences of designs, in particular how otherwise attractive solutions might or might not work at scale.

By contrast, we've found two types of personalities who tend to make poor panelists. Panelists with deep subject matter expertise fail miserably if they tend to come to quick conclusions and spend the rest of the review looking for confirming data and advocating their viewpoints to other panelists. Readers might recognize that this profile, unfortunately, fits most experts. Actually, the more expert they are, the greater the tendency to fit the profile. On the other hand, panelists who are overly empathetic also tend to add little value. Some panelists tend to quickly identify with the plight of management, accepting management's assumptions without a critical eye. This is easy to do, especially during the course of an extended review. The key is to maintain an outsider's perspective, and to continue to search for the key assumptions and ask the tough questions.

In general, all members of the devil's advocate team must understand and accept the charter of the review.

3. Focus on the strategy, not the process that produced it

Too often, quality-assurance reviews focus on process rather than content. Reviewers check that all required steps were taken, that proper analyses were conducted, and that all the right forms and documents were constructed. If all items on the checklist are crossed off, the reasoning goes, then the strategy is ready to go. While checklists are valuable, we don't think that you need to convene an independent devil's advocate review just to verify that appropriate steps were taken to develop the strategy. Nor would that be enough. At the end of the day, the strategy will stand on its own, or not. The review should focus on the strategy itself.

It won't surprise the reader that we recommend using many of the process safeguard tools that we previously recommended to assess the strategy. Here's a quick summary on how they might apply.

The first step is for the review team to develop a shared context about the business situation. Just like the strategy developers, the review team should understand “what is the problem” before venturing to assess “what we should do about it.” This requires starting at the beginning and following the issue history as far back as needed to understand the current context. It should understand the story of “how we got here,” including major events, trends, people, and organizations. It should understand this story from multiple perspectives, including current and past managers, rank-and-file employees, customers, competitors, reporters, analysts, and other industry experts.

Part of this context development is to stare into the abyss, to understand what might happen if the company does nothing, if the market takes certain turns, or if competitors make certain moves. What does the wisdom of the crowd, which might include capital markets or other kinds of prediction markets, say about the current situation?

With a shared context about the present, the review team should continue to resist diving into the proposed strategy. Instead, the next step is to understand the company’s strategic objectives and the whole range of possible options. This should just be a matter of calling up the early work of the strategy team. In any event, understanding what the company is trying to achieve and all the possible options adds more critical context. It allows the review team to ask whether the objectives are the right ones, and whether there might be other appropriate objectives. It also helps the review team assess if all the options were appropriately explored, or whether the ultimate choice was selected too quickly, with the elaboration of other options shortchanged.

With this larger context developed, the reviewers should immerse themselves into the particulars of the proposed strategy. The key documentation should, of course, come from the strategy team. We strongly suggest the reviewers start by asking for a detailed written description of the strategy. It should reject slides and spreadsheets as the principal documentation; they allow too much hand-waving and leave much to reader interpretation. As Bruce Nolop, Pitney Bowes’s CFO, observed, “I’ve been amazed at how many elements of a deal that seemed clear in PowerPoint can fall apart when they’re subjected to prose. In bullet-point format, the rationale for a deal might be summed up in a phrase, such as ‘cross-

selling.⁷ But a memorandum demands clarity about exactly who is cross-selling to whom—and how and why.”²²

Reviewers should also refrain from executive presentations on the strategy until they have had time to read and absorb the strategy documents. This avoids the problem of form over content, with the review team being swayed by the power of the presenter’s personality.

Once immersed in the strategy, the reviewers should seek to understand and examine its foundations. Surfacing and testing critical assumptions underlying the proposed strategy is one of the most important tasks of the review team. One technique that we’ve found effective for this purpose is the Strategic Assumptions Surfacing and Testing (SAST) process designed by Ian Mitroff, Richard Mason, and Jim Emshoff (and taught to us by Vince Barabba). SAST is typically used at the beginning of a strategy development process. We’ve adapted it for our devil’s advocate review, which happens nearer to the end of the planning process.

In our version, we organize divergent groups to examine and debate the assumptions that underlie the proposed strategy. Each group is chosen so there is commonality within the group but so members represent divergent experiences and perspectives relative to the other groups. There might be three groups, for example. One group could consist of key developers of the proposed strategy. Another group could be the review panel, which would represent an informed but objective group. A third group might be high-performing middle managers drawn from diverse parts of the organization.

Each group works independently through a facilitated process to identify key assumptions of the strategy. What must hold for it to succeed? Do we know enough to know whether these assumptions will actually hold if we proceed with the strategy? The result is a prioritized list of several kinds of assumptions: assumptions that are, in fact, strategic premises, and we are certain they are true; assumptions that are strategic premises, but we are uncertain whether they are true, and thus they require monitoring and research; other assumptions that might be explored if time and resources allow; and irrelevant assumptions.

Next, the groups come together to compare notes and debate differences. From this debate comes a prioritized list of the most critical issues

and a collective sense of which ones are well understood and which ones need more exploration.

We call these issues the pivot points of the strategy. They might involve key forecasts, risk assessments, possible scenarios, assessments of customer behavior, probabilities, and so on. In some cases, work will be needed to verify them. In other cases, it might be easier to try to disprove a key pivot point. In either case, some of the techniques that we described earlier might be deployed. Prediction markets could be set up to test key forecasts. Management bets could be used to crystallize odds and time predictions.

Another technique for making critical assumptions explicit is the “reverse income statement,” introduced by Rita Gunther McGrath and Ian MacMillan. The reverse income statement is part of McGrath and MacMillan’s “discovery-driven planning,” a systematic method for planning that all companies should consider. In the context of the independent devil’s advocate review, we use the reverse income statement to organize and verify the basic economics of the proposed strategy.

Ordinarily, a group determining a strategy starts with some knowledge and some assumptions about the world, then sees if it can construct a plan that would surpass some minimally acceptable level of profit. The McGrath-MacMillan approach does the reverse. It starts with an income statement that states what the minimally acceptable profit is, what revenue would need to be, and what major components of allowable cost are. The approach then asks, what assumptions must be true for the numbers to hold true? This approach can focus attention on the right issue: the assumptions that go into a strategy. Without this sort of approach, there can be too much attention on the numbers in a proposed business plan—but anybody can massage a business plan to make the numbers look good enough to pass a standard corporate screen; the question is whether the numbers are believable.³

Another important technique for examining the proposed strategy is history. Building on the earlier situational analysis using issue history, the review team should do its own investigation of history that fits. Assuming the strategy team did such an analysis, the devil’s advocate review can build on that. The key is to look at relevant historical successes and failures to understand the potential pitfalls of the proposed strategy.

An important aspect of the strategy to consider revolves around this

question: How might the future confound our plans? No battle plan survives first contact with the enemy. So a discussion needs to occur about two aspects of the future. The first aspect is tactical. What new evidence might change important current assumptions? And, as a result, what alarm systems should be designed around issues that might drift over time, both to monitor changes and to facilitate adaptation? The second aspect is more hypothetical. Can the review panel envision future scenarios that would thwart the proposed strategy? In this regard, their lateral perspectives will be particularly useful.

Another key element to explore is strategic fit. Is the proposed strategy within the operational reach of the organization? Because strategy is not organizationally neutral, the review team should explore the unvarnished truth about the fitness of the organization to implement the proposed strategy. What are the key strengths upon which this strategy will build, and what are the barriers that need to be removed to implement this strategy?

Over the course of an extended review, the review panel might have to deploy these tools multiple times, or to use different techniques to get at the same issues. It might well have to raise the same questions over and over again, to make sure that the team has not lost sight of them as the strategy, and related analysis, evolves.

4. Deliver questions, not answers

Perhaps the hardest aspect of the devil's advocate review for many to accept is that the goal is to deliver questions, not answers. Many corporate cultures espouse the attitude that one should not point to problems without also having answers. In general, planning processes try to produce the best answers. The devil's advocate review will focus on generating the most illuminating questions. This helps to ensure that members of the review team do not push their own agenda or preconceived answers but, instead, focus on examining the strategy as presented. In this vein, the review team might well leave the patient chopped up on the table, if that is warranted. It shouldn't try to sew it back up or even make it presentable. While this might sound harsh, it is important to head off any attempt to replicate the strategy process to provide "better" answers.

That's just not possible in the short window of the review. At the same time, good questions should lead to better answers.

We generally find that most of the value of the review comes from the discussions that it sparks and the questions that it raises. Still, it is important to have the discipline to develop a final report that summarizes and synthesizes the team's discussions and findings, including a prioritized list of critical issues and questions.

5. Come to closure

More than once we've seen leaders exercise the equivalent of a "pocket veto" of a review team's findings by thanking the team for its findings and disassembling the panel without formally addressing the findings. At some point, of course, management needs to make the final call and, as we've discussed, choose from the finite alternatives in front of it. That said, the work of the review team should not end with the delivery of the questions. At the very least, some type of formal response to the team's questions should be built into the process. Ideally, the review team should work with management and the board of directors until a final decision is made on the strategy.

After all, it does no good to know a strategy is bad and not be able to head it off. Covisint learned this the hard way. In the late 1990s and early 2000s, our friends at Diamond were doing some work at Covisint, a joint venture between major automotive companies that was expected to become the electronic exchange to end all exchanges. GM, Ford, and Chrysler, along with Nissan and Renault, had hoped to funnel tens of billions of dollars' worth of component and supply business to Covisint. Covisint would post the companies' needs electronically, so suppliers could bid on the work. The car companies hoped the increased competition would drive down the prices they paid. In addition, the companies hoped that Covisint, which they owned jointly, would be able to collect fees for acting as the middleman. When some of our then-partners at Diamond conducted a version of a devil's advocate review, however, they found all sorts of problems. For instance, governance was lacking, so the jockeying for position among the car companies and their bankers, consultants, and technology vendors made even routine decisions slow and painful; often, those decisions were

formally revisited several times. This fact raised all sorts of questions about whether Covisint could work as efficiently as the strategy required. In addition, there was no historical precedent for the sorts of fees Covisint hoped to collect, so revenue projections and the overall business model were dubious. But the Diamond team had no license to conduct its review, as it had been engaged in only one aspect of the business, reviewing the technology platform. Diamond could only deliver its findings to the Covisint management, which was a part of the problem. Diamond couldn't raise the issues with the parent companies. In the end, Covisint took no action. Diamond resigned, walking away from significant fees because it felt that the problems jeopardized the very success of the venture. The problems never did get fixed. Covisint did get launched and now offers some reasonable capabilities, but, when it was sold in two pieces in 2004, it was generating just \$60 million in annual revenue—a tiny fraction of what some had projected. Covisint wasn't yet profitable.

Contrast the Covisint experience with what happened at a major insurance company where it was possible to work with senior management on the results of a devil's advocate review:

A large consulting firm had helped the insurance company formulate a strategy designed to have a division double its revenue to almost \$6 billion over five years, beginning in 2006. Our friends at Diamond reviewed the strategy and found numerous assumptions that were, at best, dubious. For instance, the strategy required that the division win 50 percent of all new business that came to market over the five years, even though the company's best-performing products had never managed that high a share. In fact, a historical analysis of the competition revealed that no offering from any company had ever achieved the growth targets that the strategy required in the unit's target markets.

The CEO of the insurance company wasn't interested in having his assumptions challenged. He was ready to move on the strategy. But senior executives at the division supported the review because they didn't want to sign off on growth targets that couldn't be achieved. They eventually won a hearing with the CEO, who scaled back his expectations for the strategy. Rather than assume 17 percent growth annually, the company settled on a stretch goal of 13 percent. The CEO set up a "White Space" project to look for new ideas that would provide the additional \$1 billion

of annual revenue by 2010 that he was acknowledging would not come from the original strategy.

As it happens, one major piece of the division's revenue grew 12.5 percent in 2006 and 5.5 percent in 2007. The other major piece saw revenue rise 5.6 percent over the two years. If the original strategy had been implemented unquestioned, the division would have overinvested in search of growth and would have crushed earnings. Because targets were scaled back to more realistic levels, earnings rose along with revenue: Net income increased 12 percent in 2006 and 5.7 percent in 2007.

Even though the CEO didn't initially like the idea of a review, his willingness to listen made him come out looking like a hero, and everyone else involved was happy, too.