Spotlight

Deciding How to Decide

A tool kit for executives making high-risk strategic bets
by Hugh Courtney, Dan Lovallo, and Carmina Clarke

Senior managers are paid to make tough decisions. Much rides on the outcome of those decisions, and executives are judged—quite rightly—on their overall success rate. It's impossible to eliminate risk from strategic decision making, of course. But we believe that it is possible for executives—and companies—to significantly improve their chances of success by making one straightforward (albeit not simple) change: expanding their tool kit of decision support tools and understanding which tools work best for which decisions.

Most companies overrely on basic tools like discounted cash flow analysis or very simple quantitative scenario testing, even when they're facing highly complex, uncertain contexts. We see this constantly in our consulting and executive education work, and research bears out our impressions. Don't misunderstand. The conventional tools we all learned in business school are terrific when you're working in a stable environment, with a business model you understand and access to sound information. They're far less useful if you're on unfamiliar terrain—if you're in a fast-changing industry, launching a new kind of product, or shifting to a new business model. That's because conventional tools assume that decision makers have access to remarkably complete and reliable information. Yet every business leader we have worked with over the past 20 years acknowledges that more and more decisions involve judgments that must be made with incomplete and uncertain information.

The problem managers face is not a lack of appropriate tools. A wide variety of tools—including case-based decision analysis, qualitative scenario analysis, and information markets—can be used for decisions made under high degrees of uncertainty. But the sheer variety can be overwhelming without clear guidance about when to use one tool or combination of tools over another. Absent such guidance, decision makers will continue to rely solely on the tools they know best in an honest but misguided attempt to impose logic and structure on their make-or-break decisions.

In the first half of this article, we describe a model for matching the decision-making tool to the decision at hand, on the basis of three factors: how well you understand the variables that will determine success, how well you can predict the range of possible outcomes, and how centralized the relevant information is. We make a strong case for increased use of case-based decision analysis (which relies on multiple analogies) and qualitative scenario analysis under conditions of uncertainty.

Inevitably, the model we propose simplifies a very complicated reality in order to uncover some important truths. (That's what models do.) In the second half of the article, we explore some of the most common complications: Most executives underestimate the uncertainty they face; organizational protocols can hinder decision making; and managers have little understanding of when it's ideal to use several different tools to analyze a decision, or when it makes sense to delay a decision until they can frame it better.

Developing a Decision Profile

As you ponder which tools are appropriate for a given context, you need to ask yourself two fundamental questions:

Do I know what it will take to succeed? You need to know whether you have a causal model—that is, a strong understanding of what critical success factors and economic conditions, in what combination, will lead to a successful outcome. Companies that repeatedly make similar decisions often have strong causal models. Consider a retailer that has launched outlets for years in one country, or one that has made many small acquisitions of adjacent competitors.

One simple test of the strength of your causal model is whether you can specify with confidence a set of "if-then" statements about the decision. ("If our proposed new process technology lowers costs by X% and we are able to achieve Y% market share by passing those savings on to our customers, then we should invest in this technology.") You should also be able to specify a financial model into which you can plug different assumptions (such as how much the technology lowers costs and how much market share you are able to capture).

For the vast majority of strategic decisions, executives can't specify a clear causal model. Some managers have a reasonably good idea of the critical success factors that matter, but not a complete picture—this would generally be true of a company developing a new product, for example. Others don't even know how to frame the decision—for instance, a company being disrupted by a new technology wielded by a firm outside its industry.

Idea in Brief

THE PROBLEM

In highly complex, uncertain contexts, traditional decision-support tools such as discounted cash flow analysis are worse than useless. Nonetheless, executives use them all the time.

THE SOLUTION

Good analytic tools exist for all strategic contexts. If you use the right tool, the odds of making a good decision go way up. Here, the authors present a model for matching the decision-making tool to the decision at hand, on the basis of three factors: how well you understand the variables that will determine success, how well you can predict the range of possible outcomes, and how centralized the relevant information is.

THE EXAMPLE

The authors illustrate their framework using decisions that executives at McDonald's might need to make—from very clear-cut (choosing a site for a new store in the United States) to highly uncertain (changing the business in response to the obesity epidemic).

ASK YOURSELF:

Do you understand what combination of critical success factors will determine whether your decision leads to a successful outcome?

Do you know what metrics need to be met to ensure success?

Do you have a precise understanding of—almost a recipe for—how to achieve success?

Can I predict the range of possible outcomes? In choosing the right decision-support tools, you also need to know whether it's possible to predict an outcome, or a range of outcomes, that could result from the decision.

Sometimes it's possible to predict a single outcome with reasonable certainty, as when a company has made similar decisions many times before. More often, decision makers can identify a range of possible outcomes, both for specific success factors and for the decision as a whole. Often they can also predict the probability of those outcomes. However, under conditions of uncertainty, it's common for executives not to be able to specify the range of possible outcomes or their probability of occurring with any real precision (even in instances where they understand critical success factors and the model for success).

ASK YOURSELF:

Can you define the range of outcomes that could result from your decision, both in the aggregate and for each critical success factor?

Can you gauge the probability of each outcome?

Choosing the Right Tools: Five Contexts

As the exhibit "Diagnosing Your Decision" suggests, the answers to the questions above will point you to the best decision-support tools. (For brief definitions of each, see "Decision Support Tools: A Glossary.") In

some cases you'll need just one tool; in others you'll need a combination. Many of these tools will be familiar. However, the tool we advocate using most, case-based decision analysis, is not yet widely used, partly because the more formal, rigorous versions of it are relatively new and partly because executives typically underestimate the degree of uncertainty they face. (For more on case-based analysis, see the sidebar "Developing Rigorous Analogies: An Underutilized Tool.")

To illustrate, let's look at five scenarios that executives at McDonald's might face. (These are oversimplified for the sake of clarity.)

Situation 1: You understand your causal model and can predict the outcome of your decision with reasonable certainty. Suppose McDonald's executives must decide where to locate new U.S. restaurants. The company has or can get all the information it needs to be reasonably certain how a given location will perform. First, it knows the variables that matter for success: local demographics, traffic patterns, real estate availability and prices, and locations of competitive outlets. Second, it has or can obtain rich data sources on those variables. And third, it has well-calibrated restaurant revenue and cost models. Together that information constitutes a causal model. Decision makers can feed the information about traffic and other variables into standard discounted cash flow models to accurately predict (to a close-enough approximation) how the proposed location will perform and make a clear go/no-go decision.

TOOLS: *Conventional capital-budgeting tools such as discounted cash flow and expected rate of return*

Situation 2: You understand your causal model and can predict a range of possible outcomes, along with probabilities for those

outcomes. Imagine now that the McDonald's managers are deciding whether to introduce a new sandwich in the United States. They still have a reliable way to model costs and revenues; they have relevant data about demographics, foot traffic, and so forth. (In other words, they have a causal model.) But there's significant uncertainty about what the outcome of introducing the sandwich will be: They don't know what the demand will be, for example, nor do they know what impact the new product will have on sales of complementary products. However, they can predict a range of possible outcomes by using quantitative multiple scenario tools. Some preliminary market research in different regions of the country will most likely give them a range of outcomes, and perhaps even the probability of each. It might be possible to summarize this information in simple outcome trees that show the probability of different demand outcomes and the associated payoffs for McDonald's. The trees could be used to calculate the expected value, variance, and range of financial outcomes that McDonald's might face if it introduced the sandwich. Managers could then use standard decision-analysis techniques to make its final determination.

Alternatively, McDonald's could pilot the new sandwich in a limited number of regions. Such pilots provide useful information about the potential total market demand without incurring the risk of a full-scale rollout. Conducting a pilot is akin to investing in an "option" that provides information and gives you the right but not the obligation to roll out the product more extensively in the future. (This approach is still market research, but usually a more expensive form.) Real options analysis, which quantifies the benefits and costs of the pilot in light of market uncertainty, would be the appropriate decision-making tool in this case.

TOOLS: *Quantitative multiple scenario tools such as Monte Carlo simulations, decision analysis, and real options valuation. (These tools combine statistical methods with the conventional capital-budgeting models favored in Situation 1. Managers can simulate possible outcomes using known probabilities and discounted cash flow models and then use decision analysis tools to calculate expected values, ranges, and so on.)*

Situation 3: You understand your causal model but cannot predict outcomes. Let's now assume that McDonald's is entering an emerging market for the first time. Executives still understand

Decision Support Tools: A Glossary

CONVENTIONAL CAPITAL-BUDGETING TOOLS

These tools use the estimated incremental cash flows from potential investments to establish whether a project is worth being funded through the firm's capitalization structure. They include discounted cash flow, expected rate of return, and net present value models.

QUANTITATIVE MULTIPLE SCENARIO TOOLS

These tools analyze decisions by fully specifying possible outcomes and their probabilities. They use mathematical, statistical, and simulation methods to identify the risk/return properties of possible choices. The tools include:

• Monte Carlo methods, which use computational algorithms that rely on repeated random sampling to obtain numerical results

• decision analysis, which uses outcome scenarios and probabilities to identify the best decision to make given different decision-maker objectives

• real options, which applies the concept of financial option valuation to "real" situations and allows managers to quantify the costs and benefits of flexibility when making decisions under uncertainty

QUALITATIVE SCENARIO ANALYSIS

These tools inform decisions by developing a set of qualitative, representative scenarios of how the present may evolve into the future and identifying the likely consequences of the decision under consideration. Since these techniques don't assume a complete and fully specified set of possible outcomes, they are helpful to decision makers who face high levels of uncertainty about outcomes.

CASE-BASED DECISION ANALYSIS

These tools provide a systematic approach to aggregating and synthesizing information from analogous past experiences and examples. In general, analogies that are most similar to the decision at hand are given more weight in determining the best choice. (See the sidebar "Developing Rigorous Analogies.")

INFORMATION AGGREGATION TOOLS

These tools are used to collect information from diverse sources.

• Traditional approaches, including the Delphi method, gather information from a variety of expert sources, aggregate the responses, and generate a range of possible outcomes and their probabilities. Decision makers may then consult with the group again until a consensus is reached.

• Prediction or information markets are designed to gather "the wisdom of the crowd" by creating financial markets where investors can trade securities with payoffs linked to uncertain future outcomes (for example, the winner of an election or the release date of a new product).

• Incentivized estimate approaches involve surveying individuals with diverse information sources to estimate the outcomes of variables and then rewarding individuals with the most accurate estimates.

• Similarity-based forecasting methods involve asking individuals to rate how similar a decision or asset is to past decisions or assets. These similarity ratings are then aggregated across individuals using simple statistical procedures to generate forecasts (for revenues, completion times, or costs) depending on the goal. Since these methods rely on past decisions and outcomes, they are a form of case-based decision analysis as well.

the model that will drive store profitability. The cost and revenue drivers may well be the same, market to market. However, the company has much less information about outcomes, and predicting them using market research and statistical analysis would be difficult. Its products are relatively novel in this market, it will be facing unfamiliar competitors, it's less sure of supplier reliability, and it knows less about whom to hire and how to train them. In this situation, McDonald's can use qualitative scenario analysis to get a better sense of possible outcomes. It can build scenarios on the revenue side that cover a wide range of customer acceptance and competitor response profiles. On the supply side, scenarios might focus on uncertainties in the emerging market supply chain and regulatory structure that could cause wide variation in supplier costs and reliability. These scenarios will be representative, not comprehensive, but they will help executives assess the upsides and downsides of various approaches and determine how much they are willing to invest in the market. Executives should supplement the scenarios with case-based decision analysis of analogous business situations. They might look at outcomes from their own or other fast-food entries in developing markets or consider outcomes from a consumer goods entry in this particular market.

TOOLS: *Qualitative scenario analysis supplemented with case-based decision analysis*

Situation 4: You don't understand your causal model, but you can still predict a range of outcomes. Suppose McDonald's wants to enter a new line of business with a new business model, such as consulting services for food-service process improvements. In this case, executives probably can't define a full causal model or easily identify the drivers of success. However, that doesn't mean they can't define a range of possible outcomes for the venture by tapping into the right information sources—for example, by getting estimates of success from people who have more experience with this business model or by aggregating information about the range of outcomes experienced by others using similar business models. It's often easier to tap into outcome data (and thus define a range of possible outcomes) that define an underlying business model than to ask people to reveal the details of their business models. (That "secret sauce" is confidential in many companies.)

TOOLS: *Case-based decision analysis*

Situation 5: You don't understand your causal model, and you can't predict a range of outcomes. Even a well-established market leader in a well-established industry faces decisions under high levels of ambiguity and uncertainty. When considering how to respond to the recent concern about obesity in the U.S. and the backlash over the fast-food industry's role in the obesity epidemic, McDonald's can't be sure of what effect various moves might have on customer demand. The backlash has the potential to fundamentally rewrite the rules for leadership in the fast-food industry and to make existing decision-making models and historical data obsolete. McDonald's certainly can't accurately forecast future lawsuits, medical research,

Diagnosing Your Decision

When choosing a decision support tool for a major investment, executives need to answer three questions:
- Do I know what it will take to succeed? (or, Do I have a full causal model?)
- Can I predict the range of possible outcomes?
- Do I need to aggregate information?

The answers will point to the best decision-support tools.
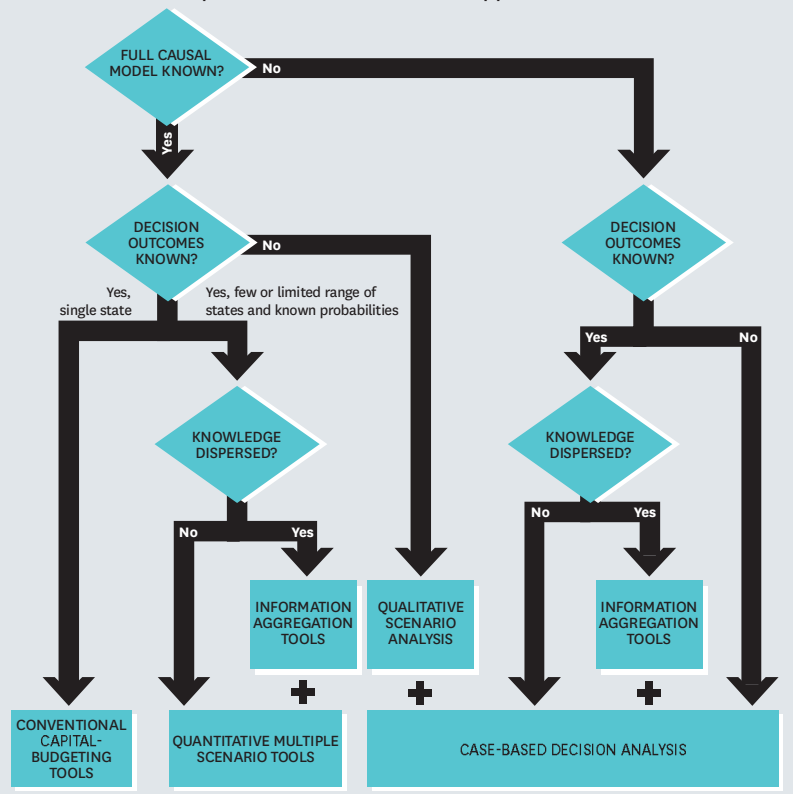
Developing Rigorous Analogies: An Underutilized Tool

Although business leaders frequently use analogies to inform their decisions, many don't do so in a rigorous, systematic way.

For example, it is natural to fixate on the analogous situation that best supports the action you would like to take, ignoring other cases that might provide a broader picture of possible strategies and their outcomes. Case-based decision making provides a structured framework for synthesizing information from multiple analogous experiences and examples. Even when decision makers don't know the exact relationship between critical success factors and outcomes, they can use this method to learn from past successes and failures.

These methods require decision makers to collect a sample of analogous cases, determine the results achieved in those cases, and assess how similar each case is to the decision at hand. The best decision, then, is the one that maximizes the similarity-weighted average of results in the analogous cases. Sometimes the analogy is close to home: Movie producers can compare a project with similar projects from the past; serial acquirers can do the same. Decision makers on less familiar terrain must look to other industries for comparisons, and those compari-

sons will take more ingenuity. (Here it's essential to use structured frameworks for comparison.) Consumer product industries facing digital disruption might look to the unbundling of music and books as an analogy. A company shifting from a product-based to a service-based business model might look at IT companies that have made this shift. When decision makers can't build their own causal models of success, the best they can do is study the successes and failures of others in analogous situations, putting greater weight on the analogies that best fit their own situation.

legislative changes, and competitor moves that will ultimately determine the payoffs of any decisions it makes. When faced with this level of uncertainty, the company should once again rely on case-based decision analysis. Relevant reference cases might include other consumer goods companies' attempts to reposition themselves as healthy or safe alternatives in an otherwise "dangerous" sector or to influence legislation, regulation, or stakeholder perceptions through public relations and lobbying campaigns. McDonald's might analyze, for example, cases in the gaming, tobacco, firearms, carbonated beverage, and baked goods industries for insights.

TOOLS: *Case-based decision analysis*

Aggregating Information

Careful readers will have noticed that the decision tree has one set of tools we have not covered: information aggregation tools. We treat these separately because, for the most part, they function independently of the decision profile questions we pose at the top (do you have a causal model, and do you know the range of possible outcomes?).

The information that managers need in order to make strategic decisions is often dispersed and context-specific. For example, if a company is trying to gauge the synergies to be gained from a prospective acquisition, it's likely that different experts (inside and outside the firm) hold different pieces of relevant information. It's reasonably easy to gather the perspectives of these experts, using tools designed to aggregate information, and to generate a range of possible outcomes and their probabilities. Standard aggregation tools such as the Delphi approach have been in use for decades.

A newer approach to gathering dispersed information is to use information markets (also known as prediction markets) to capture the collective wisdom of informed crowds regarding key variables such as likely macroeconomic performance in the next year or how a proposed product will be received. We should note two limitations of this approach: First, because information and prediction markets are structured like financial securities markets, in which participants can "bet" on different outcomes, they can be used only when executives are able to specify a range of possible outcomes (as in situations 2 and 4 above). Second, using such markets may allow information to leak out that executives would prefer to keep private (for example, the expected revenue for a new drug).

Two alternatives to information markets can get around those limitations. The first is incentivized estimates: People who have access to diverse information are asked to provide estimates of a key outcome, and the person who comes closest to the actual number receives a payoff of some kind (which may or may not be monetary). The second is similarity-based forecasting: Individuals are asked to rate how similar a particular decision or asset is to past decisions or assets. The ratings are then aggregated using simple statistical procedures to generate forecasts for revenues or for completion times or costs, depending on the goal. (This is actually a case-based decision analysis tool.)

ASK YOURSELF:

Is the information you need centralized or decentralized?

If it's decentralized, can you tap the experts you need and aggregate their knowledge?

Is it feasible and helpful to use "the crowd" for some portions of your information gathering?

Is it possible to aggregate useful information from the crowd without having to reveal confidential information?

Complicating Factors

For the sake of clarity, we've presented a simplified set of examples above. In practice, of course, all kinds of complications occur when major decisions are being made. We explore a few of those below.

Executives don't know what they don't know. The model we've developed for choosing decision support tools is dependent on managers' being able to accurately determine the level of ambiguity and uncertainty they face. This may be problematic, because decision makers—like all human beings—are subject to cognitive limitations and behavioral biases. Particularly relevant here are the well-established facts that decision makers are overconfident of their ability to forecast uncertain outcomes and that they interpret data in ways that tend to confirm their initial hypotheses.

In essence, executives don't know what they don't know, but they're generally happy to assume that they do.

Cognitive bias creeps in. Managers' biased assessments of the level of uncertainty they face might lead some to conclude that our diagnostic tool is of limited practical use and might point them toward the wrong approach. Our consulting experiences suggest that most organizations can manage those biases if, when a strategic decision is being considered, managers choose their decision-making approach in a systematic, transparent, public manner during which their judgments can be evaluated by peers. (This will require process and culture change in many organizations.)

For example, any decision maker who assumes that she has a firm understanding of the economics underlying a big decision should be challenged with questions such as, Is there reason to believe that the relationship between critical success factors and outcomes has changed over time, making our historical models no longer valid? Similarly, those who assume that all possible outcomes and their probabilities can be identified ahead of time might be asked, Why are other seemingly plausible outcomes impossible? What assumptions are you making when estimating probabilities? Finally, those who conclude that the relevant information for making the decision resides within the company or even within a small group of senior executives might be asked, If we could put together a "dream team" to advise us on this decision, who would be on it and why?

When asked these questions, decision makers are less likely to assume that their decisions are straightforward or even intuitive and are more likely to turn to tools like scenario analysis and case-based decision making. This is especially important when a relatively new or unique strategic investment is under consideration.

Organizational processes get in the way. Organizations need to develop general protocols for decision making, because political and behavioral pitfalls are rife when money or power is at stake. Here's just one of many examples we could give: We worked with a leading technology company whose forecasting group used the same decision-support tool regardless of where a product was in its life cycle. This made no sense at all. When we investigated, we learned two things: First, business unit heads demanded simple forecasts because they didn't understand how to interpret or use complex ones. Second, the company did not charge business units for the capital used in R&D investments, so unit heads pushed the forecasting team to raise their revenue estimates. As a result of these factors, the forecasts were badly distorted. It would have made more sense for the forecasting team to report to the CFO, who was more sophisticated about financial modeling and also could be more objective about business units' investment needs. It's not possible to design all of the perverse incentives out of a system, but some commonsense protocols can make a big difference.

Decision makers tend to rely on a single tool. We were moved to create the decision profile diagnostic in part because we saw so many managers relying solely on conventional capital-budgeting techniques. Most important decisions involve degrees of ambiguity and uncertainty that those approaches aren't equipped to handle on their own.

It's often useful to supplement one tool with another or to combine tools. To illustrate this point, let's imagine that a Hollywood studio executive is charged with making a go/no-go decision about a mainstream movie. Decisions of this kind are vitally important: Today, the average production cost is $70 million for movies opening at 600 theaters or more (many have production budgets over $100 million), and only three or four out of every 10 movies

Even in situations that seem unambiguous, it often pays to use tools to check for potential biases.

break even or earn a profit. Yet the decision to green-light a project is usually based solely on "expert opinions"—in other words, executives' intuition supplemented by standard regression analysis. In a recent study, two of us used similarity-based forecasting to predict box office revenues for 19 wide-release movies. Nonexpert moviegoers were asked via online surveys to judge how similar each movie was—on the basis of a brief summary of the plot, stars, and other salient features—to other previously released movies. Revenues for the new movies were then forecasted by taking similarity-based weighted averages of the previously released movies' revenues. On average, those predictions were twice as accurate as ones driven by expert opinion and standard regression forecasting. They were particularly effective in identifying small revenue-earning movies. This type of case-based decision analysis is an effective way to tap into crowd wisdom.

Even in situations that seem relatively unambiguous, it often pays to supplement capital-budgeting and quantitative multiple scenario tools with case-based decision analyses to check for potential biases. For example, if your "certain" investment project is expected to deliver a rate of return that is unprecedented when compared with similar projects in the past, that might be more a reflection of overconfidence than of the extraordinary nature of your project. A robust analysis of analogous situations forces decision makers to look at their particular situation more objectively and tends to uncover any wishful thinking built into their return projections.

Managers don't consider the option to delay a decision. Deciding *when* to decide is often as important as deciding *how* to decide. In highly uncertain circumstances—such as a fast-changing industry or a major shift in business model—it's wise to borrow from a different tool kit altogether: learning-based, iterative experimentation. For instance, colleges today are being disrupted by massive open online courses (MOOCs), and most administrators

don't know if or how or when their institutions should react. Rather than make an expensive, high-risk decision now, many colleges are undertaking small-scale experiments to test the waters and learn more about what "success" in this space will look like. (They're also using analogies, of course—for example, by trying to understand whether the unbundling of the music business has lessons for higher education.)

WHAT CAN you start doing tomorrow to become a better business decision maker? Begin by developing your decision-making tool kit more fully. There is a clear disconnect between the tools that are being used and those that should be used most often. Make it a priority to learn more about quantitative multiple scenario tools such as Monte Carlo simulations, decision analysis, and real options valuation. Get some training in scenario planning. Explore the fast-growing academic and practitioner literatures on information markets. Make more rigorous use of historical analogies to inform your most ambiguous and uncertain—and usually most important—decisions. We all use analogies, implicitly or explicitly, when making decisions. The cognitive scientist Douglas Hofstadter argues that analogy is the "fuel and fire of thinking." But it is far too easy to fall prey to our biases and focus on a limited set of self-serving analogies that support our preconceived notions. Those tendencies can be checked through rigorous case-based decision methods such as similarity-based forecasting.

Finally, and perhaps most important, make it a habit at your company to consciously decide how and when you are going to make any decision. ▱

Hugh Courtney is the dean and a professor of international business and strategy at Northeastern University's D'Amore-McKim School of Business and a former consultant at McKinsey & Company. **Dan Lovallo** is a professor of business strategy at the University of Sydney and a senior adviser to McKinsey & Company. **Carmina Clarke** is a senior manager at Macquarie Group.