

1

Introduction

What you will find in this chapter

- A useful definition of innovation that can be applied to an institution.
- Five key mistakes people make when they think about innovation in banks.
- A brief history of innovation in banks.
- An overview of the futureproofing process.
- An explanation of how it is that not all innovation is good.

There is no such thing as a bank that is innovative. At least, that is what I would believe if all I read was the popular press or the blogs of customers. Try this experiment: say the words 'bank' and 'innovation' in the same sentence to anyone in the street, and see if you get much more than a blank look in return.

Most people think of innovation in terms of breakthroughs of the sort one regularly sees coming from high technology companies. They rarely consider that, in their day, ATMs were breakthroughs. They don't think of the revolution of Internet and browser technologies combining to bring banking into the home. Nor do they realise or care that many incremental changes banks implement every day – a change to the call centre interactive voice response, or the update to queue management in the branch – are in fact innovations that other industries have, from time to time, copied.

Perhaps because their customers don't perceive the innovation all around them, bankers have started to believe they aren't very innovative as well. They accept that change will be slow. That they will react when the market demands they do so. And, in fact, that this represents the *prudent* course which will safeguard their institution and its customers.

But there is a problem with this, and that is the pace of change in financial services has accelerated markedly. When it was just regulators, competitors, and markets that were the issue, the glacial engine of *prudence* was entirely satisfactory. But the democratisation of the tools of financial services has changed that. Now *anyone* can do things banks used to think were safely behind the competitive barriers of their very special role in the economy. A savvy consumer is fully capable of using online tools to run a small loan book via emerging person-to-person lending sites. They can pick and choose from dozens of customer service experiences courtesy of the next generation of personal finance software. And they can make international or domestic payments, even to the unbanked, and do so instantly, pretty much without fees.

Many of the commercial, technical, and regulatory barriers which protected banks in the past are about to, or have already, fallen. Their fall brings a groundswell of new change which will utterly defeat *prudence* as a strategy. Prudence is simply too slow to react.

What is needed, then, is a business process which can predictably and reliably respond to all this change, and which doesn't abandon the fundamental tenet of prudence upon which banks must rely. Futureproofing, the subject of this book, is one way of doing that.

Futureproofing is the process of planning what the future might bring and doing something about it. Having read that sentence, you'd be excused for imagining these pages – as so many others at present – might concern themselves with examining doomsday scenarios in which banking no longer exists as an industry. Or if you are more positive, the happy alternative where all present threats to the special economic role of institutions have been dealt with and we continue onward indefinitely. But actually, this is a book which makes only one prediction about the future, and it is one firmly based in historical fact: change is a constant, and there is nothing that can be done to stop it happening.

Once one accepts that change is inevitable, it is only a small step further to the realisation that a business process which can systematically deal with change provides assurances against many of the challenges that might arise in the future.

This book is about building such business processes. It was born from understanding that whilst innovation might be the engine that drives progress and competitive advantage, ad-hoc innovation is, well, random. That randomness, far from providing assurances for the future, is gambling without knowing the odds in advance. Since it is possible to stack the cards in one's favour, it makes excellent business sense to do so.

So, what are the characteristics of an institution that is futureproof?

Firstly, it will have systematised a focus on tomorrow. Many organisations spend the greatest part of their operational attention seeking to optimise the business of today. A futureproof institution recognises that putting structure around future consideration is the best way to avoid surprises. This book explains how such structure can be optimised into a *futurecast*, a substantive vision of alternative futures that can be used to rehearse key strategic decisions in advance.

Secondly, it will embed a business process that actively seeks out solutions for the problems of tomorrow. A futureproof institution knows that ad-hoc, random innovation is just as likely to generate bad ideas as good ones, so it puts sophisticated tools in place to eliminate the guesswork. It recognises also that this is a process that can pay its own way, and demands that each step towards tomorrow makes good business sense.

Finally, a futureproof institution explores multiple things at once. It knows that individual innovations may be successful or may not, but taken as a portfolio, the returns can be predicted with great accuracy.

But futureproofing requires a great deal of hard work. And inevitably, there are plenty of individuals in institutions who argue that the effort, capital, and organisational bandwidth involved is better spent on core businesses. They make the point that banking has been going well since its incarnation in modern form in the late 16th century, pointing to these hundreds of years of development as proof that financial services are able to respond to change without a formal process for doing so.

They would be correct in pointing this out. But now there is an emerging body of evidence suggesting that institutions which proactively and deliberately design their future are significantly superior performers in the long term. And the interesting thing is that such superior performance is almost never about the amount of money spent. Booz Allen Hamilton, who review the top 1000 corporate spenders on R&D every year [1], found there is almost no relationship at all between spending on innovation and superior financial returns. What they did discover, though, was that those companies with a deliberate innovation process – one with links to corporate strategy and customer needs – achieved up to 40% higher growth in their operating income as a result.

With arguments such as these, it is interesting that so few financial services organisations are listed as innovative. In fact, according to Boston Consulting Group and *Business Week* [2], there are only five institutions who make the top 50 innovators globally. That *any* institution is listed alongside such famous innovators as Apple, Google, and General Electric is surprising, given the widely held view that banks aren't innovative at all.

What are those institutions doing to draw the attention of *Business Week*? What they all have in common is that they've developed robust processes to help them design their own futures, and they use them to get reliable and predictable returns from their innovation investments. They are institutions whose futures are secure.

Most banks spend years building their innovation capabilities before achieving this level of mastery. Having said that, however, the basic principles that underlie success are easily understood, and the chief concern is usually operationalising them in such a way that they become a core part of doing business. It is my hope that this book will help you do that in your own institutions.

1.1 WHAT IS INNOVATION ANYWAY?

In many financial services firms, it isn't hard to find groups that are responsible for something that is, conceptually at least, innovation. It is typical that the focus of such groups be laser-sharp on the core business operations of the organisational lines that host them. In fact, in most banks, there are many innovation teams scattered across various silos, though they might not always think of themselves as being part of the innovation function.

It is hardly unusual, for example, for a group calling itself 'Business Development' to engage in new product innovation, whilst sitting across the hall a technology team looks for innovative gadgets they can shoehorn into a banking context. Meanwhile the strategy function is undoubtedly looking at new business models and new markets, and inevitably, the CEO herself is pushing along some pet projects that have an innovative aspect to them.

Unsurprisingly, such diversity of focus leads directly to organisational confusion with respect to the corporate innovation agenda, if an institution is lucky enough to have one at all. And almost certainly, getting to an adequate definition of innovation that works for everyone is pretty much a hopeless task with so many conflicting priorities.

It is useful, then, to look first at common definitions of innovation. This will give us common language we'll be able to use throughout the remainder of this book.

With that in mind, it is possible to classify innovations in two dimensions. The first is the degree of newness incarnated in whatever-it-is. The second relates to the relationship of innovation to the competitive position of the firm. The latter of these two I'll get to in a moment, but first let us look at innovations based on the amount of uniqueness inherent in them.

Breakthrough, revolutionary, and incremental innovation

Innovations which are completely unprecedented are variously called breakthrough, radical, or discontinuous innovations depending on which book you read (I'll call them breakthroughs from now on). Breakthroughs have several attributes: they have few analogues to anything that has gone before, they change the rules of the game substantially in some way, they involve high levels of risk and reward, and they are inherently unpredictable.

History gives us a rich tapestry of breakthroughs to examine: the Wright Brothers with their first aircraft, the creation of the transistor, the discovery of penicillin. What do all these have in common? They were the result of years of thankless work with no guarantee of reward. But more importantly, they all changed the world. It is hard to imagine the inventors knew, when they started their work, how very important their efforts would be to those coming later.

A very common preconception is that innovation teams spend their days doing this kind of work: creation that is so substantially different from what has gone before that the rules of the game are completely rewritten. In fact, only unsuccessful innovation organisations spend *all* their time seeking breakthroughs, as will become evident later.

Nonetheless, there is a deceptive attraction to being first with something that completely changes the nature of a market or product. The rewards may be exceptionally large, and quite often result in a long-term sustainable competitive advantage as well. The downside, though, is that breakthrough innovations, no matter how clever they are, are extremely unpredictable. One cannot easily control when, or even if, one will make a return on what is almost certainly going to be a very large investment up front.

Breakthroughs have occurred from time to time in banking. When they have, they have substantially changed the playing field for everyone. One of the most famous was the introduction of computing to financial services by Bank of America.

As accessibility to retail banking services grew in the 1950s, especially with the rise of the credit card, banks began to struggle with the volume of paper processing required. It was becoming increasingly obvious to everyone that paper was going to put an upper limit on just how large an institution could reasonably grow. Computers seemed one answer, but the application of real computing to banking was substantially delayed by the fact that, at the time, the technology was primarily a scientific and military endeavour. Electronic machines had extremely limited input and output capabilities, which seemed to mitigate against their use in volume transaction processing environments.

Nonetheless, in 1950, Bank of America approached Stanford University regarding the possibility of an electronic machine for data processing [3]. At the time, an experienced book keeper could post 245 accounts per hour, or about 2000 per 8-hour work day. But growing volume was forcing the bank to shut its doors at 2 p.m. each day to deal with the paper backlog and checking accounts were growing at a rate of 23,000 a month. There were few alternatives but automation if the business was to continue its growth trajectory.

An early feasibility study was completed by Stanford University in 1951, leading to a first practical demonstration of a machine in 1955. This machine (called ERMA for Electronic Recording Method of Accounting) introduced several new innovations specific to banking. The first of these, Magnetic Ink Character Reading (MICR), addressed the input problem for volume cheque processing. Another parallel development was the creation of machines that could move paper at speed to the MICR reader. The use of transistors instead of valves made the machine practicable from a heat and power perspective. And magnetic memories were introduced to store instructions and intermediate data.

In 1956 the machine was tested for three months in a branch environment with loads that would be required of a central accounting facility. The tests were successful, leading to the acquisition of 32 ERMA machines by 1959.

The mechanisation of business – in which Bank of America was the pioneering innovator of the day – led to the rise of central accounting as the default mode of operation for banks globally. The breakthrough was so fundamental it was replicated by practically everyone else

in short order. By 1965, almost all banks in the UK and the USA were running automated machines similar to ERMA [4].

Following breakthrough innovation (classified, as before, by the amount of uniqueness involved) is revolutionary innovation. Revolutionary innovations are sufficiently superior to what they replace that they become the default choice for a significant percentage of the market. They offer substantial advantages over what has gone before, but do not, themselves, redefine existing categories or create new ones. The Apple iPhone is a revolutionary innovation. It does not create a new product category (high end mobile phones), but it enhances the concept of an integrated phone, player, and organiser device in such a way that it has become the default choice for many people. It is revolutionary because it is winning share away from incumbent products, rather than changing the way things work fundamentally.

Revolutionary innovations tend to be less risky than breakthroughs, but as might be expected, usually have less upside. The reason? Revolutionary innovations, arising from well understood areas, are far less likely to have the kinds of entry barriers that breakthroughs have. As a result, they are copied more easily. Less than a year after the first release of Apple's iPhone, companies such as HTC of Taiwan were already releasing phones that duplicated some of its best features, for example.

Revolutionary innovations in banking are not that common, but have occurred from time to time. The launch of ING Direct, a Canadian innovation that opened its doors in 1997, is one example. At the time, Canadians had little choice but to choose a low-interest, fee-charging savings account from one of the incumbent big five banks. ING Direct's flagship product, a chargeless, high-interest savings account, was something quite new: it offered bare bones service to low margin customers, but did so at volume. It was immediately a runaway success. Apparently customers were over-served by the features of accounts they could get at their traditional banks.

In 1999, ING Direct opened in Australia, disrupting the industry there as they had done in Canada. Once again, the successful formula was repeated: provide a bare bones service and pass on those savings to customers. I recall being in a meeting with a senior banker in Australia at this time, during which he expressed his irritation that ING was 'borrowing' the use of his institution's channels without paying for them. His complaint stemmed from the fact that ING Direct offered a branchless service, and therefore customers were forced to use the facilities of his bank in order to get funds in and out of their ING accounts. Bankers' complaints aside, ING Direct in Australia went from standing start to the sixth largest retail bank in a few short years.

The next year, 2000, ING Direct opened in the USA, again repeating its successful branchless model, and except for some trademark 'ING Direct' cafes in key markets, remains relatively bricks-and-mortar free. It has now grown to become the largest direct bank in that country.

ING Direct is now operating in the UK, Spain, Germany, Italy, France, and Japan. Its revolutionary model that cut service back as much as reasonably possible and returned customers the savings is one that is, apparently, easy to transplant across cultures and geographic boundaries.

Finally, we come to the least new of all types of innovation: incremental, also known as continuous innovation. Incremental innovation takes what is well known and makes a minor improvement with a positive payback. Incremental innovations may not be visible outside an organisation: they are characteristically small, probably very specific to an institution's individual way of doing things, and are relatively low risk.

In many countries, it is possible to sign up for 'pre-pay' mobile phone contracts. Telecommunications firms provide potential customers with a free or low-cost SIM card, which is then

‘topped up’ with credit. Customers are allowed to make phone calls up to the value of their credit before they have to ‘top up’ again. Initially, the process of adding credit was available only through the shop fronts of the mobile phone operators, but in some countries, banks were quick to spot an incremental opportunity: allow mobile phone top ups through their ATM networks.

The business model that supports ATMs is very specific: one wants as much cash dispensed as possible, in the shortest amount of time. Ideally, one wants the customers of other banks to use the institution’s ATMs as well, since this provides a rich source of fee income. Consequently, locations for ATMs are hotly contested, and the best spots are almost always already filled with one of the ubiquitous cash dispensing machines.

Now, most opportunities to use ATMs to dispense things other than cash have a huge downside: the time taken to operate the new dispensing function is generally greater than that for cash, and consequently revenues from individual machines tend to fall. But mobile phone top ups have none of these disadvantages. Customers simply enter their mobile phone number, and the credit is deducted from their account and added to their phone automatically. It is a nice piece of new fee revenue that institutions are able to acquire from telecommunications companies.

Mobile phone top ups at the ATM are an entirely incremental innovation. They take what is already in place – the ability to dispense cash and provide minimal account information to customers – and twist it just a little to provide a new service consumers find valuable. Unsurprisingly, in most countries that do top up at the ATM, it has become a ubiquitous offering from everyone who runs ATM networks.

If you go back over these three types of innovation I’ve just covered, you will probably think of examples from your own organisation. That is not unusual: it is the hallmark of an appropriate innovation strategy that things are developed from each category. But what is not generally obvious is that a much broader definition of innovation is possible: *anything* that is not presently being done by an organisation is an innovation opportunity. Market-wide uniqueness doesn’t come into it. Innovation is the process of introducing new things, certainly, but it only has to be *new* to your institution for it to be an opportunity worth exploring.

Disruptive and sustaining innovation

As I mentioned earlier, it is possible to classify innovations in two dimensions. We’ve just looked at the first dimension: how genuinely new whatever-it-is is compared to what it proceeds. The second dimension is the way institutions and markets respond to the innovation itself. Banks react to things which are new in very different ways depending on whether the new thing *sustains* or *disrupts* their current operations. This classification was first proposed by Clayton Christensen, author of the hugely influential book *The Innovator’s Dilemma* [5]. His theory of disruptive innovation has a well-established track record of explaining why it is that companies in different industries ignore some innovations and support others. We’ll look at the mechanics of disruptive innovation in detail in Chapter 2.

In any event, a sustaining innovation is one that creates additional value for a firm by enhancing the products or services already being offered. By increasing the functional capabilities of existing offers, new customers can be reached or existing ones better served. Sustaining innovations create new value for banks organically in the short and medium terms. They do this by delivering growth along established trajectories in a predictable manner.

Although this may be a much disputed point, Internet banking is a sustaining innovation. It makes it possible for existing bank customers to use their products in new ways, and certainly with much greater convenience. Internet banking is also a revolutionary innovation: it applied the existing technologies of Internet networks and web-based browsers to the problem of self-service.

Internet banking, however, did not create a parallel industry of Internet-based competitors with much possibility of eroding banking business. Netbank, the most prominent online-only bank, was established in the 1990s, and roundly hailed as a disruptive innovation likely to change the face of banking forever. It was a sign, many thought, that the branch was dying, if not practically dead already.

Netbank failed in 2007. The failure, according to prominent analyst and blogger Jim Bruene, was 'primarily from poorly underwritten loans, both prime and sub-prime, and most of those originations came the old-fashioned way, through face to face broker sales' [6]. The lesson of Netbank is that whilst the reliance on Internet delivery was a novel innovation for the time, the traditional business model was still very much centre stage.

On the other hand, ING Direct, also an Internet-based bank, is extremely disruptive in any market it enters. The difference is the business model change it copies to its direct channels, represented by reducing service to the bare minimum and passing on the savings deriving from this to customers in the form of much higher interest rates.

Breakthrough and revolutionary innovations are often confused with disruptive ones. Actually, though, disruptive innovations are usually not so much about brand new capabilities, as they are about creating new value propositions. These new propositions are deceptive to an incumbent player in a particular market, who will likely ignore them as not core to their own business. They become disruptive, however, when expansion of the entrant causes the new and old value propositions to overlap.

A disruptive innovation usually starts life as a poorly performing, inferior product compared to those of incumbents. But the fact the product or service does less, means it brings with it quite a different cost and value structure than what it precedes. This is attractive to a small segment of the market, one which is probably uneconomic to a mainstream institution. The market may, in fact, be so small or so low margin that not only is it unattractive, it is actually impossible for an incumbent to enter it at all. A disruptive innovation, being less capable, and therefore less expensive, may find attractive returns in this low end space.

Over time, the disruptive innovation improves its performance, and in the end, is as capable as anything else in the market. But this time, the disruptive innovation competes against existing products from a significantly better cost base.

When UK-based Zopa became the first company in the world to launch a peer-to-peer (P2P) lending operation in March 2005, they implemented a radical concept: the 'Zone of Possible Agreement' (from which the name is derived). The term refers to that price point where both borrowers and lenders agree that a particular interest rate is fair to both sides. Zopa links both parties up at this price and facilitates the actual transaction. No bank is involved. This disruptive innovation has now been copied by companies in many countries, including the USA, Australia, Germany, and the Netherlands, and the growth of this new model seems to be gathering pace.

Initially, the facility that Zopa offered to deposit customers was significantly inferior to that of banks: depositors had a much higher risk of losing their money than they would have done using a traditional banking account. Nonetheless, a certain fringe of depositors – those early

adopters who were familiar with social networking and had an appetite for risk – began to use the service.

As with all disruptive innovations, however, the capability of the product swiftly improved. Lenders were given more tools, and much more certainty about their returns, making the product a much better fit with a wider market.

In 2008, analyst firm Gartner forecasted, controversially, that peer-to-peer lending might take 10% of all retail lending volume by 2010 [7]. Initially, I was sceptical in the extreme of this, and said so on my blog [8], but when you examine the dynamics of disruption (see Chapter 2), it is possible to see a mechanism at work that might breathe truth into the prediction.

Disruptive innovations are needed to help banks deliver robust growth in the long term. As with any long-term strategy, execution is the problem. Disruptive innovations tend to have very small returns in the beginning, insignificant compared to the main business lines of any bank. But the right disruption, given time, can grow into a business with very substantial scale, as it looks that Zopa may do. An institution able to create and nurture such innovations has, indeed, futureproofed itself. Unfortunately, as we'll see later in this book, doing disruptive innovation is probably the hardest thing a bank can do.

The difference between innovation and invention

Before we leave the definitional part of this chapter, it is useful, also, to draw out the difference between *invention* and *innovation*. Joseph Schumpeter, a Harvard economist who rose to fame through his ground-breaking work in entrepreneurship, was one of the first to make a distinction between these terms. In his conceptualisation, an invention is largely a theoretical construct, an idea with, perhaps, some evidence to prove that it can be implemented in the real world. But *innovation* takes an invention and puts it into practice. It converts what was an initial theoretic construction into something that can do useful work. Another way of looking at this is that invention occurs whenever a concept is created for the first time. An innovation takes that concept and turns it into something real, something that can make real returns.

In general, institutions do not lack invention. Most people have experienced the situation where a few people in a room with a whiteboard create lots of solutions to a particular problem. Usually few, if any, of these solutions (inventions) make it from the whiteboard to practice. It is only the small number of cases that do that are innovative. Innovation is invention plus execution. And, as you'll read later in this book, the process of *killing off* inventions that aren't going to make it is a key part of ensuring a balanced return on an innovation portfolio.

1.2 WHAT HAPPENS WHEN YOU *DON'T* FUTUREPROOF

I described how futureproofing is a business process an institution can use to ensure it correctly recognises things that might affect it in the future, and then respond in an appropriate and measured way to generate a return. One question that's often asked, however, is what the consequence of *not* futureproofing might be. Have not banks been operating, with largely the same services and a fundamentally unchanged business model, for hundreds of years?

Certainly they have, but the pace of change in banking (and in most other industries as well) has increased markedly in recent times. The upshot is that the time between a trend being noticed and its implementation by a competitor is becoming increasingly short. There isn't time to dither around before making a decision: what is needed is a system that can respond routinely to change. Change is the only constant.

Recent research proves the value of having a process that looks forward. A study of North American financial services chief executives conducted in 2007 established a link between the time senior leaders in banks spent looking at the future, and the innovation success of their institutions [9].

The researchers began by examining implementations of Internet banking at 169 banks, starting from the moment it dawned on leaders that online banking on the web might be important, through to its eventual near-universal roll out. They then created a statistical analysis of public statements made by bank leaders to get an indicative measure of how much time each was spending thinking about the future.

The first thing they discovered was that as a result of focusing on potential future states, banks were not only better at making predictions about the future, they were also *much* better at responding. In the study, the average time to respond to the online banking opportunity was just over four years, but the worst performer took nearly nine and a half years before they had something customers could use.

The second, more startling thing, however, was that future-looking banks not only managed to respond more quickly, but the *breadth* of their response was superior. The first Internet banking sites in particular markets were, on average, delivering just over three new innovations each to their customers, but the best of them had up to five. Clearly, the bottom line impact of such a substantial functional difference between leaders and followers is exceptionally valuable.

Whilst evidence such as this is helpful justification of the value of futureproofing processes, one doesn't need to go much further than the rise of PayPal over the last decade to understand what can happen if an appropriate strategy for responding to change is *not* part of the way institutions do business.

In December 1998, a company called Confinity was founded on University Avenue, Palo Alto. The new company set out to explore whether the most popular digital organiser of the day – the Palm Pilot – might make a good electronic wallet that could be used to beam money between owners. Just down the road, another company, X.com, was founded to look at the opportunities surrounding online payments. When the two merged in March 2000, the combined entity, renamed PayPal, swiftly became the preferred means of payment for more than half of consumers who had begun using online auctions. Two years later, when auction giant eBay bought PayPal, its valuation was \$1.5 billion. At the time of writing, it operates in 197 countries, provides payment services in 17 different currencies, and has more than 150 million accounts.

The success of PayPal was the result of the confluence of several things. Existing bank products at the time did not lend themselves to person-to-person payments. Paper cheques and direct transfers (in the countries that had them) took too long to settle in a world where auctions completed instantly, and many sellers were unable to take credit cards. Later research [10] found that payment instrument choice on eBay was influenced almost entirely by the certainty of attributes of the product being acquired (i.e., colour, size, and so forth), but in the online auction space, not only did consumers have less certainty about the product they were buying, they were dealing with uncertain individuals as well. Consumers demanded something new to go with this new shopping experience, something that enabled them to reduce all this unwanted risk when it came time to pay.

The innovation of PayPal was that it created a layer atop existing financial relationships that consumers already held. The new layer made it simple, safe, and fast to send money between people. It swiftly became ubiquitous.

As early as 2001, banking journals began to report there might be interesting things on the horizon. One went as far as to note that whilst the number of alternative payment systems was in decline, payment systems associated with inherent transaction streams (such as eBay) were thriving [11]. The publically available pre-IPO prospectus issued by PayPal around this time indicated it was processing over 165,000 transactions per day with an average value of around \$50.

Competitors, not coupled with an inherent source of transaction volume (as PayPal was with eBay), swiftly found they were unable to compete. In 2003, Citibank's C2it service closed, followed closely in 2004 by the cessation of Yahoo's PayDirect offering. Then, in 2005, Western Union disposed of its person-to-person service, BidPay.

In 2006, consultancy Booz Allen Hamilton [12] made a prediction with respect to online payment providers:

'If existing providers (predominantly the card issuers and acquirers) do not find an effective counter strategy we believe they could lose 10–20% by 2008 and in the long term up to 30%.'

In my own discussions with bankers around the world on this topic, I conclude these predictions are, if not already true, very close to being so, at least in developed markets. The online payment opportunity is one that banks allowed to slip through their fingers.

Hindsight is a wonderful thing, but the attraction of a privacy and security layer above traditional payment instruments is retrospectively obvious. The problem now is that PayPal is so large (by number of accounts, the largest financial services provider in the world) that competitive responses by banks are somewhat limited.

1.3 FIVE THINGS THAT INNOVATION IS NOT

Innovation is very much on the corporate agenda of a large percentage of financial institutions. Despite that, I am surprised how often I come across misconceptions about innovation and how it is managed.

For those whose job it is to manage the innovation agenda, this can be particularly problematic. The function can often be seen in terms of way-out things new and exciting, disconnected from the core business. When that happens, innovators are liable to get labelled: geekery of any kind rarely drives business returns. Their relevance in the strategic context gets called into question.

Recently, for example, I was approached by an individual in another bank convinced an innovation programme should be all about trying to get new gadgets into the branch. His assertion was that if we weren't doing highly visible public-facing things, the whole conception of innovation in financial services was bankrupt. This individual made a mistake one often sees: he narrowed the innovation agenda to such a degree it would be hard to make a decent return no matter how good the toys were.

For institutions that are starting their innovation journey, it is critical to dispel these kinds of myths immediately. They have a negative effect that can taint efforts for ages afterwards. Such tainting is exacerbated, unfortunately, as some organisations have experienced faulty innovation attempts in the past. Getting distance from these historical issues is critical.

Here, then, are five things people believe about innovation (and about innovation teams) that should be dispelled as quickly as possible.

Innovation is only about things that are completely new

Earlier, I discussed the difference between breakthrough, revolutionary, and incremental innovations. The former two certainly get the most attention, and consequently most people believe innovators do very little incremental. Incremental is the realm of business-as-usual. In fact, people are surprised when I tell them most returns from good innovation programmes come from incremental innovation.

Perhaps the most famous incremental innovator is Toyota. The volume written about this company and its rise from relative mediocrity to global dominance on the back of small, quite basic changes is monumental. Founded in 1937, the company started commercial passenger car production in 1947, and by the 1980s was consistently ranked higher than any other manufacturer in owner satisfaction surveys. Attention to detail, and making small changes to create lasting improvements, led the company to become the largest automotive manufacturer in the world by 2007. Clearly, incremental innovation can pay, even if the individual changes aren't exciting and high profile.

Convincing people that small improvements are important is a big challenge for an innovation function. A common response to the idea that innovators should do incremental is that innovators who do so are reducing themselves to optimisers.

The emergence of instrumented methods for process improvement – Six Sigma is one – has made it easy to confuse optimisation with innovation. When product-type people create a new savings account, are they actually innovating, or are they optimising the savings category based on their expectations that making the change will result in greater market share? This thin line is the principal reason people imagine that true innovators would never concern themselves with anything which fails to change the game completely.

It's easy to understand why there is this confusion, but the difference between optimisation and innovation is really quite simple: you optimise by pulling various levers you already have to get better results. Creating a new savings product is an example: one captures a greater share by varying interest rate, fees, and terms and conditions.

True innovations, on the other hand, create *new* levers altogether, or modify the *range* of existing levers. They don't just change the position of the dials, regardless of how unique the new combination is.

ING Direct, for example, changed the range of its interest rate dial on its savings account for customers by reducing its costs and passing on the savings. This let its product managers price its savings accounts in such a way that ING was able to claim large shares of markets even when it had no traditional banking presence.

Innovation is speculative and risky

Financial institutions spend a great deal of time managing risk, and indeed, it is a core capability for any bank that wishes to lend money successfully. At the same time, banks are not known for innovation, another process which would seem to lend itself to strong risk management principles. There is a reason for this, and it is that innovation *seems* to come with risks unquantifiable in advance. Banks are good at managing predictable risk, but how is one supposed to predict the success of a radical new product? Far better to commit that capital to the lending book where, at least, the statistical return generated is well understood.

It is true that the quantitative risk associated with an individual innovation is very hard to forecast in advance. Some studies have suggested, in fact, that a new product introduction is

less than 25% likely to succeed. That situation makes it *appear* as if investing in innovation is a rather poor proposition when compared to the opportunity cost of the money.

Actually, it is possible to predict – to a degree – how risky an innovation might be, and to make determinations about the likely returns. Later in this book, some of the models and other analytic tools which can be used to do that will be discussed, but the fact is, you can't always get a picture of the total risk associated with individual innovation no matter what you do.

Single innovations are risky. Which is why a portfolio of innovation investments is required.

Such a portfolio is no different from the basket of business in a loan book: some will fail, and those most likely to do so should command a higher return. Taken as a group, though, very predictable returns can be made. In fact, given a big enough portfolio, the return can be predicted pretty precisely. One invests in a range of innovations – some more, some less risky – in order to guarantee that a particular level of return is achieved.

In most cases, it is true that incremental innovations carry the least level of risk, so they are the ones that have, most of the time, small returns. By weighting your investment strategy towards incrementalism, you get lower overall returns from innovation, but much greater certainty that you will actually achieve your numbers.

On the other hand, breakthrough innovations are typically very risky. There is a lot that can go wrong: technology might need to be invented, for example, or it might not be possible to forecast the demand curve sufficiently in advance to know that the introduction will be a success. But breakthrough innovations are also the pot of gold at the end of the rainbow: they generate windfall returns when they are successful.

Let's face it, if innovation were so speculative and risky, the entire venture capital industry would be out of a job. VCs typically invest in early stage companies, most of whom will fail. Still more will persist in a kind of living/dead state – neither making nor losing money. Nonetheless, most VC firms make substantive returns for their shareholders – returns well above those available from other investment opportunities. They do it by managing a portfolio of investments.

What return is good enough? That's a difficult question, but generally, it is necessary to achieve some multiple of the institution's internal cost of capital. Obviously, if the returns of innovation efforts are less than that, questions about opportunity costs are going to rear their head. One mantra I recently heard was 'we're in the business of lending, not spending'.

However, to achieve many multiples of the internal cost of capital might involve taking an extremely high percentage of risky – disruptive, for example – innovations. A reasonable benchmark is to examine the returns of existing business-as-usual investments, and then set the bar some material way beyond. One wants to demonstrate that innovation is a preferred investment activity, whilst not accepting so much risk that it is impossible to deliver reliably.

One final point: it is absolutely key that innovation teams have a lot of activity to show for their efforts. They must have a portfolio of things going at any one time, since most will not get through the futureproofing process and generate returns. The wider the portfolio, and the better the mix of activity across breakthrough to incremental innovation, the more risk can be controlled.

One side effect of all this – which we'll discuss later – is that the innovation team will not be likely to have enough resource to do everything. It is therefore important that the innovators involved are superlative influencers: they have to be able to win support from executives to support their innovation agenda. When you go it alone, you just can't get the breadth you need to make innovation a good proposition.

Measurement is hard

Having just read the discussion of the previous section, you'd be forgiven for thinking the only innovation measure that makes any sense at all is return on investment. As I've mentioned a number of times already, unless you are creating new revenue (even if indirectly), there isn't much point in doing innovation. And the ultimate goal of the whole futureproofing process is to protect the ability to create future revenue, after all.

The reason that people think measuring innovation efforts is hard is because it *is* hard. One big problem is this: you create something new on the basis of some future return. There is inevitably a time lag between the moment of investment and the payback. The delay between investment and payback can, in real terms, be some years. This gap disconnects the innovation effort from the actual business outcome. Waiting years to know how well an innovation programme is performing is rarely acceptable to those who make the decisions about funding.

Another problem with metrics that focus only on matters financial is that they tend to make it impossible to do anything, no matter how important, that doesn't come neatly associated with a convenient cash return. Some innovations – those which drive productivity improvements, for example – are notoriously difficult to associate with hard numbers. If you have only cash-based metrics, you are likely to get only cash-based innovation. This results in leaving a great deal of opportunity on the table.

It is necessary to have measurements that touch every stage of the futureproofing process. Without defining the futureproofing process now (it is covered briefly at the end of this chapter and much more extensively elsewhere in this book) there need to be four kinds of measures. Firstly, you need to know how good you are at spotting the trends that matter to your institution. Secondly, having come to a good idea of the future, your idea-harvesting mechanism needs to be instrumented. Thirdly, you have to be able to measure how all these ideas get coupled to execution, and finally, what actual results were achieved.

Having metrics that touch every stage of the futureproofing process is important because it enables an institution to *optimise* the futureproofing process as it goes along. It is almost never the case that everything will work efficiently first try. Pulling levers to optimise is a functional necessity to ensure that an institution makes the returns it needs from its innovation efforts.

Finding the 'mega-hit' is the best way to success

I've lost count of the number of times I've had to disappoint someone when I've told them that the role of innovation programmes is *not* to do work that results in breakthroughs because they *might* be blockbusters. Innovation programmes must be about delivering predictable returns if they expect to continue in times bad as well as good.

This is an important point. In the course of research for this book, I spent a great deal of time with innovators in banks around the world, both those that were successful and those not. One key theme emerged: the average time an innovation team exists is about 18 months. Those that last longer have done so because they've managed to deliver predictably, and we'll be talking about the processes needed to get to that point later in this book. All the rest were cancelled because they failed to generate sufficient returns to justify the resources they were consuming.

With this in mind, everyone imagines that if they were to create a mega-hit, another Google, say, their careers would be made. And they'd be right. But how often does that perfect

storm of technology, business model change, and consumer need happen? Answer: very, very infrequently.

Mega-hits, or the breakthrough disruptive new, is what many people imagine when they hear the word 'innovation'. The fact of the matter, however, is that real disruptions, the ones that result in long-term competitive advantage and windfall returns, don't happen all that often. Later in this chapter, I'll look at the history of disruptive and breakthrough innovations in banking. There have been some, of course, but they tend to be few and far between. Arguably, there have been only a few *really* disruptive innovations that have made a significant difference. The rest have largely been variations on a theme.

An innovation strategy that seeks returns only through disruptive breakthroughs is usually a very bad bet given how irregularly the financial services industry has actually experienced them. A far better approach is to concentrate on revolutionary and incremental opportunities.

But even revolutionary innovation has its challenges. I once had someone come up to me and explain that he'd just joined the bank and wanted to do 'revolutionary innovation'. The problem was that what he really wanted to do was *talk* about doing innovation and have someone take his ideas and implement them. Some of the things he was talking about actually *were* revolutionary. But the problem with all of them was that we'd have had to break through some pretty big barriers to execute.

Firstly, they were expensive. Revolutionary and breakthrough innovation almost always is, and money does not grow on trees. The more money you have to find to invest in a particular innovation, the less tolerant you can afford to be about risk.

But secondly, the more revolutionary something is, the more support one has to generate with stakeholders. As we'll see throughout this book, the process of doing this is a question of influence. Influence is the goodwill one has previously developed with potential sponsors. It is a scarce resource that is consumed during the process of shepherding an innovation to the point where it gets the green light. Influence and money, by the way, are often interchangeable.

So, even revolutionary innovations should be considered carefully. Does this mean that innovators should ignore them entirely? Of course not, but such investments have to be considered in the context of everything else the innovation programme is doing. A well-balanced innovation strategy will likely spend most of its money doing incremental work, certainly, but there is plenty of scope to invest in breakthrough or revolutionary innovation once the bills are paid.

Incidentally, I'm anticipating much feedback on this point. In many cases, at conferences and elsewhere, people have specifically challenged the view that banks should concentrate on incremental innovation, suggesting that incremental should be everyone's day job. Real innovation programmes, they opine, should be about changing the game. The view taken in this book isn't necessarily in disagreement with that. It doesn't matter where the innovation occurs, so long as it does. The formal innovation team might not specifically focus on it, but when you take an institution as a whole you'll likely find that most of the return on innovation investments is coming from incremental.

Ideas are the thing

Schumpeter really hit the nail on the head when he contrasted invention and innovation. It is way easier to invent than to innovate, and that's because inventing is an extremely creative, exciting exercise. Innovation, on the other hand, is mainly about hard work. It's about execution, or in other words, actually doing things.

In most companies I've worked with, there is never a shortage of ideas. In fact, every time a meeting is held in a room with a whiteboard, the probable result is a pile of invention. The thing is, few of these interesting squiggles get turned into reality.

Some years ago, I had a very personal experience with this. It was the time that people were implementing their first Internet banking solutions, and my team and I came up with something that we *thought* was a breakthrough: why not screen-scrape all these banking websites and put the results on a single web page? We sat on the idea, and were then stunned when Yodlee entered the market, simultaneously defining account aggregation as a category, a few months later. This story proves two things: firstly, that invention is pretty much valueless by itself if you don't do anything, and secondly, that invention tends to happen simultaneously in lots of different places at once.

Good ideas are prompted by market conditions that are rarely localised. That's why you often get them appearing in lots of places at once. Even if an idea is completely novel and unique, it will not generate a return if you let it sit. *Someone* has to do something with it.

In recent times, open innovation has become very fashionable. In open innovation, you collaborate with other organisations that have capabilities you do not to create the new product, process, or other change you want. But open innovation is really a shortcut from an execution perspective: you get to do the ideation (the fun part) and then buy in the stuff you need to do most of the execution. And there is nothing wrong with that, assuming the internal landscape of an institution permits it.

Perhaps the most prominent industrial example of open innovation is the latest product from Boeing: the 787 DreamLiner, an aircraft due for delivery (at the time of writing) some time in 2010. The idea was to create a faster, more fuel-efficient airliner. But execution required a host of new technologies, new supply chain techniques, manufacturing innovation, and a great deal more besides. That's a *lot* of execution for one company to come up with. So Boeing bought in a lot of what it needed from its partners. Mitsubishi Heavy Industries and other Japanese companies, for example, are responsible for the manufacture of the wing, whilst Saab makes most of the various kinds of doors on the aircraft. The end result? Boeing had a new aeroplane, but had to spend much less to realise it than otherwise.

Whilst the value of individual ideas is low, the process of capturing ideas and evaluating them is of critical importance. Without a systematic way of harvesting ideas from employees, customers, and partners, great creative wells are left untapped. An innovation programme has no connection to the business that supports it without some way of systematically harnessing all this creativity. Then, too, some of the simplest problems to solve might not be noticed at all if there is no process for stakeholders to report them.

Even solving a simple problem – like moving a check box on a form – is something that can contribute to the overall returns of an innovation programme. Unfortunately, people with innovation in their job title are unlikely to be concentrating on such things, which is a pity. You don't have to commit all that much execution to make them happen. Toyota, as mentioned earlier, rose to dominance in auto manufacturing by repeatedly making small changes.

1.4 150 YEARS OF INNOVATION IN BANKS

With all those preconceptions out of the way, we turn now to an examination of a few of the innovations that *have* been significant in our industry in the past. To this end, what follows is a potted history over the last 150 years, covering some of the most significant changes in that time. As I noted earlier, breakthrough innovations have been relatively rare in financial

services, though revolutionary and incremental ones occur much more frequently. Regardless of how genuinely new these introductions were, however, there is one thing in common: some institutions led, and the rest responded. Those banks that profited were clever in spotting the emerging opportunities, and actualising them as part of their ongoing operations.

As an aside, my typification of particular innovations in what follows as incremental, revolutionary, or breakthrough is highly subjective. What is breakthrough for one organisation may, indeed, be seen as incremental for another. This, by the way, is true also for the classification of some innovations as sustaining or disruptive.

Our historical tour commences in the early part of the 19th century, with the development of the telegraph.

The late 19th century

The telegraph was first demonstrated practically on an experimental line between Washington and Baltimore. By 1846 lines were available to Philadelphia, and in 1848 to New Orleans. The effect of this innovation was electrifying (pun intended): it reduced the price associated with the information differential between New York and regional stock markets [13] hugely.

But the implications of the telegraph did not stop there. Throughout the 19th century and the early parts of the 20th century, trade between cities was governed by differing exchange rates, also the result of information differentials between cities. As a result, transactions that were geographically separated had to consider the cost of settlement. In the absence of a modern payment system, this meant physical transport of gold, the value of which could vary between geographies. The telegraph changed that, and banks were quick to recognise the potential. In fact, the *New York Herald* complained on 3 March 1846 that 'certain parties in New York and Philadelphia are employing the telegraph to speculate on stocks'.

The telegraph, then, was one of the first revolutionary innovations to be adopted by banks. No new technology was created by institutions themselves, but its application to the business of finance changed the game entirely.

The transatlantic cable (an incremental innovation on the original revolution of instantaneous messaging, from the perspective of the banks) was completed in 1866. With ten or more years' experience using the domestic telegraph in the United States, banks on both sides of the Atlantic were quick to adopt this new means of integrating their markets. They had little choice but to adopt the innovation immediately: the competitive advantage of knowing a price for a security on the same day was unimaginably significant.

The first part of the 20th century

Though these communication innovations served to unify markets, they did little to change how front- and back-office procedures were run. At the time, these operated through significant delegated authority to managers in the branch network, with a system of draconian inspections carried out by head office staff from time to time. Systems were entirely paper-based: ledgers and passbook controls being the principal means.

That worked relatively well, but communication technologies caused the volume and pace of transactions to increase. The race to process the paper was one that banks were in danger of losing. In order to address this, leading banks started to look at the possibilities of mechanisation. By the 1930s, mechanical adding machines had been introduced in a number of institutions. The technology of the time was revolutionary, but from the bank's perspective,

this was an incremental innovation: making it faster to do existing work. Later, incremental innovations such as punched cards continued to enhance the speed at which paper could be processed. With eventual broad availability and declining cost of such machines, banks were poised to revolutionise their delivery of services to the public: increased worker productivity meant that more products could be offered, to more people. Both the size and number of branches expanded rapidly as a consequence.

This was the time that a new kind of financial services product started to appear (the credit card), though it wasn't, initially, noticed by banks. The rise of the automobile, and the consequent need to acquire and pay for fuel, led fuel companies to consider how they might make it easier for their customers to deal with them. In 1920, they started accepting cards, at that time little more than a place to record a reference to the customer's account, as well as cash. This innovation – perhaps the first breakthrough in recent banking history – proved popular, and by the 1930s multiple fuel companies had started to accept each other's cards.

The fifties, sixties, and seventies

The breakthrough represented by the first fuel cards went unnoticed by financial services companies for 20 years. The paper cheque was the pre-eminent non-cash payment instrument of the time, and institutions had developed mechanical means of processing them at volume. So when Diners Club invented the first modern credit card in the 1950s, banks were not expecting the effect they would have.

Diners Club was an incremental innovation on the original fuel card: rather than a card which worked for merchants of one type only (fuel stations), Diners envisaged a card which could work with any merchant at all. The Diners Club card was initially launched in New York with 14 restaurants and 200 cards in 1950. But by the end of the year, over 20,000 individuals had the card, and more than 1000 restaurants were accepting it.

Eight years later, Bank of America launched BankAmericard, which later evolved into Visa. Then Barclays, in the UK, followed some years later with a credit card of its own, largely built using systems imported from Bank of America.

Credit cards became something of a phenomenon. They were massively adopted, and suddenly, once again, banks were drowning in paper. The mechanical processing methods of that time had become the critical limiting factor that would constrain additional growth. That's when Bank of America created the first business computer, ERMA, following its initial experiments in 1950. The importance of this breakthrough innovation cannot be understated. It transformed the world of financial services completely: prior to ERMA, banks operated on the basis of paper and delegated authorities, and afterwards they didn't. It was the first time that people and manual processes were *eliminated* from the business of banking, a trend that has only recently started to reverse.

These initial steps towards electronic transaction processing didn't completely resolve the issues banks were experiencing. All credit decisions were still being delegated to skilled individuals at the branch level, a situation that put constraints on how many of the new credit cards banks could issue. There was no robust way to determine the creditworthiness of potential card applicants *at scale*.

In order to facilitate the growth of the new product, a mechanical means of determining the likelihood of default on credit needed to be found, giving rise to the breakthrough development of credit scoring. Initially pioneered by mail order and specialist finance houses in the United States during the 1950s, credit scoring was quickly picked up by a significant percentage of

institutions. It was then expanded to other kinds of lending, such as mortgages, particularly in the last 20 years. The most recent innovation in this area has been the introduction of scoring for certain kinds of business lending as well [14].

But the breakthrough of credit scoring had another effect: it started the decline of branch-based decision-making, a trend that has continued till the present day, and set the scene for fully automated self-service lending.

Whilst banks were reinventing themselves in their modern image, the pace of technological development in business computing sped up. Amongst the developments of this time were high-capacity memories and magnetic disk storage, enabling the new machines to store and retrieve data at high speed. These new capabilities, in turn, led to new software techniques, primarily in the form of various database management systems.

Bankers, again drowning in paper-based systems because more customers than ever were now using cheques, immediately saw an opportunity to automate clearing operations. The first bank clearing house was established in London in 1770, but it wasn't until 1968 that institutions found a way to automate the process with their new technological capabilities. The revolutionary innovation of interbank automated clearing was first introduced in the UK in 1968. BACS (for Bankers Automated Clearing Services), as the system became known, expanded quickly, bringing with it associated advantages for member banks, and by 1976, it was handling just under 100 million items a year. At the time of writing, the network handles over 5.5 billion payments per year [15]. In the USA, meanwhile, the first Automated Clearing House (ACH) was established in California in 1972, with other regions following rapidly.

Technology quickly became the *key enabler* for the great majority of banking innovations that followed, and the pace of change began to build. 1967 saw the next breakthrough, when Barclays deployed the first ATM in London [16]. The early machine, built by De La Rue, used a mildly radioactive ink on a paper token, since the magnetic stripe would not be developed by IBM until the following year, and was capable of dispensing fixed amounts only in paper envelopes. But whilst the technology of this first ATM wasn't that advanced, it heralded momentous changes in service delivery for banks: the breakthrough was the idea that customers could *serve themselves*. The modern conceptualisation of the networked ATM was invented by Don Wetzel in 1968 in Texas [17], and included the IBM magnetic stripe innovation as well as modern cash-dispensing apparatus.

The key insight of the ATM – that customers valued convenience over relationship – led to rapid adoption of the machines globally, and forced banks to make additional investments in telecommunications. Early ATMs were entirely offline, so no balance checks were made before dispensing cash. Connecting the terminals to the centralised accounting function became essential as the customer base using the machines broadened.

Then, too, providing the capability to access the central accounting systems in an online mode required further innovations in the back office. As usual, those banks with their eyes most attuned to the new opportunity made windfall gains as they implemented – for the first time – self-service offerings. In the UK, for example, there were more ATMs than branches by 1994, and in the USA, there were 400 or more machines for every million people in the country by the same time [4].

Not all innovations of the 1970s were technological, however. In 1974, Bangladesh was struck by a famine devastating the poor of that country. Professor Muhammad Yunus, then a professor of economics at the University of Chittagong, realised that loans of very small amounts might enable the poor to bootstrap themselves out of crushing poverty. Trying an

experiment, he lent less than US\$1 to each of 42 impoverished bamboo stool makers in a nearby village so they could buy raw materials. He was repaid – with interest – rapidly.

Microfinance – the breakthrough discovery that lending at the bottom of the pyramid was not only practical, but could be profitable – led to the creation of the Grameen Rural Bank soon after. The founding principle of Grameen Rural Bank is that credit availability should be made on the basis of the potential of a person, rather than assets held. In order to assess such potential, a group of five potential borrowers are assembled who provide a morally binding guarantee for loans issued to two of the members. Depending on repayment performance, loans are granted to two subsequent members and, eventually, the fifth member as well.

Grameen Rural Bank presently has more than 1000 branches, and a repayment rate in excess of 98%. The economic hardship the institution has relieved, though, is incalculable. In recognition of its pioneering efforts, Professor Yunus and the bank were awarded the Nobel Peace Prize in 2006, the first and only time a financial innovation has received arguably the world's most prestigious honour.

The eighties

Although technology was the key driver for innovation in banking throughout the 1960s and 1970s, it wasn't the only thing forcing change. Distribution changes were also afoot. A breakthrough innovation in the 1980s was the entry of US retailing giant Sears Roebucks to full-service banking using its extensive stores network. This departure from traditional bank-led distribution of products was not entirely successful at the time, and led to eventual divestiture of the financial services businesses at Sears Roebucks in 1994 [18]. Nonetheless, the innovation (coupling financial services with other, only marginally related business lines) was swiftly copied by many other retailers, often with significant success. As the chairman of the American Banking Association later remarked, the evolution of non-bank players in retail financial services was leaving banks in 'the unenviable position of trying to keep up with less-regulated, non-bank competitors, while also maintaining a much more expensive infrastructure' [19]. Though Sears were unsuccessful with their play for retail financial services, the model has subsequently been proved out by major retailers such as Wal-Mart and Tesco.

Meanwhile, banks were also broadening their distribution through a range of revolutionary self-service innovations. In the late 1980s the most significant of these was the development of telephone banking. This had first been tried unsuccessfully by Banc One Corp of Ohio in 1979, and pilots continued through to 1982 [20]. The adoption of phone banking was relatively speedy after that, something that institutions could readily have expected after their previous self-service success with ATMs. Naturally, institutions such as the innovative Sanwa Bank of Japan reaped windfall rewards from being first to market: one report was that 40% of deposits and 70% of card transactions were performed through self-service channels by 1984 [4].

The rise of telephone banking next led banks to question whether they could couple telephones and televisions together to create an interactive self-service channel for customers. At the time, the technology to do this existed in the form of Videotex, an early predecessor to dial up services that used the then-developing personal computer. Introduced in 1982, Minitel, the French version of the service, actually achieved relatively high levels of adoption, but in almost all other countries, the success of Videotex was extremely limited. This did not deter banks, who could see the potential of self-service after their initial experiments with ATMs and telephone banking. By 1985, at least 37 banks in the USA offered services using

the technology, but these never achieved critical mass. The vision of home-based financial self-service via personal computer had to wait more than a decade for the Internet.

The fact that Videotex was a commercial failure for most institutions does not detract from its status as a revolutionary innovation that drove banks to consider many of their operations in a fresh light. Customers with Videotex could quite reasonably demand access to their accounts in the small hours of the morning, a time when customer service centres would be closed, and even ATMs might be cycled down. Such behaviours were the result of the modal operation of the systems of the time: they could either be processing customer transactions (during the day), or carrying out various administrative tasks such as routine account maintenance and interest calculations (during the night). Most banks were simply not set up to enable both kinds of activities simultaneously. The introduction of Videotex forced institutions to think about the consequences of needing 24×7 access to systems.

Another incremental innovation of the late 1980s was the smartcard. The initial promise of the smartcard was that its embedded computer chip would be more reliable and secure than the venerable magnetic strip first pioneered by IBM more than 30 years before. The first mass use of these cards was their adoption by French authorities as a means of payment for public phones in 1983. It wasn't, however, until the early part of the millennium that the chip-based cards came into wide usage, when the chips were capable of storing not only data but applications as well.

Not all relevant innovation from this period was bank-led, however. Intuit, founded in 1983, was a small software company that sought to help households balance their chequebooks. Initially, Intuit tried to sell its new personal finance management software through banks, but quickly realised that it would have much greater success touching its customers directly. Neither was the personal finance software category unique to Intuit: a large number of competitive programs were quickly brought to market at the same time. Nonetheless, during the late 1980s and early 1990s Intuit's core product, Quicken, went from strength to strength, proving that there existed a segment of bank customers who wanted more control of their finances than that provided by paper statements.

The nineties

In 1992 Intuit added bank direct connect capabilities to Quicken, when it signed a deal with Visa to allow statement download directly into the product. Direct connect was an incremental innovation that enabled personal customers to replicate capabilities that business customers of banks had had for years. With it, they were able to reconcile their transactions electronically from home, something that wouldn't be available to the majority of customers until three years later with the rise of Internet banking.

The first precursor of modern Internet banking launched on 6 October 1995, when Presidential Savings Bank allowed customers to open new accounts over the Internet using a standard web browser. On 18 October, Security First Network Bank, launched by Kentucky-based Cardinal Bancshares organisation (backed by investments from Huntington Bancshares and Wachovia), opened its doors for business. Security First was the world's first virtual bank: it didn't have branches at all, and customers interacted with it via the Internet and call centres. During the bank's first two weeks of operation, it opened 750 accounts from customers in 32 states, more than three times the number that the founders achieved when they opened a traditional branch-based bank in Louisville a year earlier [21].

The software assets of the bank were held by a wholly owned technology development company. This was spun off as Security First Technologies, later S1 Corporation, when the banking portion of the business was acquired by the Royal Bank of Canada three years later in 1998. S1, continuing from then on as a pure technology vendor, began to deploy pioneering Internet banking solutions with other banks, operations it continues to this day.

Security First Network Bank was a traditional bank with revolutionary distribution. Although it was not especially profitable in the early years, the institution was a template subsequently copied by hugely successful direct plays such as ING Direct, launched in Canada two years later. As we've discussed elsewhere in this chapter, ING introduced its own business model innovations to enhance the direct model first pioneered by Security First.

Although it would need to wait some years for large-scale commercial success, 1998 was the year that X.com, the precursor to PayPal, was founded. PayPal was not the only alternative payment system to come into existence around this time, though it has proved to be the most successful. Other systems founded on the concept of secure digital payments, such as the much-hyped C2iT service from Citibank or BitPass from Western Union, have all subsequently closed through lack of adoption. When one examines the ostensible differences between the successful service and those that failed, the true innovation in PayPal's approach becomes clear: it coupled itself to one of the largest sources of online transactions around (online auctions), ensuring it had volume to sustain it, whilst its competitors did not. Subsequently successful online payment systems, such as Google's Checkout service, have followed the same strategy.

The closing years of the 1990s brought one final innovation based on the revolution of Internet banking: screen-scraping-based account aggregation. Technologists, recognising that if people could drive web browsers, it must be possible for machines to do so as well, quickly built automated routines that could log into Internet banking sites and extract account balances and transaction details. These were then merged together to provide a single view to customers, regardless of the number of different banking relationships they had. Yodlee Corporation, founded in 1999, was a pioneer in the industry.

Account aggregation was both an incremental innovation from a technology perspective (being based on the developing capabilities of Internet banking) and a disruptive one from the perspective of distribution. Bankers, fearful that they'd lose control of the customers they were only beginning to attract as a result of their own online efforts, spoke of 'disintermediation', whilst analysts were fulsome in their praise of the new opportunities available to those with a full picture of the 'wallet' of customers. Regulators and customers meanwhile began to worry about the legal consequences of disclosing their account information to a third party. In the end, adoption of account aggregation was slow, and had to wait for the second Internet boom surrounding the emergence of the (so-called) Web 2.0 to attract real consumer interest.

The present day

From 2000 onwards the pace of change, though not wholly driven by technological innovation, has continued. Better clearing systems have seen the emergence of decoupled debit, for example, a disruptive innovation if ever there was one. Account aggregation has led towards high-end personal finance management solutions that are entirely web-based and use the wisdom of crowds to exceed the best efforts of banks in the customer experience domain. And Internet banking has led naturally towards mobile phone banking, which will develop,

inevitably, towards phone-based payments based on emerging techniques like near-field communications.

With the pace of change picking up speed, there's little hope of providing any comprehensive account of significant financial services innovations that are beginning to emerge right now. Neither would it be a sensible effort: any attempt would be completely out-of-date almost the moment the words were written.

Nonetheless, the remainder of this book will examine, from time to time, some contemporary financial services innovations, as well as those appearing just over the horizon. The reader, it is hoped, will bear in mind that their retrospective knowledge of many of these will be superior to mine at the time of writing. But that's par for the course, I suppose, when you're writing a book on innovation. No matter what you say, it is invariably out-of-date the moment the manuscript lands with the publisher.

1.5 THE INNOVATION DOWNSIDE

The historical perspective we've just explored offers an interesting insight: few single innovations described remained unique for very long. Although many were breakthrough or revolutionary, very few have generated sustainable competitive advantages. The reason is that most banking innovations are sustaining in nature. Sustaining innovations tend to be extremely replicable in financial services.

When Barclays deployed the first ATM in London, it did not retain any persistent competitive advantage, even though the breakthrough of cash self-service changed the game for the rest of the industry. And when Security First Network Bank of the United States created the first Internet banking site – a revolutionary use of web technologies – it led other major banks into the market by a few short months at the most.

Being first with an innovation does, often, mean a commercial case can be made for a limited amount of risk taking. Bank of America, for example, invested millions to build its ERMA system, and was then able to leverage year-on-year cost decreases that left its competitors scrambling to catch up. But sustaining innovations in financial services wind up being a cost of doing business eventually, since they are replicable by competitors. There is a limited period only in which the innovation will provide any differentiation.

Whatever is introduced today, especially if it is a breakthrough or revolutionary, will almost certainly be tomorrow's legacy that must be supported into the future. Over the life of an innovation (which will likely be measured in years for smaller investments, or decades for larger ones), those costs can mount up. For a sustaining innovation, financial services companies are rarely well served by going far out on a limb. This is a tricky optimisation problem that will be covered later in this book: what is the *right* time to invest in something new? The earlier one enters, the larger the windfall gains that must be available to balance the cost of supporting the innovation throughout its life.

What is the rationale for delaying the implementation of sustaining innovations (a *fast follower* strategy as some call it)? As something new becomes commoditised, its costs of implementation decline. This has the effect of reducing the barriers to entry for any particular new idea for following institutions. The later the institution comes to a particular innovation, the less its costs and risks.

This leads us to three quantitative questions which are important in the innovation selection process. Firstly, how large are the development and operating costs likely to be? Secondly, will the windfall gains for early entry plus the value of any residual advantages be enough to cover these costs? And finally, what is the risk associated with each of these two questions?

These arguments are not necessarily true, however, for disruptive innovations. A disruptive innovation becomes important, from a revenue and strategic perspective, in a much longer timeframe than a sustaining one. If the investment to deliver such an innovation can be controlled appropriately, a rational argument can often be made for early (or even first) entry with a particular innovation.

Making such arguments is one of the primary reasons that coupling a structured means of analysing the future with the innovation process is so very important. Such an analysis allows the innovator to determine whether or not a particular innovation should be framed in sustaining or disruptive modes. This, in turn, provides an important mechanism which guides the innovation selection process.

1.6 AN OVERVIEW OF THE FUTUREPROOFING PROCESS

And now, at last, we come to an overview of the business processes with which the rest of this book concerns itself. The overall process is illustrated in Figure 1.1, which shows the four main things a successful innovation function does as part of its ordinary day-to-day activity. The first, futurecasting, we'll come back to in a moment. The second, ideation, is about inventing new concepts that can reasonably be converted into new activity. Innovation is the process of converting these inventions into a funded reality, and execution, the final stage, is the set of things one does to make sure the idea finally gets into the hands of users and customers.

Figure 1.1 is quite different from the usual representation of innovation processes, which invariably uses stage gates to guide investment decisions. The fundamental concept of stage gates is that there is a 'funnel' of activity. Ideas enter the top of the funnel – often in large numbers – and proceed through various decision points. At each decision point a significant number of ideas are rejected, imposing a Darwinian 'survival of the fittest' approach to selection. At the other end of the funnel, a small number of ideas will 'graduate' to active development, with a smaller number still going on to commercial success (or, for those innovations which don't face customers, internal implementation).

Stage gates do not, by themselves, deal with some fundamental issues facing bank innovators, though. Are the ideas coming into the innovation funnel the right ones? What steps

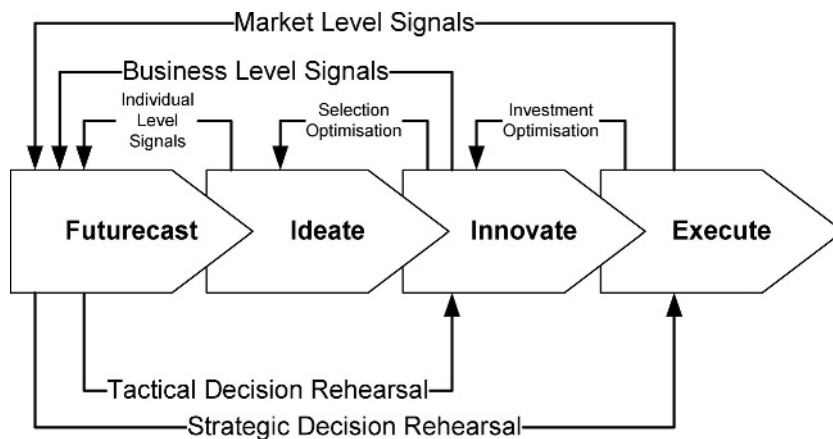


Figure 1.1 An overview of the futureproofing process

need to be taken to rehearse the decisions of senior executives before an innovation decision is made? Are the innovations that *are* being implemented the ones that the bank needs to ensure it has a future? Something beyond a robust stage-gate process is needed to answer questions such as these.

Futureproofing is probably more correctly referred to as a *phase-gate* process: it surrounds a traditional stage gate with additional mechanisms which help to ensure that what goes in (and what comes out) is connected to the strategic reality of a particular institution. And whilst stage gates send ideas through a pipeline in only one direction – towards implementation – a *phase-gate* process sends messages up and down the steps in order to optimise the predictability and returns of the whole process.

Let's now look at each of the phases shown in Figure 1.1.

The futurecasting phase

Futureproofing is a phase-gate process that adds structured mechanisms for dealing with strategic questions affecting the institution's future. This is relatively new territory for most bank innovation programmes. Many practitioners imagine that as the likely shape of the future is not amenable to accurate analysis, there is little value in devoting much attention to it at all.

But whilst one cannot predict the future with *certainty*, it is definitely possible to plan for various likely outcomes given a particular strategic scenario. For most institutions, such planning is de rigueur in any case. Economic questions, such as the likely shape of the economy in the medium term, are routinely forecast in the ordinary course of business. This is possible because the theoretical principles of economic forces are well understood.

This is increasingly the case for innovation processes as well, and the last decade has been extremely busy for innovation theorists. The consequence is that we have working models of the way that companies and their people react to innovation, and these can be used to make reasonable guesses about what is likely to happen. Such theory will be covered in the next chapter of this book. We'll also be looking at other tools – such as scenario planning – that can help guide our thinking about the likely shape of the future.

Scenario planning is a key part of the futurecast. One seeks to create several stories about possible futures. They guide decision-maker thinking, helping to inform the critical judgements that are needed to make important decisions. A scenario doesn't necessarily ascribe a particular likelihood to any particular future, but it does draw out the fundamental forces that ought to be considered before committing an organisation to a particular direction.

This mix of good theory leading to great stories about the future is the principal tool of the futurecast phase, but how does having good futurecasts help the innovation process?

Imagine that peer-to-peer lending becomes significantly successful – successful enough, in fact, that it starts to cannibalise the best customers of traditional banks. The mechanics involved in the set of circumstances which could cause that to happen are developed by the futureproofing process as part of the futurecast stage. This analysis leads the team to develop stories about likely consequences and possible responses, which are given to decision-makers to inform their thinking. They are plausible, and based on the facts that are known *now*.

Perhaps one of the stories is that P2P cracks the mortgage market and swiftly disrupts what remaining margins there are for banks in this segment of the business. A possible response is to spin off a new business to take on the entrant that utilises the superior money buying power of the institution to fend off the attack, leading to market-wide structural changes in the way that mortgages are sold and administered.

Providing such stories about the future informs the thinking of decision-makers. Then, when the innovation team presents a more realistic innovation for go/no-go approval, the issues and constraints will be well understood by those involved. The futureproofing process enables senior executives to rehearse their decision without consequences before they are actually called upon to *make* a decision.

But the benefits do not stop there. Stage-gate processes must be fed continuously with ideas, and in order to get to predictability at the other end, the gating processes need to be rigorous. The overhead in dealing with volumes of ideas is considerable.

When you use futurecasts as an input to the ideation process, the result is that you have a way to connect invention with the strategic questions facing the bank. The stories and scenarios developed by the futurecasting step deal directly with the big imperatives lying ahead. Ideas coming into the pipeline inspired by these analyses are less random, and much more likely to be helpfully applied to the big challenges.

Even during the innovation stage, when ideas are turned into reality, the futurecast is a helpful diagnostic tool. Does the implementation decision being taken now close down options that might be important in the future? Does eliminating an idea (see 'drowning the puppy' later in this book) have strategic consequences that might affect the future of the institution? Indeed, does failing to eliminate what appears to be a great idea from a stage-gate perspective lead to dead-ends later?

The ideation phase

Referring again to Figure 1.1, the second part of futureproofing is ideation. Ideation is the set of things that innovators do to marshal collective creativity around specific things that might make an institution better.

There are two main parts of the ideation phase. The first is exploratory: having a large volume of suggestions coming in from customers and staff is an excellent means of spotting opportunities for futurecasts. These individual-level signals, if they are constantly about a particular issue, are often indicative that the dynamics affecting the business are either undergoing change, or have already altered. Noticing them (and starting to write stories about consequences) results in a new wave of ideation activities that can find ways of remediating the change.

The second part of ideation, though, is collating the collective suggestions of the organisation, finding a way to prioritise and classify them, and feeding them through the futureproofing process. The great thing about ideation, as I said earlier in this chapter, is that it is a great deal of fun. Dreaming up new ways to do things is intensely creative and far less complicated than actually making things happen. Consequently, the ideation step is an excellent means of gaining broad support for the innovation agenda. Anyone can participate, and with that participation comes support, without which innovation can be like swimming uphill in treacle.

The innovation phase

Ideation naturally leads to innovation, the third part of the futureproofing process. In innovation, one seeks to create something real from the ideas that were previously prioritised during the ideation phase. Real, in this context, means everything that has to be done to get an idea from the point that it seems interesting enough to explore further to the go/no-go decision that leads to a funding decision. The set of things that must be done, practically speaking,

varies from innovation to innovation. But in financial institutions, the innovation step is most likely to involve a great deal of influencing and salesmanship. Most of the time, innovation programmes themselves simply don't have control of sufficient capital to launch something in their own right. They must win the money from benefactors in the business in order to move forwards. Innovation, then, is about creating the tools necessary to convince stakeholders that whatever-it-is is a superior investment choice compared to whatever else is on the table.

The mechanism through which these tools get built necessarily involves a great deal of interaction with executives outside the innovation team. In the process of this interaction, another set of futureproofing signals surfaces: those that have to do with the responses of particular business lines to internal and external stimuli. For example, a particular innovation may be rejected on the basis that it is 'distracting' from the core business. Such rejection is informative. If the innovation has to do with an expanding business, it may be that management does not have its eye on the disruptive threats to future revenues that will inevitably emerge in the future. An appropriate futurecast would then be one which dealt with this potential for a small player to take away the future.

On the other hand, if the business line was struggling to meet its growth targets, it might indicate that management, responding rationally to a market event (such as the entry of a competitor with a cost advantage), might determine that the best strategy is focusing on protection of the existing customer base and improving margins, leaving the future hanging in the balance to ensure the present. In that case, a relevant futurecast might be one that explored the opportunities for growth inherent in a *downmarket* rather than an upmarket move.

The specific response to signals is not as important as capturing them as evidence to support the futurecasting process. The thought leadership inherent in the futurecast can then be used to improve the quality of thinking in the institution and rehearse future decisions before they must be made. In the innovation step, the future decisions of most importance are essentially tactical. Which innovations should one pursue rigorously, and which should be eliminated? Not all innovations, when subject to detailed analysis, are equally good. Selecting those most likely to be successful requires an understanding not only of the present situational state of the business, but its likely state in the future as well. The outputs of the futurecast stage enable the rehearsal of these tactical decisions as well as the ones that senior executives must make later in the process.

During the innovation phase, the quality of ideas that have arrived from the ideation step will be subject to quite deep analysis by the innovation team. The purpose of this analysis is to answer three key questions. Firstly, is the timing of the innovation right, given market and other conditions at present? Secondly, is it possible to actually make whatever-it-is work? And finally, what will it mean if the institution does? (I always paraphrase these questions as 'Can we? Should we? When?') These three questions inevitably have to be answered in the lead up to a funding decision anyway, but are also useful as a means of optimising the mix of things arriving from the ideation phase. This optimisation is a key part of balancing the innovation portfolio – a necessary step leading to predictable returns on innovation investments.

One final point on the innovation phase is that it does not often involve *actually* delivering things that can be used: prototypes and business cases may be created, certainly, but these will not be the ones that would be relied on once the actual funding has been committed to the innovation. The whole point is to create whatever needs to be done to *win funding*. It is better to do this with throwaway artefacts than fully scaled out solutions. The former are inexpensive to produce, whilst the latter aren't. Since the only way to de-risk the overall return on innovation

is to have a large and diversified set of innovation investments, delaying as much of the spend as possible is a functional necessity.

The execution phase

And finally, we come to the execution phase. This commences immediately after funding is won, which is the ultimate point of the innovation stage. At this point, decision-makers have been convinced the innovation represents a superior investment choice compared with other opportunities they have had presented to them. They've had their questions satisfied about the opportunity, the technical capabilities of the team and the firm to execute, and are sure that *now* is the right time to execute.

But what happens next is determined by the kind of innovation being considered. Earlier in this chapter, I explained that innovations are either sustaining or disruptive. The former are well suited to development by core business lines using their ordinary delivery processes. But the latter are much more difficult. In most cases, existing businesses are incapable of disrupting themselves. This is because the rational choice is generally doing things in a way integrated with current business being done. But by so doing, the cost and value decisions of the main business are imposed on the disruptive one, *whether that makes sense or not*. For something truly disruptive, this usually results in failure.

Here, then, is one of the most important reasons for using futurecasts extensively: a senior manager with financial responsibilities is being asked to invest in an innovation which will – eventually – cannibalise his or her own business. Furthermore, they are asked to do so without imposing their own controls. That is a strategic decision of the first order, and most executives will simply not make such a decision without extensive preparation. Such preparation must necessarily include education about the possible consequences of *not* investing in the disruptive innovation.

The actual decision to commence investment is the first of many activities that characterise the execution phase. Along the way to launch, an innovation is subject to many compromises. Some of these are reasonable and others not. For example, a security team well versed in the reality of large systems may impose constraints on an emerging innovation (such as special infrastructures, expensive authentication devices, or comprehensive penetration tests) which make the innovation an expensive exercise. Sometimes these compromises lead to innovation trauma, a situation where no further innovation is possible. Everyone concludes that the pain of innovation is much greater than any benefits that arise from it.

It is the role of executives (with the coaching of the innovation team inevitably) to help those assembling the parts of the innovation to understand they are doing something genuinely *new*. Old processes and procedures might work in the traditional context, but may not be appropriate now. This is especially true for disruptive innovations, but can also be the case for sustaining ones as well. Without careful nurturing during the execution phase, an innovation can simply fail under the weight of previous (even unrelated) legacy.

Why does this occur? As I will discuss later in this book, it is primarily the result of an authority asymmetry. In most institutions, there is a group of individuals who are empowered to say 'no'. The bank will be happy to let them do so, because the certain present is far less risky than the unpredictable future. Saying 'no' is the safe option.

There are far fewer people, however, who are actually authorised to say 'yes'. These individuals are the ones that give an innovation the life-support it needs during its birth. The executive who initially approved the investment in the first place is likely to be one, but

he or she will not have the bandwidth to supervise every detail of implementation. Finding and engaging people on the right side of the authority asymmetry is a critical aspect of the execution phase.

Authority asymmetries are a natural consequence of all organisations, and do not reflect poorly on any particular bank. In fact, it is probably true to say that the greater the asymmetry, the more stable the bank! The natural response of any manager who has to meet a particular number in a profit and loss statement is to shoehorn anything new into a controlled, certain framework.

The strategic decision that resulted in the funding decision is not the only decision that futurecasts help rehearse. During implementation, a large number of choices must be made which have significant effects on the future success of a particular innovation, and potentially, on future ones as well. For example, a decision to implement an innovative new payments system based on the existing legacy infrastructure in place may *seem* like an architecturally sound decision. It may, even, be less expensive initially. But such a decision commits one to the legacy cost base of the entrenched infrastructure. This may be disastrous if competitors have a cost advantage resulting from, say, a brand new implementation. In this case, the futurecast enables decision-makers to see legacy cost drivers are one of the key strategic inputs they have to control in order to ensure success. IT, and other powerful interest groups, would have to be managed in a (for them) unnatural direction.

In the end, if everything has gone well, the innovation has been built into something internal users or customers can get their hands on. The response of the market (or internal user base) provides interesting signals that are input to the futurecasting process. Is the innovation interesting to the press or competitors? Do customers respond immediately, or is the take up more gradual? What is the analyst response?

These are questions which are market level in nature. They provide coarse-grained material that helps to optimise the set of questions that ought to be asked about the future. In this respect, an implemented innovation (even if it is ultimately a failure) may be seen as a probe that helps reveal interesting information about the future, and is therefore a hugely important input to determining the likely set of innovations that may be needed later.

In 2007, ING announced its Our Virtual Holland experiment – a major exploration of the use of virtual worlds. According to the company's website on the topic, 'ING is taking the virtual-world phenomenon seriously and expects that part of its future Internet activities will become three dimensional. . . . ING seeks to gain a better understanding of improvements it can make to products and services by assessing valuable feedback from inhabitants who will experience ING in Second Life' [22]. By 2008, it had closed down the innovation, again remarking that it had chosen to do so because of its need to focus on merger activity it was undertaking at the time, even though 'virtual worlds continue to be an interesting and for ING a relevant development'.

That the experiment lasted only a year does not reduce its value as a lens illuminating the future, both for ING and any other institution that might be interested. Our Virtual Holland was the subject of considerable press and analyst activity, a signal of a potential market trend. Customer adoption was relatively low, but a number of substantial companies joined ING on their virtual island. These facts are indicative of an innovation released before its market was appropriately developed on the one hand, or of a technology not good enough to fulfil customer functions on the other. But regardless of the interpretation, ING has useful data that will likely make its further guesses about the future of online financial services superior to those of its competitors.

Optimisations in futureproofing

Before we leave this broad brush overview of the futureproofing process, a bit more expansion on a comment I made at the beginning of this section is in order. Whilst an ordinary stage-gate process moves ideas through to implementation in one direction only, the phase-gate process described here seeks to optimise that flow by adjusting the preceding stage using the experiences of the latter one.

The input to the execution phase is the set of things that have been funded through the activities of the innovation team. Throughout the development of the innovation, however, much will be learnt about the capabilities of the institution and its reaction to various situations, both from an internal and an external perspective. This learning provides essential *optimisations* for the innovation phase, and will almost certainly change the way the innovation team selects and pursues specific innovations.

If an innovation enters the execution phase, for example, and it is then discovered that basic technological requirements cannot be easily met within schedule and cost constraints, the effect on innovations in earlier phases of the process can be significant. Some innovations which may have been selected as incremental on the first innovation will have to be dropped. Others, which may have been unattractive given simpler opportunities available, will now become priorities. The innovation team will naturally balance its choice of what to put forward and what to ignore based on the reality of what can and cannot be done. It is usually the case that all ramifications of a particular innovation will not be understood until very late in the process. Optimising the set of investment choices presented to stakeholders with this new information is a sensible and reliable way of ensuring the overall innovation portfolio remains stable.

Similarly, input to the innovation phase is the set of ideas that have passed through whatever collective creation process has been implemented. Many of these ideas will be excellent when standing independently, but difficult or impossible to get through to funding. In order to guarantee a stable innovation portfolio, however, the innovation team will necessarily need a good flow of ideas that it can reasonably hope to develop. Over time, the team will discover additional knowledge and experiences that let it determine more accurately the chances of getting a particular concept through the process. It is therefore able to adjust selection criteria from the ideation process to ensure that this flow is adequate consistently.

1.7 WHERE TO GO NOW

Much of the rest of this book covers futureproofing processes in detail. But before we examine those, the next chapter looks at a range of innovation theories and models in the context of banking. Before you nod off to sleep at the thought of that, there is a very good reason for examining these now: theories give us a way to explain the reasons particular innovations may or may not be successful. Coupled with appropriate story telling resulting from the futurecasting tools we'll be looking at later, theory gives us a means for being robust in our expectations of the future.

Chapter 3 takes theory and examines specific innovation issues in the banking context. What sorts of innovations are available to banks? What innovation strategy is appropriate? And, in general, how do bank innovation programmes develop?

Theory aside, the main body of this book is about processes that innovation teams and their sponsors can use to get some predictability out of their innovation investments. There is a

chapter on each of the phases I talked about a moment ago that provides this information in detail. You'll find this material starting at Chapter 4.

Whilst I'll mainly cover the innovation processes that lead to futureproofing, a process in isolation is just paper and abstraction. Such things have value if you are a consultant, but making things real requires people. And people, in the context of futureproofing, are just about the most critical resource there is. Time and time again, it has been shown that the difference between a successful innovation programme and the rest is the quality of the leader in charge of the programme. In Chapter 8 you'll find material that explains the roles and responsibilities of the innovation leader.

And a leader is nothing without a great team. Innovators, being creative non-conforming types (if they are any good at all, that is), are a special breed. There are many different archetypes, and it is functionally necessary to have them all in the team to ensure that the overall innovation function is a successful one. Chapter 9 has a lot to say about innovation teams.

Chapter 10 explores the systems and processes that surround an innovation team. Inevitably, the values and processes of an institution have a very big impact on the way that innovation can develop. Read this chapter if you're interested in reward systems, measurement and metrics, and the other things an innovation programme must do to fit in with the rest of the bank.

And finally, Chapter 11 presents 39 steps you can use to build an innovation team, and rounds out the rest with some final recommendations and conclusions that haven't been presented elsewhere.

I hope you enjoy this book, the result of the very hard work of many people. Many practitioners of innovation, both in financial services and from other industries, have been involved. They have been kind enough to share both their successes and failures in these pages, and I owe them a debt of gratitude. But there is another group whose work I have borrowed: the generations of scientists and academics who have been studying the problem of innovation for decades. Without their rigorous observation and theorising, the processes we'll be using in futureproofing would be more art than science. It is to their work that we will turn in the next chapter.