CONTENTS

Preface xi
How It Came About xii
Format xvi
Structure xviii
Terminology xxi
Disclaimer xxii

1. The Purpose of Forecasting 1

What IS a forecast? 2
Forecasts are not the same as budgets or plans 4
Different kinds of forecast 6
All operational forecasting is based on the same principles 8
Why operational forecasting is important 10
The quality of forecasts matters...a lot 12
Forecasting is a (very) good investment 14
Forecasting is not compulsory 16

2. Understanding Demand 19

Data Series are different - and it matters to forecasters 20
Forecasting is not easy 22
It is not possible to predict the future 24
It is not easy to spot noise 26
Types of Signal 28
The buckets used to measure demand changes the characteristics of the data 30
Characteristics of operational demand data 32

3. Forecasting Methods 35

The forecasting challenge 36
Three ways to produce forecasts 38
Forecasting by extrapolating from historic time series 40
Forecasting using causal models 42
Forecasting using judgement 44
Three types of bias in judgemental forecasting 46
The better the fit to history - (sometimes) the worse the forecast method 48
How to avoid overfitting 50
Sophisticated methods are (often) no better than simple ones 52

The methods used in different
software packages are not unique 54
The methods used in different
software packages are not always correct 56
Autoselect algorithms do not
always select the best forecasting method 58
Machine learning technology will not 60
eliminate the need for human input into forecasting 60
Combining forecasts produces better results 62

4. Understanding Forecast Performance 65

The measurement challenge 66
Conventional error metrics 68
Period errors are not meaningful 70
The denominator is not important 72
Value Added - the ultimate measure of forecast quality 74
Forecast Value Added (FVA) - the measure of the
effectiveness of steps in the forecast process 76
Some demand series are more forecastable than others 78
Comparing to the naïve forecast error adjusts for
forecastability 80

How to quantify the level of avoidable errors - the limits of forecastability 82
Errors always increase over time 84
Allow for different lags when measuring the value added by forecasts 86
Bias and variation - two types of error 88
The different business impact of bias and variation 90
The causes of bias 92
The causes of variation 94
Conventional error metrics treat bias and variation in an inconsistent way 96
Forecasts do not need to be equally accurate across the horizon. 98
Sometimes a forecast is not a forecast 100
The best forecasts do not always have the lowest errors 102
It is not possible to compare forecast error across levels 104
Error levels cannot be compared and targeted 106
It is not impossible to benchmark forecast quality 108
It is not possible to measure the quality of long term forecasts 110
Forecastability does not describe how easy it is to forecast demand 112

5. Managing Forecast Performance 115

Forecast improvement is based on understanding error 116
Forecasting is not mandatory 118
The objective should not be to increase forecast accuracy 120
Forecasting does not always add value 122
Few things are unforecastable 124
Good forecasts don't always 'look right' 126
Reacting to noise makes things worse 128
Focussing on the biggest errors is not the best way to improve forecast quality 130
You can only measure and manage forecast errors trends 132
Improving short term forecasts is not always worthwhile 134
Judgemental forecasts are not always bad or alwaysgood 136
Consensus forecasting is not (necessarily) 'best practice' 138
Experts are (often) no better at forecasting than non-experts 140
The level of aggregation affects the ease of forecasting 142
When to use 'top down' and 'bottom up' forecasting methods 144
To disaggregate or aggregate? 146
Why we clean history and when to do it 148
Judgement and when to use it 150
More frequent re-forecasting is not (necessarily) better 152
Segment the product portfolio to get better results 154

Standard deviation is not the way to measure volatility 156
Better forecasts do not automatically generate business benefit 158
A business should always have 'one set of forecast numbers' 160
The forecast numbers do not always need to be the same 162
It doesn't matter who forecasters report to - providing... 164

Terminology 166
Calculating Value Added Metrics 172
Recommended Reading 174
Acknowledgements 176

PREFACE

This is a guidebook about short term Operational Forecasting – the sort that is done to determine how much product you need to source or how many people you need to draft in to meet customer demand.

My aim is to produce something that provides a useful introduction to operational forecasting for both practitioners and their bosses by filling in the gap that lies between a naïve 'common sense' intuitive grasp of the topic and the complex technicalities of mathematical forecasting techniques.

I think my background as a self-taught business orientated forecasting nerd with limited mathematical expertise who is not scared to tell it the way I see it puts me in a good position to fill this gap.
I have designed the book to be simple but not simplistic, using short and to the point learning points supported by clear graphics. It is technically sound and practical.

I hope it will create a common language to help people talk intelligently about forecasting and help stop people doing dumb stuff – which is where most of the potential for improvement lies. I hope it will help people design good forecast processes and informed software purchasing decisions. And I hope that it will in some small way help people realize that forecasting is important and that investing in people as well as software will generate enormous benefits for many businesses.

HOW IT CAME ABOUT

This is a book I had never intended to write but it was one that I felt needed to be written.

Just before I left Unilever after a career in Finance I designed and led a change initiative that aimed to replace traditional bureaucratic budgeting practices with more agile planning processes, based on the Beyond Budgeting model, that would enable businesses to better adapt to the world.

As part of this effort I chose to focus on improving forecasting, as I reckoned that only when people had a future orientated process that they could rely on would they let go of budgeting – however much they hated it.

As I did more and more research on this topic I came to realize that no-one in the finance community – inside or outside of Unilever – really understood business forecasting. And I included myself in this number despite having been responsible for forecasting for nearly 20 years.

What I learned on this journey eventually led to me writing my first book 'Future Ready', which is still the only book on the market covering this important topic.

In this work on business forecasting I drew heavily on the insights of academic and practitioners working to improve operational forecasting in the supply chain.

It shouldn't be surprising that so much energy and expertise is devoted to the challenge of forecasting demand since it provides the foundation upon which supply planning sits.

Typically, any business selling from stock will need to produce many thousands of forecasts very frequently, and the results have a direct and immediate bearing on the business. Get it wrong and you can either

end up with too much stock that costs you a lot of money or too little which leads to disappointed customers and losing what is left of your fortune. And it is possible to do both at the same time in different parts of the product portfolio!

I came to develop an enormous respect for the expertise and professionalism of people working in this field, but before long I had an experience that was just as shocking and surprising to me as the realization that I had when I discovered how little I knew about business forecasting even after 20 years of practice. None of these really smart people could answer what I thought was a pretty obvious question: 'how do I know how good my forecast is?' The existence of this gaping hole in our knowledge intrigued and frustrated me in equal measure and so I set out to find the answer for myself.

My quest for an answer to this question ultimately led to me co-founding a software business, CatchBull, and over time I became an accidental expert on the topic, with a handful of academic publications to my name.

Naively, I assumed that selling this product would be easy, because the need and the business benefits were so obvious. Instead when I spoke to business people I found a blank wall of incomprehension. And I got a similar reaction from many academics and software suppliers I encountered.

What was going on?

What I slowly came to realize is that most people in the supply chain simply do not appreciate the nature of the challenge of forecasting demand at the very granular (detailed) level at which it takes place, and the scale of the benefits if you can get it right. My problem was that because I had immersed myself in the problem for years I had lost touch with my old ignorant self and I assumed that the professionals I spoke to would appreciate the nature of the challenge – and the opportunity – in the same way as me, which they clearly did not.

This was forcibly brought home to me in a conversation I had with the leader of the demand management process for a major multinational.

The conversation went something like this:

> Me: '…of course it simply isn't possible to forecast the future perfectly'
> Head of Demand Planning 'Why not?'
> Me: '……..' (stunned silence)

The reason why I had been rendered uncharacteristically mute was that it suddenly became clear to me that he and I had a completely different view of the world – of the reality of trying to anticipate future demand. And that I didn't know how to bridge the gap.

I believed – I would say I knew – that because the world is volatile, uncertain, complex and often ambiguous ('VUCA'), anticipating the future is difficult, and that there is an inherent limit to how accurate any forecast can be. As a result, I believe the only way to proceed is to do the best job you can of predicting the future but then, with humility, measure how successful you have been, learn and adjust your methods accordingly.

If, on the other hand, like my friend, you think the world is perfectly predictable, you must believe that there is a clockwork like mechanism hidden from our sight that governs what happens in the world. Forecasting therefore simply involves trying to discover and replicate the working of this cosmic machine where free will and chance play no part. If that is what you believe then producing good forecasts is purely a matter of building a sophisticated mathematical model, and errors are a purely a manifestation of our ignorance.

But even supply chain leaders who acknowledge that the clockwork model of the world is defective are not immune to the allure of 'clever' software, because you are spared the task of finding and retaining clever people. Instead you can employ relatively unskilled or inexperienced people – like graduate recruits – to tend the machine.

Problem solved.

Except it isn't. You don't have to take my word for it.

REL, a consultancy that specialises in working capital reduction, estimate that across the top 1000 US companies **$423 billion** is tied up in excess inventory (that's over $400 million per company on average). The number for European companies is E350 billion. And it's not got better over the decade that REL have been running the numbers…if anything, the recent trend is deteriorating.

A whole host of 'reasons' are given for this. 'Financing is historically cheap, so there is no economic incentive to reduce working capital', for instance. But if that's true how come (according to Gartner) companies worldwide are spending $8.3 billion a year on supply chain software? This doesn't square with the story that they are making a rational choice not to make improvements.

My hunch is different.

Producing forecasts and using them to drive supply chain efficiencies is hard work and requires skill. And for many people sophisticated software promises to make the problem go away…and no-one selling this software is going to tell you any different.

I am not anti-software at all – after all I own a software company! It's just that to me, many companies who invest heavily in 'clever' software are doing the equivalent of buying a car without being able to drive, not knowing where they are heading and why, what terrain they will face or how to plot their position and navigate a course. If you are in this position, buying a car with a 'better' engine is not going to get you closer to your goal.

But most smaller businesses face the challenge of forecasting without sophisticated software. Indeed, the most widely used forecasting model in every type of company sits between your ears. It is called judgment. And because we all use this model every moment of the day to make forecasts there is a tendency for them to think that forecasting is easy – little more than common sense.

But the idea that forecasting is just 'common sense' is just as just as dangerous as 'maths will solve everything'.

FORMAT

Because this book is aimed at a broad audience I have tried to make what can be a complex and dry subject accessible and easily digestible. The content is arranged as a series of bite sized 'lessons' supported by simple graphics to illustrate and reinforce the point being made.

There are two types of lessons.

There are positive lessons: concepts and ideas that you can rely on picked out in blue type like this.

But the territory of forecasting is host to a rich population of myths and misinformation. **So, there are also 'negative' lessons about what not to do, and these are shown in a red type**.

I have given these 'don't do' lessons equal billing…after all it is usually quicker and easier to stop doing stupid stuff than it is to learn how to do clever stuff.

Also, there is no getting away from the fact that forecasting in business is largely a mathematical exercise I have tried to avoid using formulae in the book. Many readers of this book will not be forecasting practitioners and so will not need to understand the mathematics but even forecasters would do a better job if they better understood the problem they were trying to solve before hitting the keyboard. In places, to make the topic accessible I have deliberately oversimplified complex ideas in a way that will make academic forecasters cringe. These are indicated as follows:

For those who need to explore these topics more deeply there are many excellent technical books on forecasting, some of which I have listed in the appendix, along with some of the more basic formulae you might need.

Elsewhere, where I feel a level of technical detail may be necessary for some readers, I have issued a 'nerd alert', which means unless you are one you can skip this bit for the time being.

Finally, in the interest of simplicity I have based the content on the scenario of a 'sell from stock' business. I recognize that service businesses also create operational forecasts, but the same principles apply so learnings based on a stock scenario can almost always be simply applied.

STRUCTURE

I have structured the book around five key themes.

Most of the books written on forecasting in business focus on forecasting techniques. For the reasons I have explained I will do not more than give an overview of the various approaches to producing forecasts – forecasting models – leaving the technicalities to people better qualified than I am. I have tried to be agnostic on the issue of technique. My favourite approach is simply the one that works best. My aim it to give the reader enough background to make informed decisions about the best approach to use – simple or sophisticated, judgement or statistical.

Technique is clearly central to the forecasting challenge, so forecasting methods comprises one of the five themes.

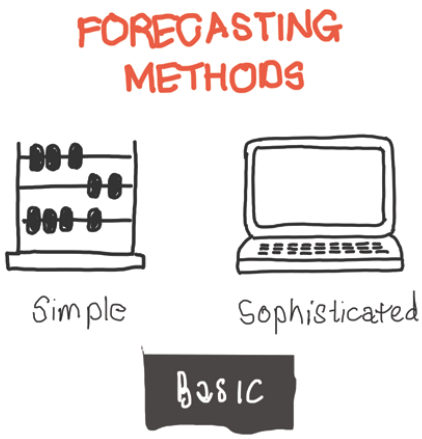

However, I think the obsessive focus on this to the exception of everything else has been unhealthy for the practice of forecasting. The other four themes will, I hope, help redress this balance.

First – and most important for anyone with a stake in operational forecasting – there is the question of purpose.

Forecasting is not an end in itself, its purpose is to help businesses make

better decisions. Only by understanding this purpose can we make the right choices. And we can only judge whether we have succeeded or not by how well it has fulfilled this purpose. In businesses that sell physical products operational forecasts are used to make replenishment decisions and to help determine how much stock to hold. In service businesses, there are close analogues to orders and stock (e.g. the number of staff needed to service forecast demand).

Success is therefore providing supply planners with the information they need to give customers the desired level of service with the lowest possible level of inventory. This, as we will discover, is not necessarily the same thing as 'high forecast accuracy.'

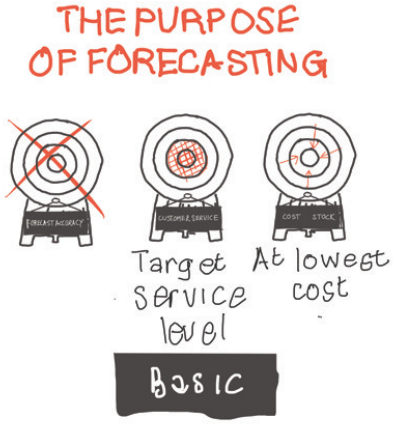

Next there is the question of understanding the behaviour of the variable you are forecasting. The forecasting of items with stable demand needs to be tackled in a different way to those that have a complex or volatile demand pattern, for instance. You need to make sure that you are properly equipped to cope with the forecasting 'weather conditions'.

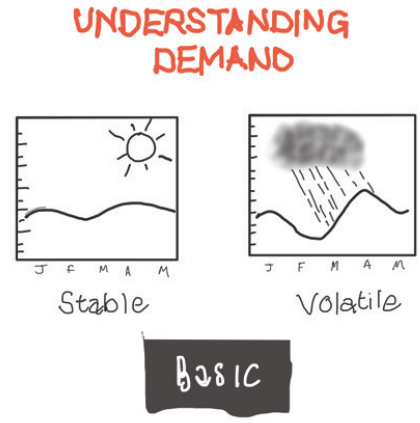

These first three themes provide a basic grounding for anyone associated with forecasting.

The final two are aimed at practitioners wanting to improve their performance, and so are inevitably slightly more technical.

Measuring and interpreting forecast error – i.e. understanding forecast performance – is the fourth theme. Effective measurement is necessary to be able to identify the best forecasting method and understanding where and why performance falls short of the practical optimum is a prerequisite for improving performance. Critically, how forecast performance is measured must relate to the practical purpose it serves in the business.

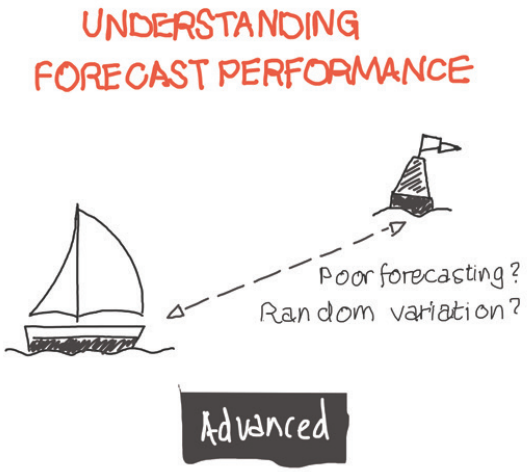

Finally, like any process, there is always scope for improving forecasts… unnecessary forecast error is waste that should be eliminated wherever possible. This final theme tackles some of the challenges facing any forecaster in the field and provides guidance on how to deal with them.

TERMINOLOGY

I have tried to avoid jargon wherever possible, but some 'buzzwords' are a helpful 'shorthand' that prevent the language getting too convoluted.

The 'specialized' terms that you will come across most frequently are:

Variable: The thing being forecast. Usually assumed to be sales.
Bucket: A unit of time used to measure of actual or forecast values, e.g. a week or month
Demand Series: A sequence of actual sales values.[1]

See the appendix for a list of more technical terms.

1 Strictly speaking sales may not be the same as demand – if, for example, there is not sufficient capacity, or if you run out of stock. In the interest of brevity the issue of "constrained demand" is not addressed in this book, but forecast practitioners should always bear it in mind.

DISCLAIMER

In writing this book I am sharing what I have learned, sometimes painfully, in my journey from naïve ignorance to a place where I can claim to be knowledgeable about forecasting in general and an expert on the subject of measuring forecasting performance.

The content of this book is based on academic findings, generally accepted 'good practice' and my own research, but some of it is my opinion based on my personal experience observing and talking to forecasters in the field. To the extent that I cannot 'prove' some of my assertions I beg the forgiveness of my friends in the academic community who have to operate to higher evidential standards.

Most of what I describe is what the best forecasting practitioners already know and do – even if they don't always realise it…with perhaps one important exception.

Few books on forecasting address the most fundamental question: what is the purpose of forecasting? This is important because if we don't know what we are trying to achieve we can't know whether we have succeeded.

For me the answer is clear: to help make better decisions. That sounds simple and straightforward, but it begs more questions: 'what decisions?' And better than what?
Again, the answer seems to me to be obvious:

- Operational forecasts are created in order to help business make decisions about what to order from manufacturers or suppliers and what they need to hold in stock[2].

- Good forecasts have better business outcomes than simply using actual demand to decide what to order and what to hold in reserve.

- A better business outcome is the total benefit of having lower stocks and better customer service by forecasting, less the additional cost of running the process. This is how forecasting adds value.

Consequently, the picture of operational forecasting painted in this book – particularly the final two sections on measurement and management – is viewed through this lens, because the tried and trusted ways of doing these things often don't work very well and so cannot be trusted. To duck this issue by adopting a more conventional stance would be short changing the reader.

My excuse for the absence of academic validation in some places is that given the state of ignorance about the craft of forecasting – even amongst people who have the word in their job title – it is better to have some guidance than none at all. I also think that my background in business gives me a valuable perspective on the practice of forecasting that isn't available from any other source.

I'm striving to provide something that people find useful, not an authoritative textbook, and I am happy to be judged on that basis.

Also, you won't find academic references. This is a book for practitioners not academics. There is a list of recommended books (that do have academic references) at the end if readers would like to learn more.

2 Throughout this book I refer to 'product' and 'inventory' or 'stock' but the concepts and approaches I describe can be applied in exactly the same way to operational forecasting in any domain, such as a service operation, for example. Whether you deal in goods or services you need to estimate what you think the demand is likely to be and how much you need to hold in reserve to deal with the inevitability of your forecast being wrong.

SECTION 1
THE PURPOSE OF FORECASTING

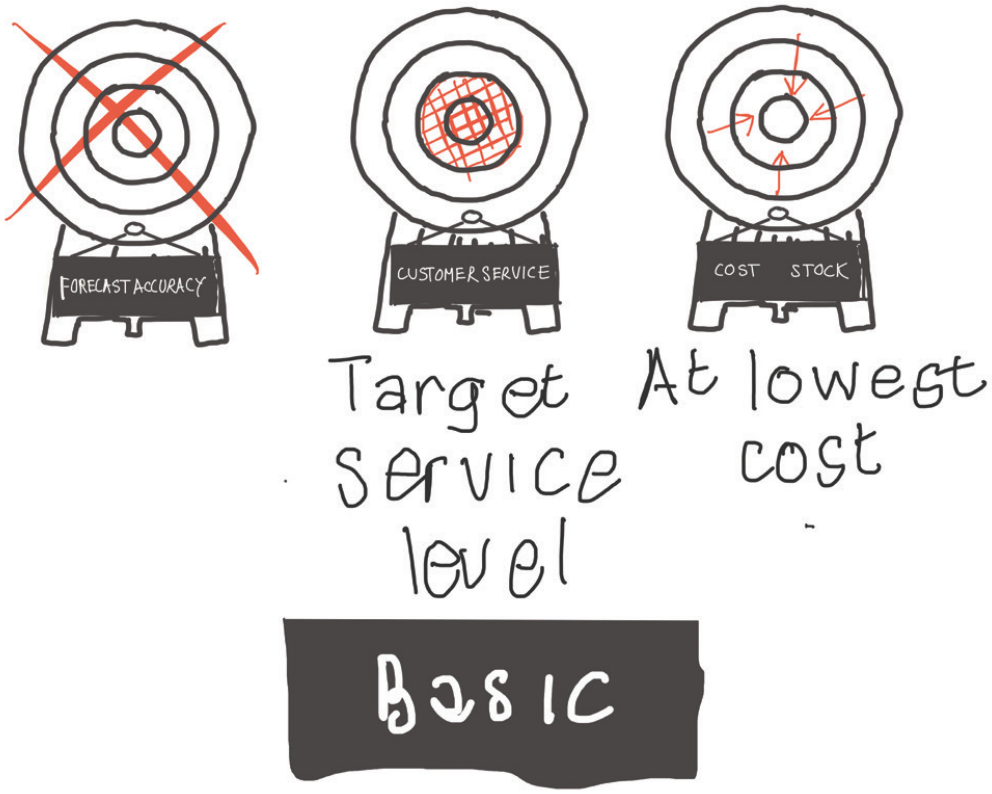

Forecasting is not an end in itself - it is the means to an end.

Success in forecasting, like any other endeavour, can only be determined by how well the process fulfills its purpose.

The purpose of forecasting is explained in this section.

What IS a forecast?

First of all, we need to be absolutely clear what a forecast is – and what it isn't.

A forecast is a best estimate of future outcomes. An expectation of what you think WILL happen.

This is different to a target, which is what you would LIKE to happen. An aspiration.

Ideally, you would like to bring your expectations in line with your aspirations, but more often than not there will be a gap between them. Gaps are important because they tell you that you need to do something different.

Problems start when people get mixed up between forecasts and targets. Very often this happens without anybody ever realizing it.

For example, when people feel good about 'beating their forecast' they are using the forecast as a target… and as a result the forecast will stop doing the job it needs to do. It will become a bad forecast.

If this is what you are doing, stop it!

TAKEOUT

Don't get forecasts confused with targets. A forecast is an expectation - a best guess of what is really going to happen. A target is an aspiration - what we would like to happen.

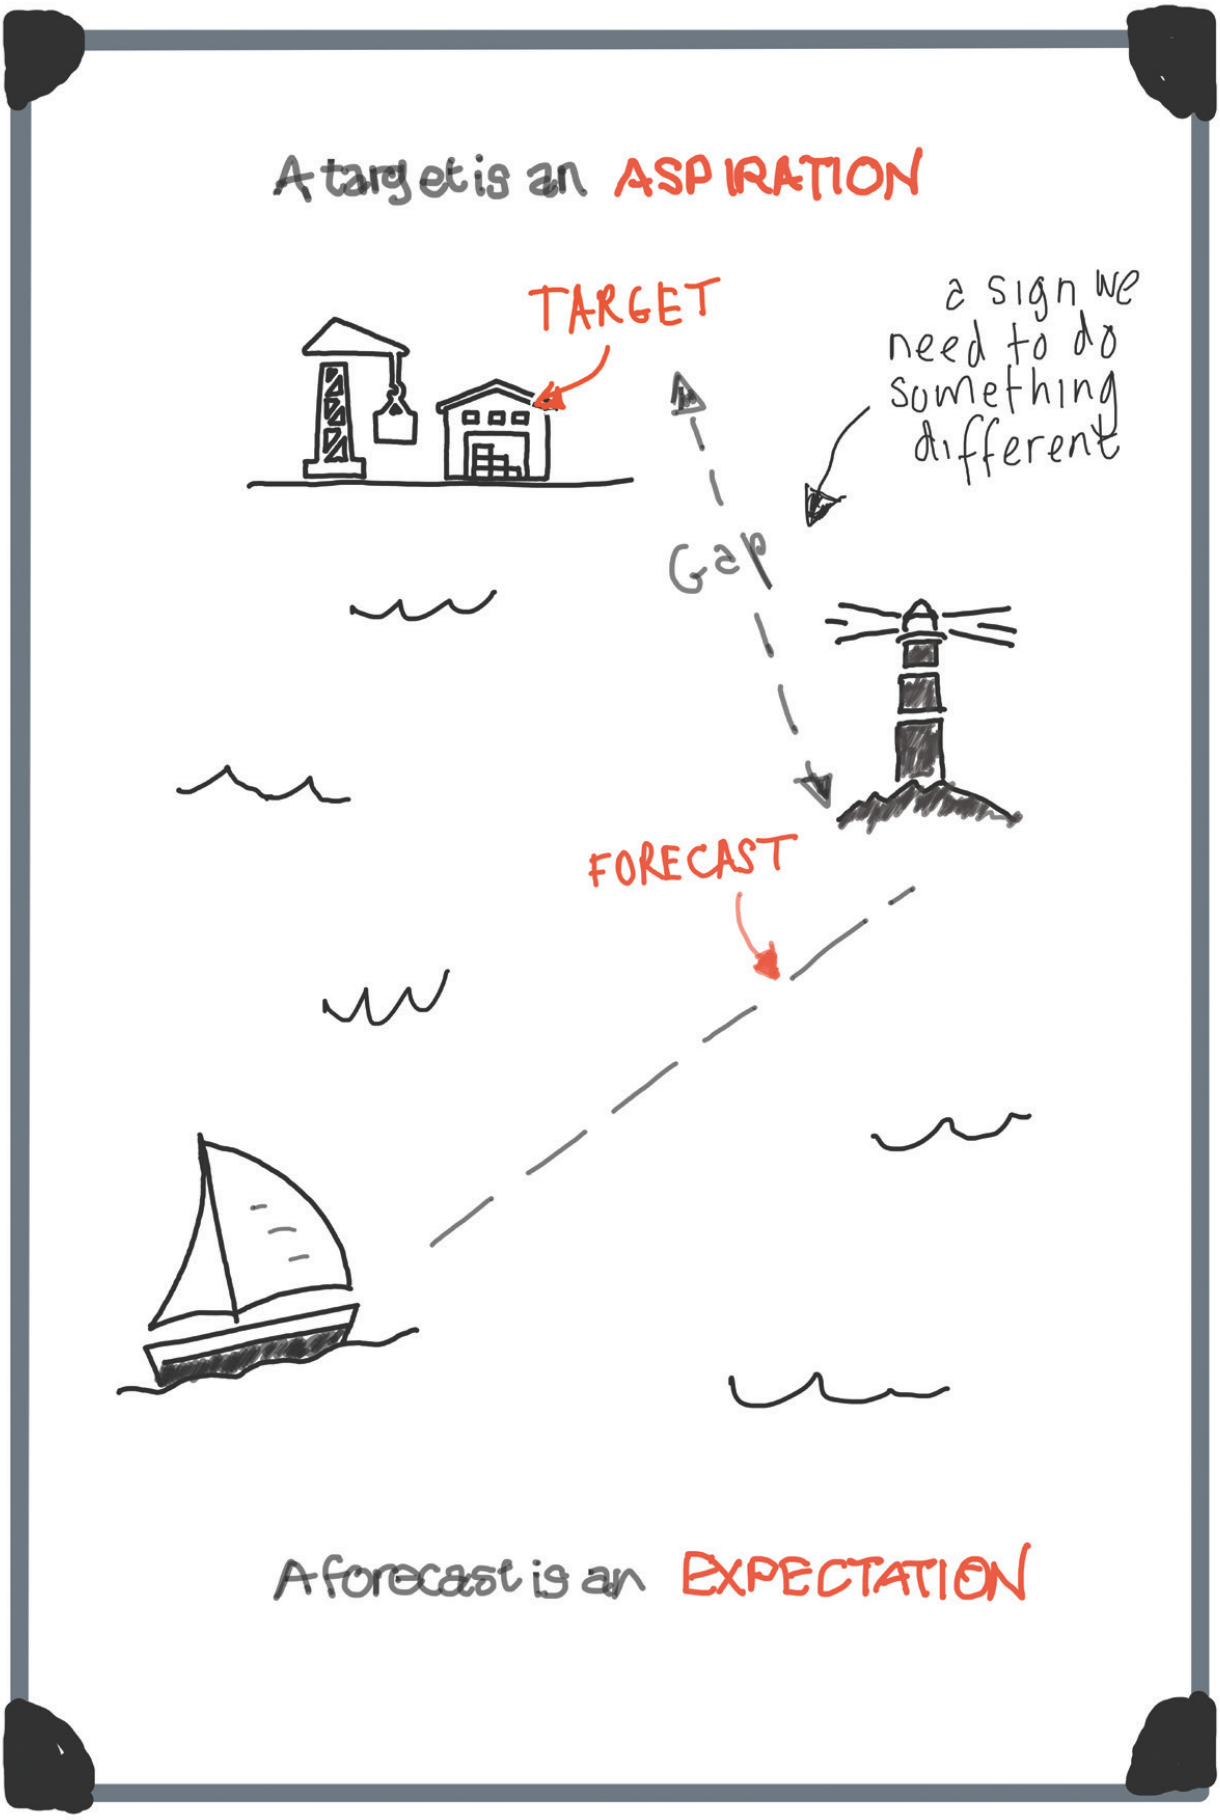
A target is an ASPIRATION
TARGET
a sign we need to do something different
Gap
FORECAST
A forecast is an EXPECTATION

Forecasts are not the same as budgets or plans

You can't have a conversation with someone about anything unless you have a common language, and conversations about forecasting suffer from the inconsistent and undisciplined way certain key terms are used.

Forecasts describe expectations about the future, and these are likely to be different to **budgets**, as these are often the source of targets – an aspiration for the future.

Forecasts and budgets are not the same as, but are based on, **plans,** which are sets of intended actions (and their anticipated impact) and **assumptions** about the future. Those underpinning forecasts will differ from those that informed the budget as new information becomes available and plans are changed in response.

Forecasts, budgets, plans and assumptions are closely related, but they do not describe the same thing.

TAKEOUT

Don't confuse forecasts with budgets, plans or assumptions. Forecasts describe expectations but budgets are a source of aspirations. A plan is a set of intended actions that - along with data about the past and assumptions about the future - are used to create forecasts.

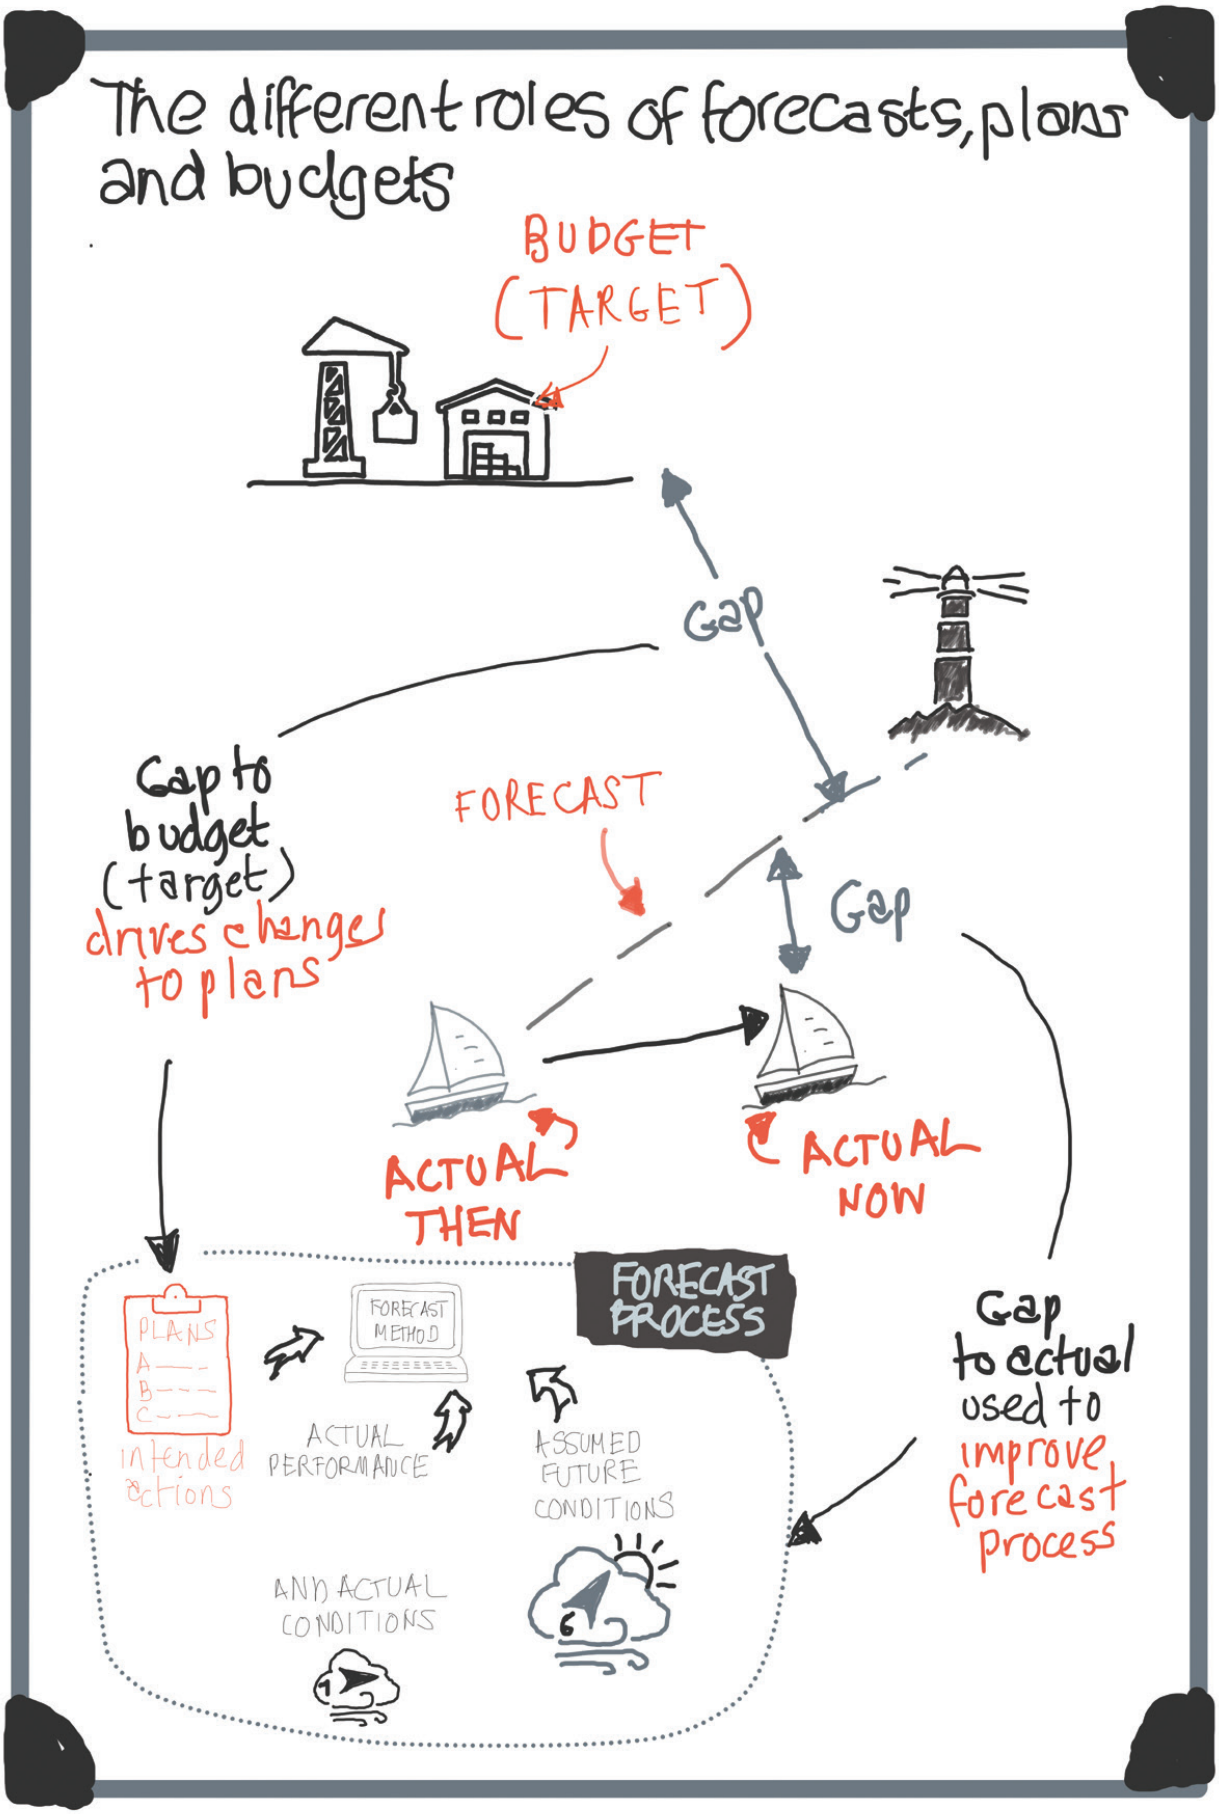
The different roles of forecasts, plans and budgets
BUDGET (TARGET)
Gap
Gap to budget (target) drives change to plans
FORECAST
Gap
ACTUAL THEN
ACTUAL NOW
FORECAST PROCESS
PLANS
A
B
C
intended actions
FORECAST METHOD
ACTUAL PERFORMANCE
ASSUMED FUTURE CONDITIONS
ANY ACTUAL CONDITIONS
Gap to actual used to improve forecast process

Different kinds of forecasts

This book is focused on operational forecasting – the stuff you do to determine what you need to buy, produce, hold in stock or otherwise give your customers what they need.

This differs from other kinds of forecasting that you might do in your business.

Long term forecasting may be needed to help an organization make strategic choices about where and what to invest in. This kind of forecasts help organizations to ADAPT.

Medium-term business forecasts help STEER the business by managing the allocation of resources in order to deliver business performance.

Operational forecasts help businesses RESPOND to demand in the short term, and because the purpose of operational forecasting differs from medium and long terms forecasting, the process tends to be different as well.

In particular, operational forecasts are done in greater detail. Since your customers want to buy a particular kind of say, shoe, in a particular size and colour, producing a forecast for shoes in general won't help you provide good customer service.

Also, most businesses (at least the ones that survive) respond to demand reasonably quickly, so operational forecasts need to be refreshed very frequently.

In summary, operational forecasts are short term, very detailed and frequently refreshed and used to help businesses respond effectively and efficiently to customer demand. Because operational forecasting is a high volume, variety and velocity process it tends to be more reliant on automated mathematical processes than other kinds of forecasts.

TAKEOUT

Don't get operational forecasts confused with other kinds of forecasts which have different purposes and processes. Operational forecasts are short term, detailed and used to help businesses respond to changes in demand.

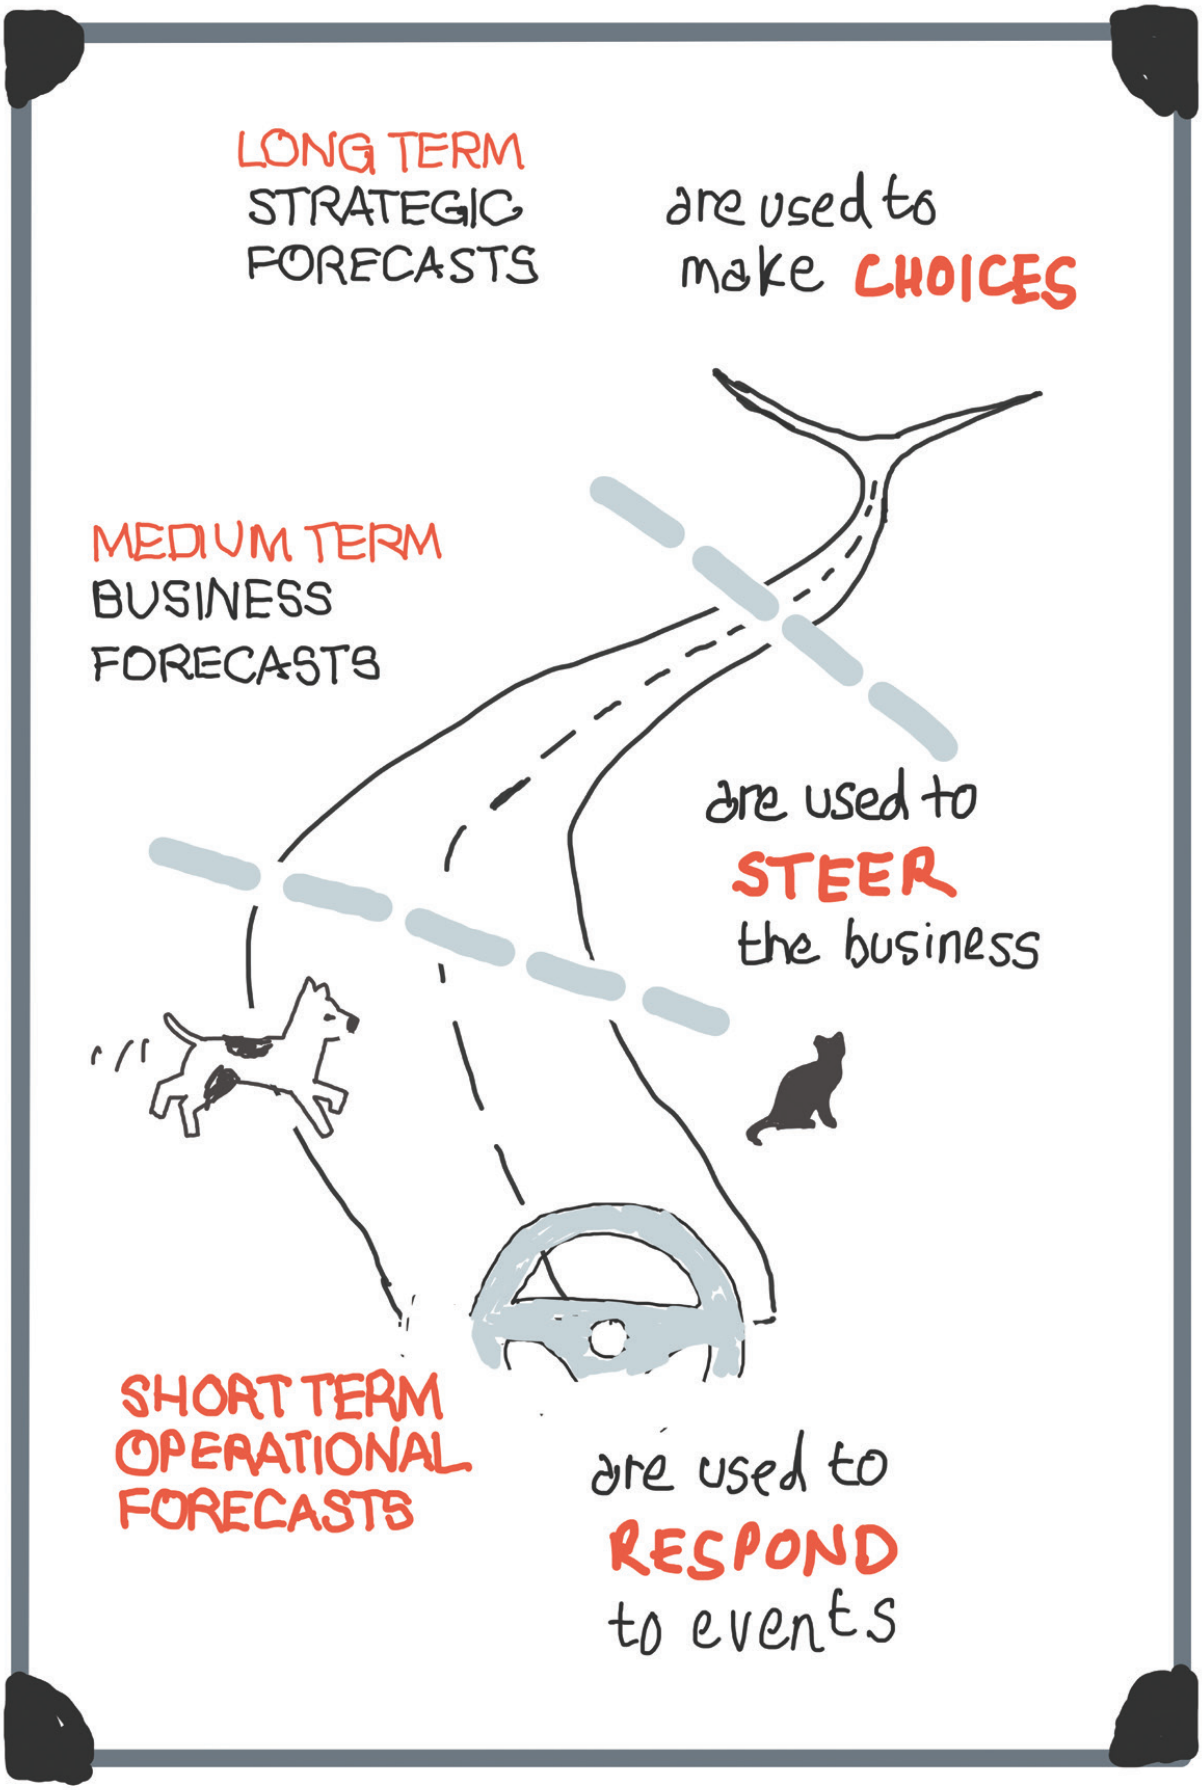
LONG TERM
STRATEGIC
FORECASTS
are used to
make CHOICES
MEDIUM TERM
BUSINESS
FORECASTS
are used to
STEER
the business
SHORT TERM
OPERATIONAL
FORECASTS
are used to
RESPOND
to events

All operational forecasting is based on the same principles

There are many different kinds of operational forecasts – forecasts of sales of products, of demand for services, for power…the list is endless. And there are many different ways that operational forecasts can be generated.

But all operational forecasts serve the same purposes, irrespective of the domain. They are required to determine:

First, how much 'X' is needed in the system to meet anticipated demand within a lead time that is acceptable to the customer. Second, how much spare 'X' is needed to provide a defined level of customer service given the inherent uncertainty in any forecast.

'X' could be products, people or power, for example. And 'spare X' could be safety stock, people on call or power plants ready to be fired up on demand. But whatever the 'X', the basic principles underpinning the forecast process, and how it is used to help make decisions, are the same.

The examples in this book is based on a business selling from stock but the lessons can be directly applied to all other kinds of operational forecasting.

TAKEOUT

Any kind of operational forecasting serves the same purpose: what needs to be supplied to meet demand and what needs to be kept in reserve to ensure that unforecast demand can be met, consistent with desired service levels. Learnings are transferable between domains.

The principles of Operational Forecasting are the same

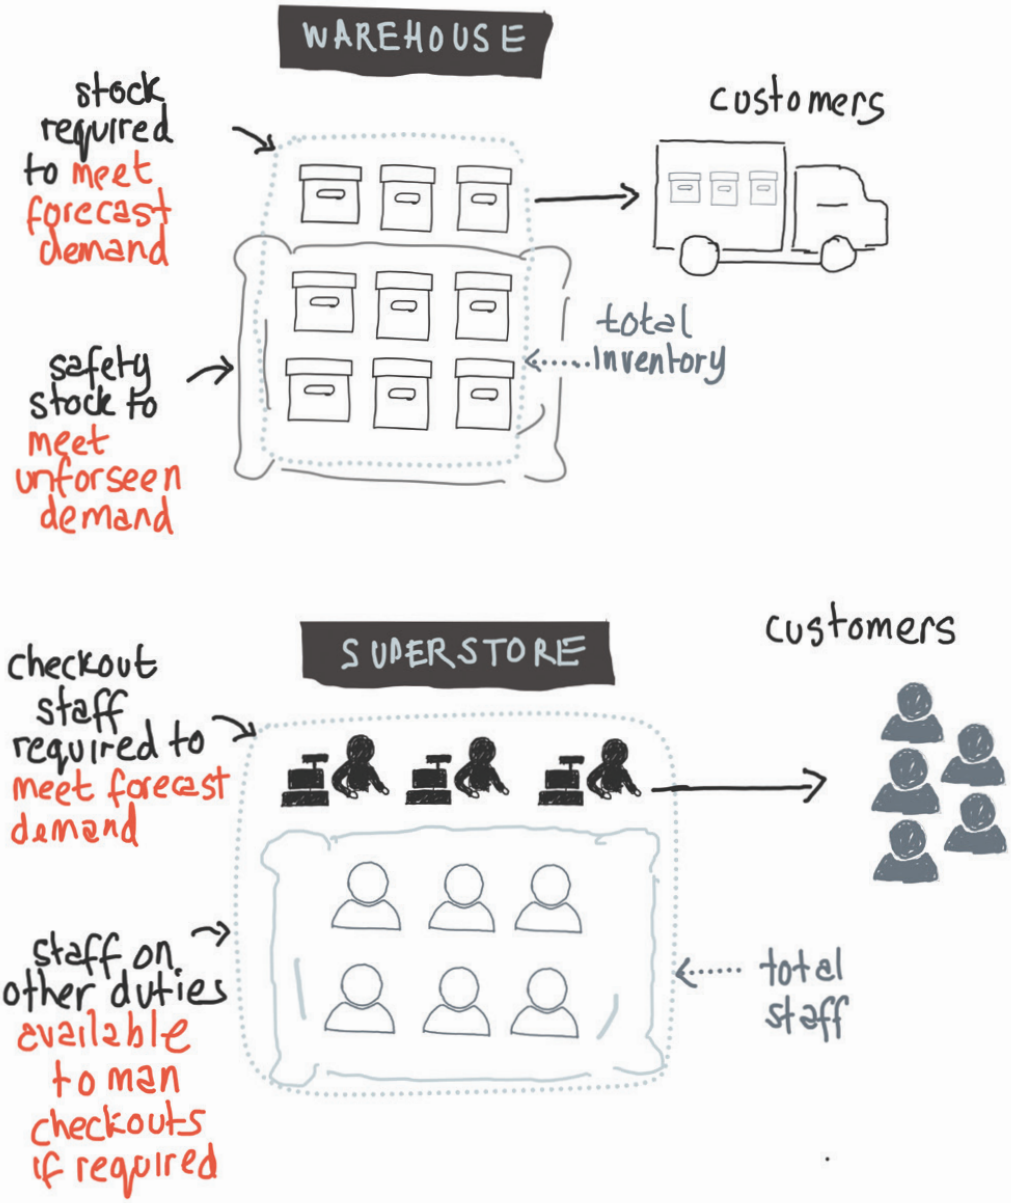

Why operational forecasting is important

In an ideal world, we wouldn't need forecasts. Things would happen and businesses would simply respond to them.

In the real world, this often just isn't possible.

If you have invested millions in a factory to make Product X, you can't instantly switch to making Product Y.

And even if it were physically possible to respond, it might not be economically sensible to do so. It doesn't make sense to manufacture say, a single loaf of bread, every time a customer asks for one.

Instead, many – perhaps most – manufacturers and retailers build stock in anticipation of future need – and this requires a forecast.

The forecast could be very simple like: 'I forecast I will sell the same number of units in this coming period as I did in the last' or it could be very sophisticated, based on complex statistical algorithms.

Whatever the method used to produce a forecast it is important to get it as right as you can.

Having too much stock is very costly. Cash is tied up in unproductive assets, and it costs money to store and transport it to where it is required.

But not having enough inventory means that customer service will suffer and sales may be lost.

TAKEOUT

Pay special attention to operational forecasting since it is one of the few business processes which impacts cash, costs AND customer service.

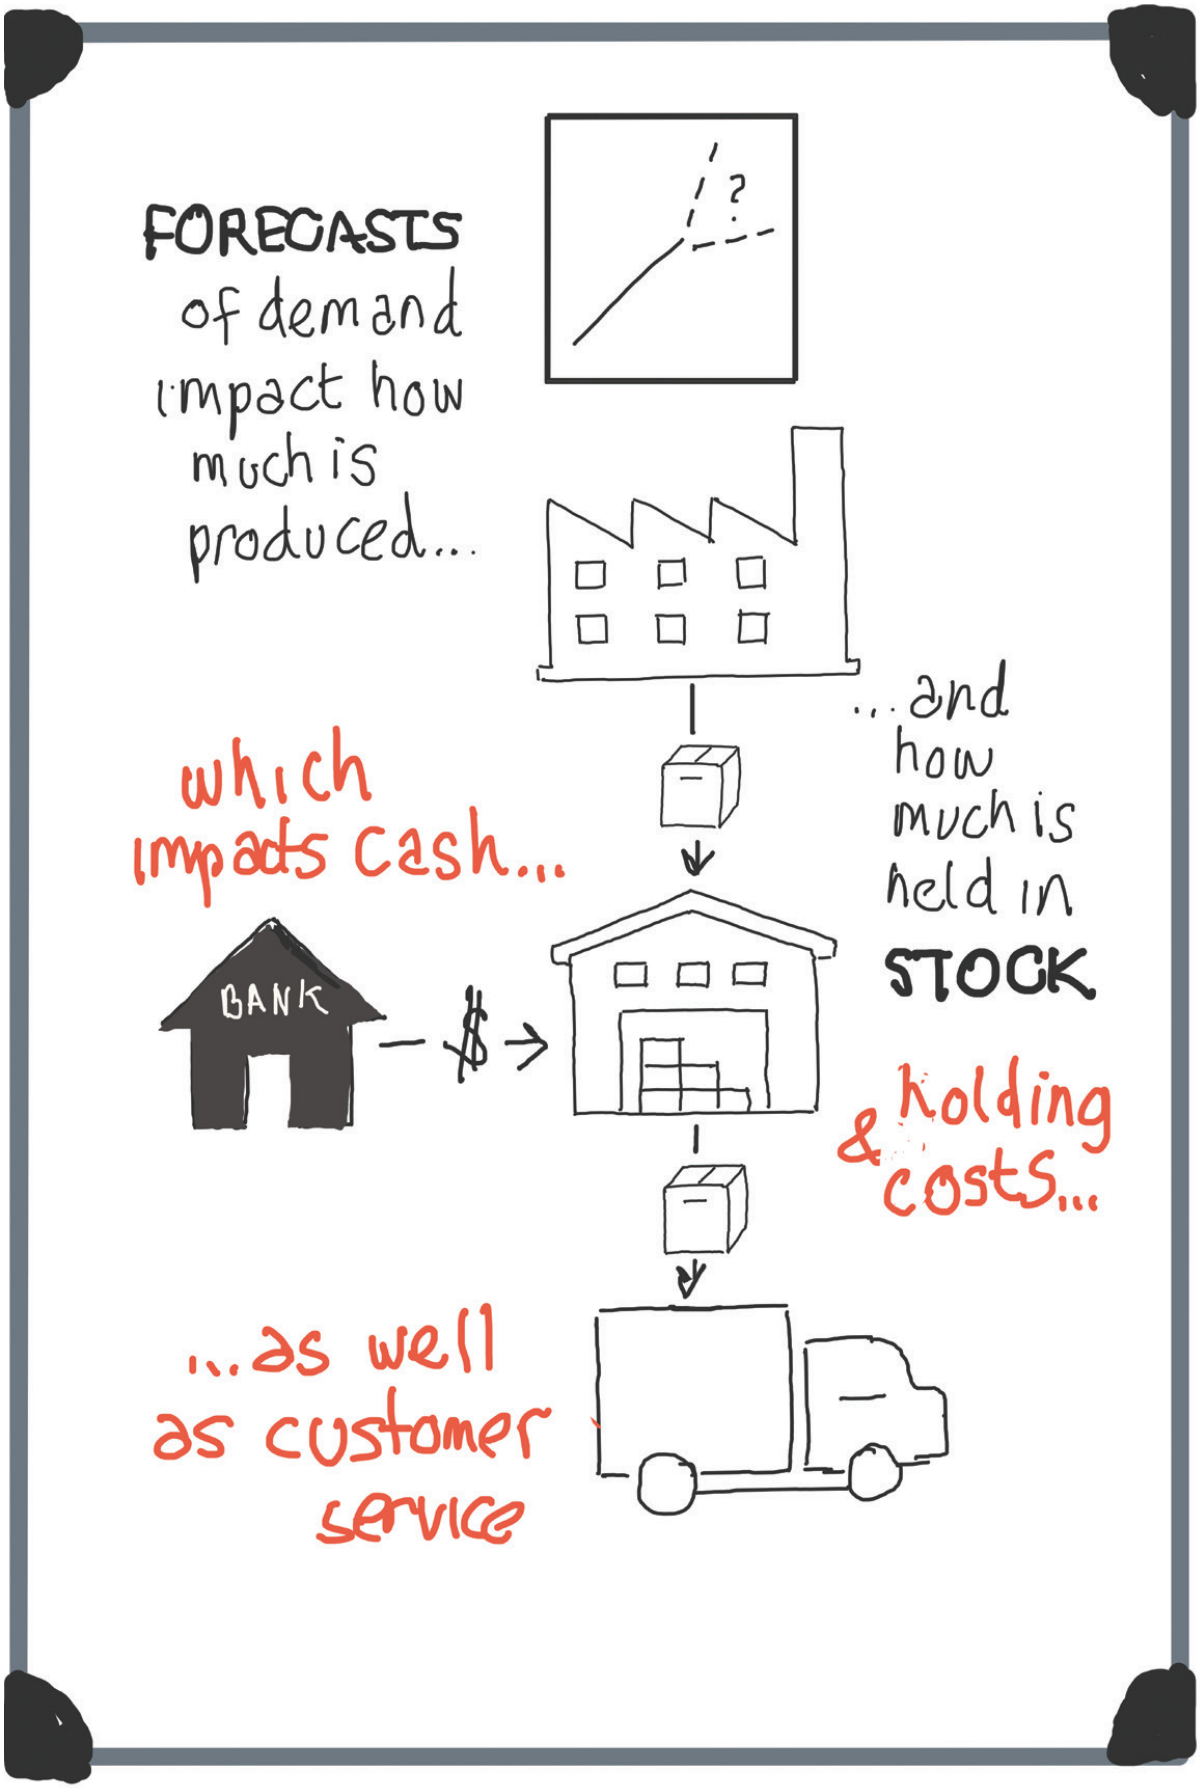
FORECASTS
of demand
impact how
much is
produced...
?
...and
how
much is
held in
STOCK
& holding
costs...
which
impacts cash...
BANK
$
...as well
as customer
service

The quality of forecasts matters...a lot

It is difficult to precisely estimate the business impact of forecast quality partly because it impacts so many variables in ways that are not easy to isolate:

- Inventory levels
- Financing costs
- Warehousing costs
- Transport costs
- Expediting costs (incurred when extra product is required at short notice to meet unforecast demand)
- Lost sales

In addition, poor forecasting can impact the service performance of an organisation, which can have a long-term reputational impact that is difficult to quantify.

Also – as we will discover – it is not straightforward to measure the quality of forecasting in a meaningful way. As a result, most businesses have no idea whether their forecasts are good, bad or indifferent.

What is for sure is that poor forecasting is wasteful and the business impact is huge.

I estimate the unnecessary operating costs of poor forecasting account for 2% of the cost of sales – perhaps 10m for each billion of revenue (assuming a margin of 50%) – and that's not counting the benefits associated with better customer service.

But the 'low hanging fruit' comes from reduced inventory, that across the top 1000 US companies is estimated to be worth $423 billion, with the best performers operating with less than half the level of stock than the median (14 days inventory cover compared to 35).

TAKEOUT

Focus on operational forecasting to improve the bottom line. Poor quality forecasting is the largest unrecognised source of waste in many businesses.

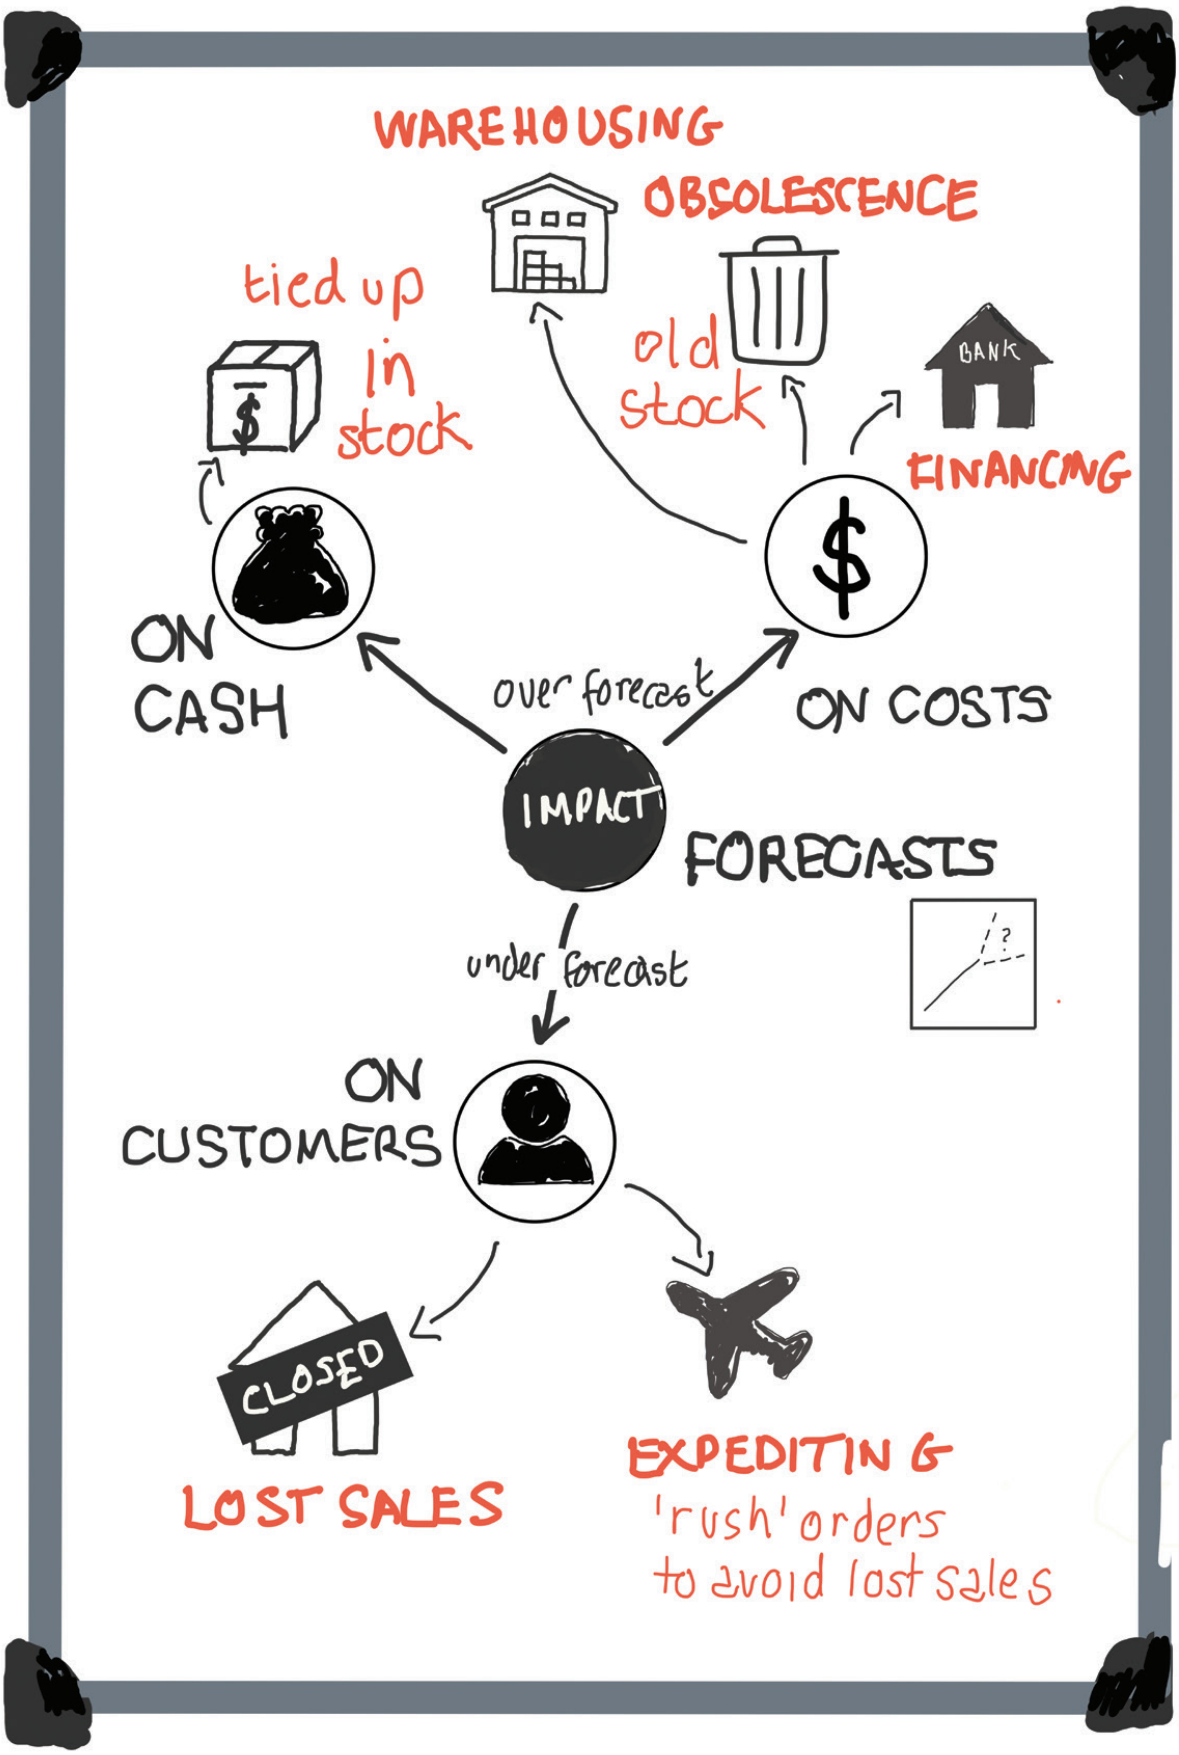
WAREHOUSING
OBSOLESCENCE
tied up
in
stock
old
stock
BANK
FINANCING
ON
CASH
over forecast
ON COSTS
IMPACT
FORECASTS
under forecast
ON
CUSTOMERS
CLOSED
LOST SALES
EXPEDITING
'rush' orders
to avoid lost sales