

The background is a composite image. The top half shows a pilot in a white uniform sitting in a cockpit, looking out. The bottom left shows a large shipwreck with a red hull. The bottom right shows several surgeons in green scrubs and blue masks performing an operation in an operating room.

James Reason

THE HUMAN CONTRIBUTION

**UNSAFE ACTS, ACCIDENTS AND
HEROIC RECOVERIES**

An Ashgate Book

THE HUMAN CONTRIBUTION



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

The Human Contribution

Unsafe Acts, Accidents and Heroic Recoveries

JAMES REASON

Professor Emeritus, The University of Manchester, UK



CRC Press

Taylor & Francis Group

Boca Raton London New York

CRC Press is an imprint of the
Taylor & Francis Group, an **informa** business

CRC Press
Taylor & Francis Group
6000 Broken Sound Parkway NW, Suite 300
Boca Raton, FL 33487-2742

© 2008 by James Reason
CRC Press is an imprint of Taylor & Francis Group, an Informa business

No claim to original U.S. Government works

Printed on acid-free paper
Version Date: 20160226

International Standard Book Number-13: 978-0-7546-7400-9 (Hardback) 978-0-7546-7402-3 (Paperback)

This book contains information obtained from authentic and highly regarded sources. Reasonable efforts have been made to publish reliable data and information, but the author and publisher cannot assume responsibility for the validity of all materials or the consequences of their use. The authors and publishers have attempted to trace the copyright holders of all material reproduced in this publication and apologize to copyright holders if permission to publish in this form has not been obtained. If any copyright material has not been acknowledged please write and let us know so we may rectify in any future reprint.

Except as permitted under U.S. Copyright Law, no part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, microfilming, and recording, or in any information storage or retrieval system, without written permission from the publishers.

For permission to photocopy or use material electronically from this work, please access www.copyright.com (<http://www.copyright.com>) or contact the Copyright Clearance Center, Inc. (CCC), 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. CCC is a not-for-profit organization that provides licenses and registration for a variety of users. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

Trademark Notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Visit the Taylor & Francis Web site at
<http://www.taylorandfrancis.com>

and the CRC Press Web site at
<http://www.crcpress.com>

Contents

<i>List of Figures</i>	<i>ix</i>	
<i>List of Tables</i>	<i>xi</i>	
<i>About the Author</i>	<i>xiii</i>	
PART I	INTRODUCTION	1
Chapter 1	The Human Contribution: Hazard and Hero	3
	Introduction	3
	The structure of the book	4
	About the book	8
Chapter 2	A Mind User's Guide	9
	Tip-of-the-tongue state	9
	The conscious and automatic modes of control	10
	Three levels of performance	13
	Interacting with the long-term knowledge base	15
	Intentions and the retrieval cycle	17
	Concurrent processing	18
	The relationship between memory and attention: the blob-and-the-board model	20
	Summing up	22
PART II	UNSAFE ACTS	27
Chapter 3	The Nature and Varieties of Human Error	29
	Defining and classifying error	29
	Error myths	36
	Slips and lapses	38
	Rule-based mistakes	45
	Knowledge-based mistakes	45
	Conclusion: a general rule	46

Chapter 4	Violations and the Varieties of Rule-related Behaviour	49
	Chernobyl and Zeebrugge	49
	Violations considered as unsafe acts	51
	Who is most likely to violate?	55
	Why do people violate safety rules?	55
	The 'mental economics' of violating	57
	Bad procedures	58
	Procedure-usage	58
	Testing two models of violating behaviour	59
	The varieties of rule-related behaviour	60
	Great improvisers	66
	End piece	66
Chapter 5	Perceptions of Unsafe Acts	69
	The plague model	70
	The person model	71
	The legal model	90
	The system perspective	92
	Person and system models: getting the balance right	102
PART III	ACCIDENTS	105
Chapter 6	Error Traps and Recurrent Accidents	107
	Accident proneness: a quick survey	107
	Everyday error traps	111
	Recurrent accident patterns	113
	The elements of recurrent scenarios	124
	Cultural drivers	126
	Conclusion	127
Chapter 7	Significant Accident Investigations	129
	Problems with the past	129
	Changes in accident investigation	131
	The Mahon and Moshansky reports	133
	Has the pendulum swung too far?	136
	Conditions and causes	137
	Counterfactual fallacy	137
	The current view	138

PART IV	HEROIC RECOVERIES	141
Chapter 8	Training, Discipline and Leadership	143
	The Light Division's retreat at Fuentes de Onoro (1811)	143
	The withdrawal of the 1st Marine Division from Chosin Reservoir (1950)	149
	Concluding remarks	159
Chapter 9	Sheer Unadulterated Professionalism	161
	Captain Rostron and the rescue of the <i>Titanic</i> survivors (1912)	162
	Saving <i>Apollo 13</i> (1970)	169
	British Airways Flight 09 (1982)	177
	The BAC 1-11 incident (1990)	180
	Surgical excellence (1995–97)	184
	Minor and major events	185
	Concluding remarks	193
Chapter 10	Skill and Luck	195
	The Gimli Glider	195
	Captain Al Haynes and United 232	200
	Final word	207
Chapter 11	Inspired Improvisations	209
	General Gallieni and the Paris taxis	209
	Captain Gordon Vette and the rescue of Jay Prochnow	213
	Conclusion	218
Chapter 12	The Ingredients of Heroic Recovery	221
	Coping with expected hazards	221
	Coping with unlikely but possible hazards	223
	Generic qualities	229
	Conclusion	236
PART V	ACHIEVING RESILIENCE	237
Chapter 13	Individual and Collective Mindfulness	239
	Consistency versus variability	239

A dynamic non-event	239
Collective mindfulness	240
Individual mindfulness	243
Aspects of resilience	248
Foresight training at the UK National Patient Safety Agency	250
Organisational support	251
Looking towards the future	253
Mindfulness and resilience	262
Chapter 14 In Search of Safety	265
Introduction	265
What does the term 'safety' mean?	265
The two faces of safety	267
The 'safety space' model	268
What does a resilient system look like?	276
The knotted rubber band model	280
Defining the nature of positive safety	284
Final words	287
<i>Index</i>	289

List of Figures

Figure 2.1	Three levels of performance control	13
Figure 2.2	The 'blob-and-the-board' model	21
Figure 5.1	First version of the Swiss cheese model (though it hadn't yet taken on its Emmenthole appearance. Various human contributions to the breakdown of complex systems are mapped on to the basic elements of production	97
Figure 5.2	Part of the earliest version of the Swiss cheese model. The diagram shows a trajectory of accident opportunity penetrating several defensive layers, and begins to have Emmenthole-ish features	98
Figure 5.3	The early 90's variant of the Swiss cheese model	99
Figure 5.4	A mid-90s variation with two interacting pathways to an accident	100
Figure 5.5	The latest version of the Swiss cheese model	102
Figure 6.1	Showing the unequal liability to accidents in a group exposed to the same hazards over the same time period. The curve represents the chance prediction of event likelihood (Poisson distribution). The columns represent the actual number of accidents	108
Figure 6.2	Kanizsa Triangle	122
Figure 7.1	Tracing the shifting emphases of accident investigations	131
Figure 8.1	The manoeuvres required in preparing to receive cavalry on the march.	144
Figure 8.2	The Battle of Fuentess de Onoro, 5 th May 1811	147
Figure 8.3	Showing the area over which the Chosin Reservoir campaign was fought	151

Figure 9.1	A cutaway diagram of the service module, It was the round cryogenic oxygen tank (second layer) that exploded and nearly destroyed the spacecraft	171
Figure 9.2a	Surgical defences in the arterial switch operation represented as Swiss cheese slices in which the holes do not line up	188
Figure 9.2b	Using the Swiss cheese model to represent a typical major event scenario	189
Figure 9.3a	The slice of Cheddar at the left hand end represents coping abilities	189
Figure 9.3b	The mouse nibbling at the slice of Cheddar represents how adverse events, both major and minor, can eat away at limited coping resources	190
Figure 11.1	Flying the aural box pattern into the VHF range circle	216
Figure 13.1	Combining individual and collective mindfulness to enhance patient safety	242
Figure 13.2	The 'three-bucket model for assessing high-risk situations	246
Figure 13.3	How the bucket contents might be interpreted	247
Figure 13.4	Tracing the possible progress of developments in patient safety	255
Figure 13.5	Showing the reduction in variability as the cycle progresses	262
Figure 14.1	Illustrating vulnerability and resistance. The arrows at the top of the figure represent perturbing forces	268
Figure 14.2	Showing a number of hypothetical organisations within the same hazardous domain distributed throughout the safety space	269
Figure 14.3	Summarising the driving forces and navigational aids necessary to propel an organisation towards the region of maximum resistance	275
Figure 14.4	Three states of the knotted rubber band	280
Figure 14.5	Showing the resource implications of the knotted rubber band model	283

List of Tables

Table 2.1	Comparing the properties of the conscious workspace and the long-term knowledge base	12
Table 3.1	Comparing the properties of exceedances, incidents and accidents	36
Table 4.1	Summarising the violation ‘balance sheet’	58
Table 4.2	Summarising the 12 varieties of rule-related behaviour	62
Table 5.1	Comparing the CANDU and RBMK safety features	81
Table 7.1	Three examples of the distinction between cause and condition	137
Table 14.1	Summarising the interactions between reactive and proactive measures	272
Table 14.2	Combining the 3Cs and the 4Ps to produce 12 sets of indicators	276



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

About the Author

James Reason was Professor of Psychology at the University of Manchester from 1977–2001, from where he graduated in 1962. He obtained his PhD in 1967. From 1964–76, he was Lecturer then Reader in Psychology at the University of Leicester. He has also worked at the RAF Institute of Aviation Medicine, Farnborough and the US Naval Aerospace Medical Institute, Pensacola.

His primary research interest has been the human and organizational contributions to the breakdown of complex, well-defended systems. He has written books on absent-mindedness, human error, aviation human factors, on managing the risks of organizational accidents and, most recently, on error management in maintenance operations. He has researched and consulted in the fields of aviation, railways, nuclear power generation, maritime safety, oil exploration and production, mining, chemical process industry, road safety, banking and health care.

He received the Distinguished Foreign Colleague Award from the US Human Factors and Ergonomics Society (1995), the Flight Safety Foundation/Airbus Industrie Award for achievements in human factors and flight safety (2001), and the Roger Green Medal from the Royal Aeronautical Society for contributions to human factors as applied to aerospace (2001), and the Flight Safety Foundation/Boeing Aviation Safety Lifetime Achievement Award (2002). He is a Fellow of the British Academy, the Royal Aeronautical Society and the British Psychological Society. He received an honorary DSc from the University of Aberdeen in 2002, and was awarded a CBE in 2003 for contributions to patient safety. In 2006, he was made an honorary fellow of the Royal College of General Practitioners.



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

PART I

Introduction



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

Chapter 1

The Human Contribution: Hazard and Hero

Introduction

The purpose of this book is to explore the human contribution to both the reliability and resilience of complex well-defended systems. The predominant mode of treating this topic is to consider the human as a hazard, a system component whose unsafe acts are implicated in the majority of catastrophic breakdowns. But there is another perspective, one that has been relatively little studied in its own right, and that is the human as hero, a system element whose adaptations and compensations have brought troubled systems back from the brink of disaster on a significant number of occasions.

After studying human unsafe acts within hazardous enterprises for more than three decades I have to confess that I find the heroic recoveries of much greater interest and – in the long run, I believe – potentially more beneficial to the pursuit of improved safety in dangerous operations. Since most observations of people in high-risk systems are event-dependent; that is, they generally emerge from well-documented accident investigations, it is inevitable that we should know far more about the hazardous human than the heroic one.

But there is a stark contrast between unsafe acts and these intrepid recoveries. Errors and violations are commonplace, banal even. They are as much a part of the human condition as breathing, eating, sleeping and dying. Successful recoveries, on the other hand, are singular and remarkable events. They are the stuff of legends. But these abilities are not unattainable. A few people are born heroic, but most of us can acquire the skills necessary to give a better than an evens chance of thwarting a disaster scenario. However, it must be acknowledged that such

heroism is not necessarily an enduring characteristic. Even the best people have their bad days.

The Structure of the Book

The book is organised into five parts. The present chapter introduces the subject matter. [Chapter 2](#) offers an inside-out guide to being a mind user. We all know what it feels like to have a mind, but we don't always appreciate that 'feelings-of-knowing' can sometimes be very misleading. There are things you think you know about how your mind works, but don't. And there are things you think you don't know, but actually do. To realise your mind's full potential, you need to understand when these feelings of knowing (or not knowing) are useful and when they are deceptive. I will be discussing these mental processes from a user's perspective: from the inside looking out, not from the outside looking in – hence an inside-out view.

I'm not denying the intimate connection between the mind and the brain. It is just that this book is not about brain scans and neural wiring diagrams. While modern techniques can tell us a great deal about the brain's structure and function, they never wholly capture the moment-to-moment experiences of being a mind user, and none can track all of the subtle interactions between the conscious and automatic control processes. And that is what interests me here.

[Part II](#) is concerned with unsafe acts: errors and violations, and how they are perceived by those upon whom they impact. Unsafe behaviour may be less fascinating than heroism, but it is no less important.

[Chapter 3](#) focuses on the nature and varieties of human error. In order to limit the damaging occurrence of errors and improve their chances of detection and recovery, we need to understand something of their cognitive origins and the circumstances likely to promote them. This understanding can translate into 'error wisdom' at the sharp end – what has been termed 'individual mindfulness', to be discussed in [Part V](#).

[Chapter 4](#) deals with rule-related behaviour. I begin by considering the various types of violation, and then discuss the social, emotional and systemic factors that lead people to

choose not to comply with rules, regulations and safe operating procedures. However, such acts of non-compliance are not universally bad. They can have beneficial as well as unwanted consequences. This becomes evident when we look in detail at the 12 varieties of rule-related behaviour.

Chapter 5 examines a number of different perceptions of human unsafe acts, of which the two most dominant are the person and the system models. Each has its own theory of how these unsafe acts arise, and how they might be remedied and managed. The person model, asserting that errors originate within the minds of the people concerned, is intuitively appealing and still holds sway in many domains. However, over recent years, the system model has gained increasing ascendancy. This argues that the people on the frontline are not so much the initiators of bad events as the inheritors of long-term system failings. My thesis is that the extremes of both views have their shortcomings. We need to strike a balance between the two.

Part III deals with accidents and their investigation. One fact that lends strong support to the system approach is that similar situations keep provoking the same kinds of unsafe acts in different people. These recurrences, discussed in **Chapter 6**, indicate that a substantial part of the problem is rooted in error-provoking situations rather than in error-prone people. A primary function of error and incident reporting systems is to identify these 'error traps'. Eliminating them becomes a priority task for error management programmes.

Complex hazardous systems are subject to two kinds of bad event: individual accidents, resulting in limited injury or damage, and organisational accidents that occur relatively infrequently but whose consequences can be both devastating and far-reaching. One of the features that discriminates between these two kinds of adverse event is the degree of protection available against the foreseeable hazards. Whereas individual accidents usually result from the failure of very limited safeguards (or their absence), organisational accidents involve a concatenation of breakdowns among many barriers, safeguards, and controls. It is this combined failure of the many and varied 'defences-in-depth' that characterises the organisational accident and it is this type of event that will be the main concern of this book.

[Chapter 7](#) focuses on two pioneering accident investigations that fundamentally changed the way the human contribution to bad events is regarded. In particular, they spelled out how unsafe acts and latent organisational conditions (resident pathogens) interact to breach the multi-layered system defences. It also traces how the emphases of investigations have shifted from technical and human failures at the sharp end to examining the effects of ‘upstream’ factors such as organisational processes, safety culture, regulation and even the economic and political climate. It is suggested that perhaps the pendulum has swung too far towards identifying causal factors that are remote in time and place from the local events. This chapter also looks at some of the problems facing accident investigators, and others who seek to make sense of the past. One such problem is the failure to differentiate between conditions and causes, thus falling foul of the counterfactual fallacy

For these and related reasons, it is argued that continual tensions between production and protection lead to resident pathogens being seeded into the system, and this is true for all systems. But such organisational shortcomings are conditions rather than causes. Although they contribute to defensive failures, they are not in themselves the direct causes of accidents. The immediate triggers for such bad events are local circumstances: human and technical failures that add the final ingredients to an accident-in-waiting that may have been lurking in the system for many years. All systems, like human bodies, have resident pathogens. They are universals. It is usually only the proximal factors, immediate in both time and space to the accident, that distinguish between a system suffering a catastrophic breakdown and those in the same sphere of operations that do not.

Up to this point, the book deals mainly with the human as a hazard. In [Part IV](#), we look at the other side of the coin: the human as hero. Eleven stories of heroic recovery are told. They are grouped into four chapters:

- [Chapter 8](#) (training, discipline and leadership) examines two military case studies: the retreat of Wellington’s Light Brigade on the Portuguese–Spanish border in 1811; and the retreat of the US 1st Marine Division from the Chosin Reservoir in 1950.

- [Chapter 9](#) (sheer unadulterated professionalism) deals with Captain Rostron and the rescue of the *Titanic* survivors in 1912; the recovery of *Apollo 13* in 1970; the Boeing-747 Jakarta incident; the recovery of the BAC 1-11 in 1990; and surgical excellence as directly observed in 1995–96.
- [Chapter 10](#) (luck and skill) looks at the near-miraculous escapes by the ‘Gimli Glider’ on the edge of Lake Winnipeg in 1983 and United 232 at Sioux City in 1989.
- [Chapter 11](#) (inspired improvisations) discusses General Gallieni and the ‘miracle on the Marne’ in 1914; and the saving of Jay Prochnow lost in the South Pacific by Captain Gordon Vette in 1978.
- What, if anything, did these heroes have in common? [Chapter 12](#) seeks to identify the principal ingredients of heroic recovery.

[Part V](#) (Achieving Resilience) has two chapters. [Chapter 13](#) elaborates on Karl Weick’s notion of ‘mindfulness’. In its broadest sense this involves intelligent wariness, a respect for the hazards, and being prepared for things to go wrong. Mindfulness can function both at the level of the frontline operators and throughout the organisation as a whole. The former we term ‘individual mindfulness’ and the latter ‘collective mindfulness’. Both are necessary to achieve enhanced systemic resilience. We can’t eliminate human and technical failures. And no system can remain untouched by external economic and political forces. But we can hope to improve its chances of surviving these potentially damaging disruptions in its operational fortunes.

The last chapter deals with the search for safety, the broadest part of the book’s spectrum. Two models of safety are described: the safety space model and the knotted rubber band model. The former operates at the cultural and organisational levels; the latter has a more tactical focus and deals with keeping some continuous frontline process within safe boundaries. Together, they have implications for re-engineering an existing culture to improve safety and resilience. Or, to put it another way, this concluding chapter is concerned with the practical measures necessary to achieve states of both individual and collective mindfulness.

About the Book

I should end this introduction by saying something about the readership and style of the book. Perhaps I should begin by saying what it is not. It is not a scientific book, even though it touches upon scientific and technological issues. It is written in the first person. That means that these are my personal views – and prejudices. It is written for real people, even though academics and students may find parts of it of interest. But it does not require any prior knowledge of academic psychology, although these issues are touched upon in the [first chapters](#). Nor is it a ‘how-to-do’ book. If there is any way of describing the content, I would say it was about the philosophy of managing complex hazardous systems. Philosophy is a daunting word, but in this book it simply means a way of thinking about the issues. In short, it is a way of confronting the problems of conducting a hazardous operation so that you keep your risks as low as reasonably practicable and still stay in business. It is this latter injunction that, for me, is the most important one.

Chapter 2

A Mind User's Guide

After using a mind for seventy years, I realise that I know very little about it, and the older I get the more this conviction grows. It is true that after nearly forty years of researching and teaching psychology, I do have an inkling of what I do and don't know. It is also the case that I have some understanding of what I think I know, but really don't. And, if I ask myself the right questions, I can occasionally come up with things that I didn't think I knew, but actually did. For all that, though, much of my mental functioning remains secretive, seemingly out of reach and full of surprises.

But – you might be thinking – Sigmund Freud told us that over a hundred years ago.¹ So what is new? Quite a lot, as I will hope to show later. Freud is closely linked with the idea of an unconscious mind, but he did not invent the term, nor is his rather narrow view of the unconscious widely accepted by contemporary psychologists. I do not reject the idea of an unconscious mind – indeed, its existence is the main reason for including this chapter – but I do dispute the strict Freudian interpretation of its role in our mental lives.

My purpose here is to make you, the everyday mind user, more familiar with the mysteries of your own mental life, and – once in a while – to tell you something that you did not already know. This chapter is also intended to act as an introduction to the discussion of errors and violations in the next two chapters.

Tip-of-the-Tongue State

Let's start with a commonplace experience. There's nothing quite like the 'tip-of-the-tongue' (TOT) state to expose the subtleties

¹ Freud, S. (1914) *Psychopathology of Everyday Life*. London: Ernest Benn. (Originally published in 1901).

of knowing and not knowing the things that go on in your own mind. A TOT state begins with an attempt to retrieve from memory something that you are sure you know, but then the search fails to yield an immediate, felt-to-be-correct response. Instead, it produces a name or a word that you recognise as being 'warm', but you also know that this is not the sought-for item.

When you persist with the search, the same wrong item keeps coming to mind in an irritatingly obtrusive fashion. What makes the whole experience so frustrating is that this recurrent blocker is felt to be very close to the target item. We appreciate that it might have similar properties, like sound, structure or meaning, but yet we are certain it's wrong. How do we know these things when we can't directly access the right word or item? Some part of the mind knows, but it's not the conscious part.

Here is an actual example. I was searching for the name of D.W. Griffith's silent film *Intolerance*. Every time I made an active memory search I came up with 'intemperance'. I knew this was wrong, but it felt very close. I had some conscious information about the sought-for word – I knew it was a single-word title, I had a rough idea of the number of syllables, and I had a strong feeling that it began with 'I'. But what I did not know at the outset was that the word had an '-erance' ending, though it was clear that some part of my mind knew this by the repeated retrieval of 'intemperance'.

The search experience was rather like standing on a small stage – equating to my conscious awareness – with large wings on either side. I would go to one of the wings and call out the search cues: it's the title of a classic film; it begins with 'I'; it has many syllables; it covered a wide historical sweep of man's inhumanity to man; and so on. Then 'intemperance' would be repeatedly thrown on to my conscious stage, and a voice from the other wing would call 'no, but you're very close'. Somewhere off-stage was the correct template for the sought-for word, but I couldn't access it directly. I only recognised it when the search finally produced the right answer, *Intolerance*, which I immediately knew was correct.

The Conscious and the Automatic Modes of Control

This TOT experience shows that one of the main problems with the human mind is that the user is only in direct conscious

contact with a fraction of the whole. The conscious part seems to be located somewhere between the ears and behind the eyes. At any one moment in this very limited space the larger part of our current waking thoughts and feelings are experienced, our sense data interpreted and our present actions planned, initiated and monitored. It is also this tiny space that feels at that instant most closely identified with our innermost selves – our personal beliefs, attitudes, values, memories, likes and dislikes, loves and hates, and the other passing clutter and baggage that goes to make up one's mental life. But we are only aware of a very restricted amount at any one time. The ideas, feelings, images and sensations seem to flow like a stream past a blinkered observer standing on the bank. We can't see far upstream or downstream, but we can take in between one to two seconds worth of what goes past. This is what comprises the conscious workspace, the experiential here and now.

Beyond this present awareness lies a vast and only partially accessible knowledge base. Some of the information contained there is in the form of previous life events (episodic memory), though this becomes very patchy for the time before we were five years old, maybe even later. Other parts of it are used to make sense of the world (semantic memory). And yet other knowledge structures (called schemas) control our routine perceptions, thoughts, words and actions.

We have a rough idea of the contents of this long-term knowledge base – not all of them, of course, but enough to be aware of the general headings. But what we don't know is how stored items are called to mind. Such retrievals can be so accurate and so immediate as to convince us – incorrectly, as I hope the TOT example has shown – that we have direct voluntary access to all parts of the store. The reality is that while we are conscious of the products of such retrievals – the words, feelings, images, thoughts and actions – we have little or no awareness of the processes that seek them out and call them to mind. Understanding this is very important since most of our mental lives involve a continuous interaction between the conscious workspace and the long-term memory store. Sometimes we deliberately call items to mind, but at other times they simply pop up unbidden.

The continuous interplay between the conscious workspace and long-term memory or knowledge base will occupy us for the remainder of this chapter. These two co-existing, and sometimes competing controllers of mental life have markedly contrasting properties. These properties are summarised in [Table 2.1](#).

Though these two mental components work in harmony for much of the time, they can also compete for command of the body's output mechanisms, both in the observable physical world, through unintended words and actions, and in the conscious workspace, into which items may be delivered without conscious intent. This is hardly surprising given their radically differing properties and the power of familiar environments to evoke habitual responses.

Table 2.1 Comparing the properties of the conscious workspace and the long-term knowledge base

Conscious Workspace	Long-term Knowledge Base
Accessible to consciousness. Closely linked with attention and working memory.	While the products (actions, thoughts, images, etc.) are available to consciousness, the underlying processes are largely outside its reach.
Selective and resource-limited.	Apparently unlimited in both the amount of stored information and the length of time for which it is retained.
Slow, laborious and serial (one thing after another).	Fast, effortless and parallel (many things at once).
Intermittently analytical. Sets intentions and plans and can monitor their progress at the various choice points. But this often fails.	Automatic in operation.
Computationally powerful. Accepts inputs from nearly all senses. Vision dominates.	Behaviour governed by stored specialised knowledge structures (schemas) that respond only to related sensory inputs and do their own thing.
Accesses long-term memory by generating 'calling conditions' or retrieval cues.	Two basic retrieval processes: similarity-matching (like with like), and frequency-gambling (resolving possible conflicts in favour of the most frequent, recent or emotionally charged items).

Three Levels of Performance

The extent to which our current actions are governed either directly by conscious attention or more remotely by pre-programmed habit patterns gives rise to three levels of performance: knowledge-based, rule-based and skill-based. Their contrasting characteristics are summarised in [Figure 2.1](#).

All human performance – with the exception of what comes ‘hard-wired’ at birth – begins at the knowledge-based level in which our actions are governed online by the slow, limited, and laborious application of conscious attention. This level relies very heavily upon conscious images or words to guide our actions, either in the form of inner speech or through the instructions of others. While this type of control is flexible and computationally powerful, it is also highly effortful, enormously tiring, extremely restricted in scope and very error prone – and we don’t like it very much.

Although we all know what attention feels like, its precise function in mental life is not at all obvious. An optimum amount of attention is necessary for successful performance in all spheres of activity; but too little or too much can be highly disruptive. The consequences of inattention are clear enough, but if you need an example of over-attention, try using a keyboard while thinking about what the index finger of your right hand is doing

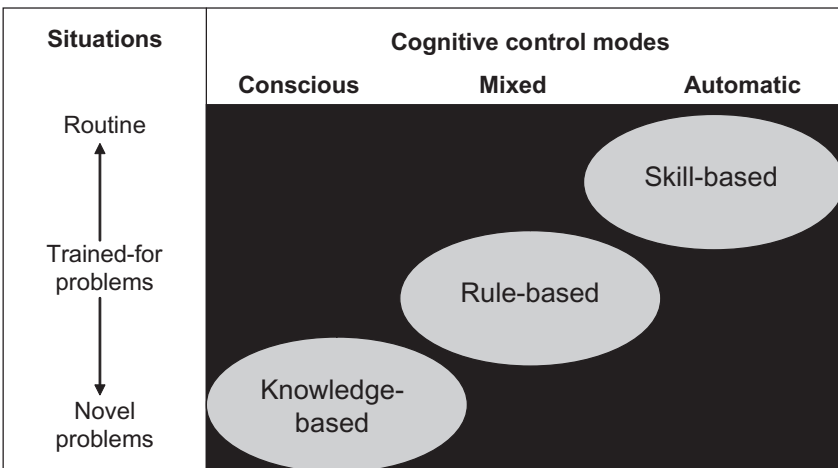


Figure 2.1 Three levels of performance control

– the greater your typing skills, the more likely it is that this will cause problems. As in many other areas of psychology, we can learn a good deal about the function of attention by observing its occasional failures.

At the other end of the spectrum there is skill-based performance. By dint of practice, self-discipline and the reshaping of our perceptions, we can gradually acquire the rudiments of a skill – that is, the ability to mix conscious goal-setting and guidance with the largely automatic control of our individual actions. Instead of agonising over each separate movement or word, we are able to run them off in pre-packaged sequences. This is the stuff of which habits are made, and it is the essence of skill-based performance. William James wrote: ‘Habit diminishes the conscious attention with which our acts are performed.’²

This gradual automatisisation of performance is universal; it occurs in all areas of mental performance. Even our social interactions become more and more automatic with time and experience. Without this ability to delegate control to non-conscious habit sequences, or motor programmes, we would consume all our limited attentional resources in dealing with the present moment and have nothing left to review the past or plan for the future. Such perpetual ‘present-mindedness’ would be insupportable; we could spend all day trying to tie our shoelaces. But nothing comes free. Automatisisation carries the penalty of occasional absent-mindedness when our actions do not go as planned.

Intermediate between the knowledge-based and skill-based levels lies rule-based performance. This arises when we need to break off from a sequence of largely habitual (skill-based) activity to deal with some kind of problem, or in which our behaviour needs to be modified to accommodate some change of circumstances. The commonest kinds of problems are those for which we have a pre-packaged solution, something that we have acquired through training, experience or some written procedure. We can express these solutions as ‘rules’: *If [problem X] then [apply solution Y]*, or *If [indications A and B are present] then [it is a type-C problem]*. These ‘rules’ are acquired as the result of experience and training and they are the stuff of which expertise is made.

2 James, W. (1889) *Principles of Psychology*. New York: Holt (p. 114).

However, as we shall see in the next chapter, this level of performance is associated with a variety of errors. We can misapply a normally good rule (i.e., one that usually works in this particular situation) because we have not noticed the contraindications; we can apply a bad rule; or we can fail to apply a normally good rule – a mistaken violation.

When we run out of pre-programmed problem solutions, as in some novel or unexpected situation, we are forced to resort to working out a solution 'on the hoof' using the slow, effortful, but computationally powerful conscious control mode. This is a highly error prone level of performance and is subject to a range of systematic biases. These knowledge-based mistakes will be considered in the next chapter.

Interacting with the Long-term Knowledge Base

The secretive properties of long-term memory and, in particular, the processes by which stored items are called to mind lie at the heart of the mind-user's misunderstandings of his or her mental function. There appear to be a number of different mechanisms involved in memory retrieval and two of them – similarity-matching and frequency-gambling – are automatic, unconscious and continuously operative. Since we cannot introspect upon these processes directly, we can only guess at their nature by observing the recurrent patterns shown by our errors, and by which of these processes, similarity or frequency, dominates in different kinds of memory search.

When the initial search cues are detailed or highly specific, matching these calling conditions to the characteristics of stored memory schemas on a like-to-like basis is the primary retrieval process. However, when the search cues match several stored knowledge structures, the mind gambles that the most frequently used knowledge item in that particular context will be the one that is required. Two examples will make the point clearer. If we were asked what is it that barks, has four legs, wags its tail, cocks its leg at lampposts and is regarded as man's best friend, most of us would quickly retrieve the knowledge item that matches all of these characteristics uniquely – namely a dog. Here the retrieval is based almost entirely on similarity matching and converges

upon a specific stored item. The process is so rapid that we feel we have reached out and retrieved the item in a conscious and deliberate fashion.

However, if we were asked to generate exemplars of the category 'four-legged animal' in no particular order, it is highly likely that, on average across a group of people, the first items coming to mind would be dog, cat, horse and cow. Once again, the item 'dog' is called to mind, but the search process was not similarity-matching. In this case, the order of retrieval is dominated by the familiarity of the animal. Familiarity is a function of frequency of encounter, so, in this divergent memory search, frequency-gambling is the primary search process.

Memory searches are strongly influenced by 'feelings of knowing'. We do not strive to retrieve things that we know we don't know. But we doggedly continue searching for something that we are sure we know, even though, as discussed earlier, we keep coming up with the wrong item. So, from the mind-user's point of view, these feelings about the contents of memory are of considerable value. They are not always right, of course, but they are correct often enough for us to treat them as handy guides to whether or not we should invest mental effort in a memory search. There are also things we don't always realise we know, but actually do – if only in a very approximate way.

A good example of this is frequency of encounter. The research evidence suggests that people automatically log how often they come across any reference to a particular person or topic. This logging process does not involve an actual numerical count; it is best captured by asking people to rate how often they have encountered some person or thing on a graded 0–6 scale ranging from 'never' (0) to 'nearly all the time' (6). By the same means, we can also obtain a fairly good assessment of co-occurrences. That is, we can gauge in very general terms how often X has occurred with Y with a moderate degree of accuracy. Understanding and exploiting these feelings of knowing is very useful, though not always appreciated by the mind user.

The human mind is exceptionally good at simplifying complex information-handling tasks. It does this by relying as far as possible upon the automatic mode of control and using intuitive 'rules of thumb' or heuristics. These are unconscious