

Coⁿtr^olling Uncertainty


**Decision Making and
Learning in Complex Worlds**

Magda Osman

Controlling Uncertainty

Decision Making and Learning in Complex Worlds

Magda Osman



 **WILEY-BLACKWELL**

A John Wiley & Sons, Ltd., Publication

This edition first published 2010

© 2010 Magda Osman

Wiley-Blackwell is an imprint of John Wiley & Sons, formed by the merger of Wiley's global Scientific, Technical, and Medical business with Blackwell Publishing.

Registered Office

John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, UK

Editorial Offices

The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, UK

9600 Garsington Road, Oxford, OX4 2DQ, UK

350 Main Street, Malden, MA 02148-5020, USA

For details of our global editorial offices, for customer services, and for information about how to apply for permission to reuse the copyright material in this book please see our website at www.wiley.com/wiley-blackwell.

The right of Magda Osman to be identified as the author of this work has been asserted in accordance with the UK Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, except as permitted by the UK Copyright, Designs and Patents Act 1988, without the prior permission of the publisher.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners. The publisher is not associated with any product or vendor mentioned in this book. This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Library of Congress Cataloging-in-Publication Data

Osman, Magda.

Controlling uncertainty : decision making and learning in complex worlds / Magda Osman.
p. cm.

Includes bibliographical references and index.

ISBN 978-1-4051-9946-9 (hardcover : alk. paper)

1. Uncertainty. 2. Decision making. 3. Cognitive psychology. 4. Control (Psychology) I. Title.

BF463.U5O86 2010

003'.5-dc22

2010011928

A catalogue record for this book is available from the British Library.

Set in 10.5 on 13.5 pt Palatino by Toppan Best-set Premedia Limited

Printed and bound in Malaysia by Vivar Printing Sdn Bhd

01 2010

Preface: the master puppeteer

The ringmaster of a travelling circus told a promising young puppeteer of a mysterious string puppet that would present a challenge to anyone who tried to command it. The puppeteer was intrigued.

‘Though it could be mastered, it would take all the skill and determination of even the most maven puppeteer.’ Then, in answer to the young puppeteer’s question, the ringmaster replied, ‘The reason why it is so difficult to work the puppet is that inside the puppet is a mechanism which follows its own rules and logic.’

This was enough to inspire the puppeteer. At last he had found his ambition. He would seek out and learn all there is to know of the puppet and how to rule it. But, before leaving, as a warning to the young apprentice, the ringmaster’s final words on the matter of the string puppet were these:

Be very careful. You can’t always be sure if the puppet acts the way you want it to because you’ve pulled the right strings or because of what is inside it. Sometimes you’d be forgiven for thinking that it acts as if it has a mind of its own.

When eventually the young puppeteer found the mysterious string puppet, what he saw was a large glass box supported by an elaborately carved wooden frame. The puppeteer paced around inspecting the box. On one of the sides of the box, there was a panel with many levers of different shapes and sizes, and beside that was a chair. The puppet inside the box looked fairly ordinary suspended in space, its limbs heavy, holding the posture that all string puppets have before they're brought to life. Though the puppeteer was told that this was a string puppet, having peered hard into the glass, he just couldn't find the strings. The strings must be either very fine or transparent to me, the puppeteer thought to himself.

The young puppeteer sat at the chair. He tried out various levers in a haphazard way at first, but not much seemed to happen. It was hard to tell which lever moved the head, and which the limbs. There were more moving parts to the puppet than there were levers. This brought moments of doubt into the puppeteer, making him wonder if he could ever hope to find the right means to master the puppet. However, when there was a purposeful order to operating the levers, the puppeteer could make the puppet's arm jolt, and even bring about the tapping of a foot.

Once in a while the words of the ringmaster surfaced in the mind of the puppeteer. As the ringmaster had warned, the puppet moved on its own, sometimes quite dramatically, but other times it would just be an almost imperceptible twitch. The puppeteer tried to ignore this, hoping that it would be enough to keep pulling and pushing at the same familiar levers. Though this did seem to work, it was only an illusory sense of command. The puppeteer knew that his rehearsed actions couldn't alleviate the scepticism that the puppet's behaviour brought about. 'I wonder whether the mechanism in the puppet truly is erratic, or is it simply that I haven't yet learnt all there is to know about it?'

Perseverance and will helped to keep the puppeteer's resolve to struggle through the doubting times. The belief that he was the one who could ultimately rule the puppet and not the other way around gave him hope. The puppeteer became wise to the fact that the mechanism that made the puppet move of its own accord wouldn't

always keep in time with the music that was played. On some occasions it seemed like the mechanism would go faster, and on other occasions it would seem to go much slower. There were even times when the limbs would move and dance in sequence, but not when anticipated.

It did take a long time, but the young puppeteer became skilled in conducting the mysterious string puppet. The ringmaster had followed the progress of the puppeteer, and found an opportunity to watch a performance. The tales were true: the puppet danced elegantly before the audience. The ringmaster hoped to know what valuable lessons the puppeteer had gained in their time of apprenticeship.

'All I know is this,' said the puppeteer. 'To make the puppet dance the way I wanted, I had to know how supple the limbs and head were. The suppler it seemed to be, the more elegantly I could make it dance. I had to know when the mechanism in the puppet was running fast or slow, and when it failed to behave as I'd expected from what I did. This helped me to decide the times when it simply seemed that it was indeed behaving of its own accord, from times that it was behaving as I wanted. For all this, I needed to keep in mind two very important details: that I was the one who could choose what levers to operate, and I was the one who could choose when to operate them. This is how I gathered what the puppet might do from one moment to the next, and this was what helped me to understand what more I needed to know to make it dance to different tunes.'

It was coming to know these things that helped to make the apprentice into the master puppeteer.

Acknowledgements

I begin my thanks to the Department of Cognitive, Perceptual and Brain Sciences University College London, where I spent the long summer months of 2009 working on this book. I would also like to acknowledge the support of the Centre for Economic Learning and Social Evolution (ELSE) University College London; the Psychology Department at Queen Mary University of London; the Centre for Vision, Speech and Signal Processing at University of Surrey; and the Engineering and Physical Sciences Research Council (EPSRC) grant EP/F069421/1.

There are many people to thank for their various contributions in helping me prepare for this book, as well as just helping me prepare myself in the first place. First and foremost, the individual who I would like to give thanks to is David Shanks. He has been a great mentor to me. His patient and methodical manner as well as his incisive style of thinking have been a great inspiration. Moreover, his continued support and endorsement have meant that I have been able to take advantage of one of the best intellectual environments in which to write this book. On this point, my thanks also go to Gabriella Vigliocco, Shelley Channon and John Draper, who have made the practicalities of working at UCL possible.

Of course, the intellectual environment I have been lucky to find myself in includes my friends. Dave Lagnado, Mark Johansen, Maarten Speekenbrink and Ramsey Rafaat have been most pivotal in introducing me to ideas that have been instrumental to the development of this book, as well as making suggestions that have helped structure the way the work has been presented. Their expertise and ideas have been invaluable. There are also many other lively discussions that I've had with a number of other good friends and colleagues that have helped shape the ideas that are meted out in this book. While you may not have realized, I hope this acknowledgement will now make this known. So, thanks to Ellen Seiss, Adam McNamara, Lorraine Attenbury, John Groeger, Shuichiro Taya, David Windridge, Bill Christmas and Josef Kittler. Also, thanks to Erica Yu, Andrea Smyth, Leonora Wilkinson, Marjan Jahanshahi, Nick Chater, Geoff Bird, Anna Cox and Paul Cairns. My thanks and appreciation also extend to those further afield: Axel Cleeremans, Stéphane Doyens, Wim De Neys, Bjoern Meder, Momme von Sydow, York Hagmayer, Franziska Bocklisch, Michael Waldmann, Ruth Stavy, Reuven Babai, Bob Hausmann and Brandon Weeks.

I would also like to thank those friends who have played various other important roles. They've encouraged me throughout this process, or helped to motivate me to embark on such an adventure in the first place, and others were there at the right place and right time to distract me when I needed it. So, thanks to Charlie Mounter, Belen Lopez, Kulwinder Gill, Christiane Waleczko-Fernandes, Susan Brocklesby, Madeline McSorley, Piotr Ptasinski, Mandy Heath, Sara Spence, Robert Coleman, Rikki Wong, Simon Li, Darrell Collins and Tracy Ray, and to Billy Nelson who has also helped coach me through some hard times in preparing this. Also, to a special friend who is no longer with us, Rüdiger Flach.

Along with my family members including little Layla, I reserve final thanks to two people who have gone through most of my trials and tribulations. They endured much of me, and this has not gone unnoticed. While I may have selfishly used their goodwill, and benefited from their exceptionally well-chosen, tailor-made

Acknowledgements

diversions (you know me far too well), I suspect that during this time I was less than entertaining in return. I hope that you can see this book as a return of favour. Thanks Gill Ward and Christopher Berry.

Contents

Preface: the master puppeteer	vii
Acknowledgements	x
1 Introduction	1
2 Causation and agency	16
3 Control systems engineering	52
4 Cybernetics, artificial intelligence and machine learning	80
5 Human factors (HCI, ergonomics and cognitive engineering)	115
6 Social psychology, organizational psychology and management	147
7 Cognitive psychology	179
8 Neuroscience	216

Contents

9	Synthesis	254
10	Epilogue: the master puppeteer	270
	References	273
	Index	318

Chapter 1

Introduction

There are two key characters in the story of control, namely, the master puppeteer and the puppet. Assume now that the master puppeteer is us, and our goal is to control the behaviour of the puppet. The puppet can represent any system that requires our control. For instance, the puppet could represent a system that is *biological*, like keeping our physical fitness levels up so that we can run a marathon. It could be *economic*, like a stock market in which we are maximizing our profit by buying and selling shares. It could be *organizational*, like marshalling a troop of soldiers to protect a safety zone. It could even be *ecological*, for instance trying to sustain an endangered ecosystem like a coral reef. More typically, when we think of control systems, what comes to mind is something *industrial* like operating a nuclear power plant, or *mechanical* like driving a car, or *safety critical* like flying a plane.

Clearly, then, there are many examples of control systems, some of which we are likely to experience on a regular basis in our daily lives. Control systems, therefore, are a rather broad subject and, as

will be made clear in this book, can encompass almost everything. So how are we able to do it; how are we able to exert control over these various types of systems? For those who have thought of this question, this book will be a guide to some answers.

Puppets, Puppets Everywhere ...

The reason that control systems can be seen everywhere is that almost anything can be thought of as a system, and most things we do in our lives are driven by our need to control. But, while the systems may be incredibly pervasive, they do adhere to some basic characteristics. Control systems are complicated. This is because they have a number of elements that will vary all at once and from one point in time to another – like puppets. Puppets can have multiple components; their moving parts (e.g., arms, legs, feet, hands fingers and head) are more often than not interconnected (e.g., fingers to hand, hand to arm etc.), and can vary all at once (e.g., performing a jump) as well as singularly in one point in time (e.g., raising a hand to wave). Thus, given these capacious characteristics of control systems, they are quite literally everywhere.

... It Takes All the Running You Can Do, to Keep in the Same Place¹

In these systems we are often required to manage the events that occur in a way that leads to something predictable, and desirable. This can be incredibly difficult to achieve and takes many years of training (e.g., becoming a pilot of a passenger jet), because the system varies of its own accord, as well as because acting upon it makes it change in some way. Or, more often than not, it is a combination of us acting on it and it doing something itself that produces changes in events. To bring the analogy of the puppet to

¹ Lewis Carroll, *Through the looking glass* (1871).

bear more obviously with control systems, in the story the puppet had its own internal mechanism that also made it move. Imagine how hard it is to control a malleable puppet and make it dance in time to a tune without an internal mechanism that can make it move on its own. Now imagine how much harder it is when it can move on its own and not always predictably. As hard as it seems, we are capable of achieving this. So we return to the question again: how are we able to exert control over such a complicated situation?

**Vladimir: 'Say Something!' Estragon: 'I'm Trying.
... In the Meantime Nothing Happens'. Pozzo:
'You Find It Tedious?' Estragon: 'Somewhat'.²**

To answer the question 'How?', we need to find a better way of asking it. First of all, finding some way of describing how these different types of systems work is of great importance, particularly if they can, on a basic level, be thought of in a similar way. Second, to complement this, our ability to control what happens in these systems should reduce to some basic psychological learning and decision-making mechanisms. They should do this because we need psychological mechanisms in place that enable us to predict the behaviour of the system and coordinate our own behaviours to effect a specific change in it. Therefore, finding some way of describing our psychological processes, along with describing the control system itself, is crucial to having an understanding of control (i.e., the scientific pursuit) and being able to improve our ability to manipulate our environment (i.e., the applied pursuit).

Given the extensiveness of both objectives, typically at the start of books like this there is a tendency to spell out at the beginning what things will not be included and what can't be achieved. I am going to avoid this. The aim of this book is to be as inclusive as possible. If you've flicked through it already, you will have

² Samuel Beckett, *Waiting for Godot* (1954/2009).

noticed that there are chapters spanning subject areas that include philosophy, engineering, cybernetics, human factors, social psychology, cognitive psychology and neuroscience. In order to get to the answer of 'How?', we need to consider the various contributions that each of these subjects has made. The issue of control invites attention from many disciplines that don't always speak to each other. Putting them side by side in chapters in a book is also a way of showing how they in fact do relate. Moreover, they also provide the groundwork for my answer to the question of 'How?' which is presented at the end of this book.

There are two important ideas that will help to carry you along this book: (1) all the themes introduced in this book are reducible to five basic concepts: control, prediction, cause-effect associations, uncertainty and agency; and (2) all of the issues that these basic concepts raise are ultimately, and will in this book be, directed towards addressing one question, which for the purposes of this book is *THE question: how do we learn about, and control online, an uncertain environment that may be changing as a consequence of our actions, or autonomously, or both?*

To understand the issue of control psychologically, and to understand the control system itself in all its various guises, we have to become familiar with these five core concepts and how they are tackled through the eyes of each of the aforementioned subjects. However, I am not alone; this endeavour has been embarked on by many,³ and throughout the different chapters of the book it will become apparent that there are various ways of understanding the psychological and objective characteristics of control systems. Therefore, I will take the opportunity here to qualify why this book is not a reinvention of the wheel, by stating what it hopes to do differently.

³ Cybernetics is the best-known example of an interdisciplinary movement designed to examine all issues related to control and self-regulation (see Chapter 4). More recently, machine-learning theorists have also attempted to draw work from engineering, biology, psychology and neuroscience to develop formal descriptions of behaviours associated with learning and controlling outcomes (e.g., Sutton & Barto, 1998).

The Aim of This Book

Role 1: catalogue

At its most humble, this book serves the purpose of being an inventory of sorts of what we currently know in a range of disciplines (e.g., engineering, artificial intelligence [AI], human factors, psychology and neuroscience) about control systems and control behaviour. Though not ever seriously taken up, an appeal of this kind was made in the late 1940s by Wiener, the self-proclaimed father of cybernetics – a discipline designed to study all matters related to self-organizing systems. Wiener (1948) hoped to bring together many disciplines to understand common problems concerning control. However, Wiener (1948) proposed that ‘the very speed of operations of modern digital machines stands in the way of our ability to perceive and think through the indicators of danger’ (p. 178). The effort in understanding all matters related to control came with a warning that technological advances may be such that artificial autonomous agents would be controlling our lives. That is, in the future the puppet would eventually rule the puppeteer, and not the other way around. Though the worry that control systems will reach a level of self-organization that may challenge our mastery of the world is perhaps unwarranted, surveying the most recent advances in theory and practice should give us a better understanding of what control systems can do, and our place with respect to them.

As suggested, an overhaul of this kind has yet to be undertaken, and so this book is an opportunity to do just that. For instance, due to the increasing complexity of the systems under our control (e.g., systems that identify tumours in X-ray images, voice recognition, predicting stock market trends, creating game play in computer games and profiling offenders), there are in turn ever increasing demands placed on them to achieve optimal performance reliably. Even something as prosaic as the car now includes an increased level of automation. This is generically classified under the title of *driver assist systems* (DAS). DAS now include electric power-assisted

steering (EPAS), semi-automatic parking (SAP), adaptive cruise control (ACC), lane departure warning (LDW) and vehicle stability control (VSC). All of these things now influence the ride and handling of vehicles we drive. So we might ask ourselves, if we have handed over so much autonomy to the car, what control do we have?

More to the point, disciplines such as control systems engineering present us with ever growing challenges because the control systems (e.g., car) that are part of our everyday interactions continue to increase in their capabilities and complexity. If complexity is increasing, then surely we need to know how we cope with it now, especially when things go wrong. Increasing complexity in our everyday lives doesn't just come from controlling devices such as cars. There has been a charted increase in the complexity of the decision making involved in economic, management and organizational domains (Willmott & Nelson, 2003). We can spot this complexity because some of it filters down to our consumer choices. For instance, take shopping. We have to adapt to the growing complexity that we face in terms of the information we have to process (e.g., more available product information), the choices we are presented with (e.g., more products to choose from) and the changing goals that we are influenced by (i.e., desires, aspirations and expectations). At the heart of adapting to the increasing level of complexity in our lives is our ability to still exert control. So, given the new challenges and demands that are placed on us in our lives right now, this book may be considered a sort of stock take of relevant and current research in the study of all things control related.

Role 2: solving the problem of complexity

A broader aim of the book is to help clarify what we mean when we say an environment is complex, and what it is about control systems that invites researchers from different disciplines to refer to them as complex. The complexity issue is important for the reason that there needs to be a cohesive idea about what makes